

Patch-based Progressive 3D Point Set Upsampling

Wang Yifan¹ Shihao Wu¹ Hui Huang^{2*}
Daniel Cohen-Or^{2,3} Olga Sorkine-Hornung¹

¹ETH Zurich ²Shenzhen University ³Tel Aviv University

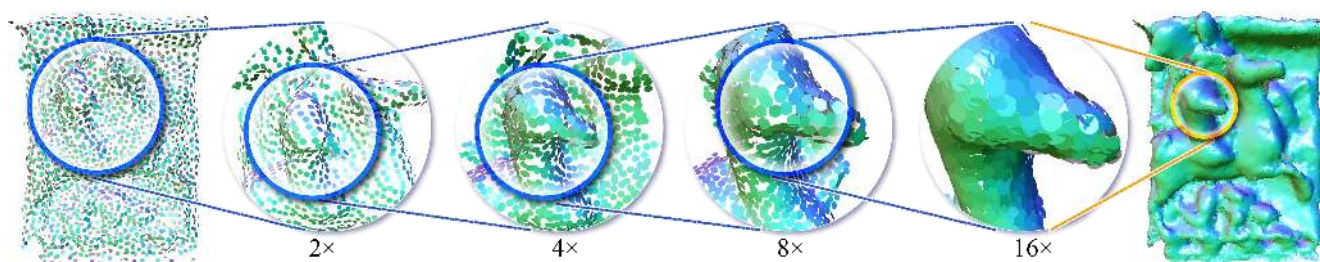


Figure 1: We develop a deep neural network for 3D point set upsampling. Intuitively, our network learns different levels of detail in multiple steps, where each step focuses on a local patch from the output of the previous step. By progressively training our patch-based network end-to-end, we successfully upsample a sparse set of input points, step by step, to a dense point set with rich geometric details. Here we use circle plates for points rendering, which are color-coded by point normals.

Abstract

We present a detail-driven deep neural network for point set upsampling. A high-resolution point set is essential for point-based rendering and surface reconstruction. Inspired by the recent success of neural image super-resolution techniques, we progressively train a cascade of patch-based upsampling networks on different levels of detail end-to-end. We propose a series of architectural design contributions that lead to a substantial performance boost. The effect of each technical contribution is demonstrated in an ablation study. Qualitative and quantitative experiments show that our method significantly outperforms the state-of-the-art learning-based [58, 59], and optimization-based [23] approaches, both in terms of handling low-resolution inputs and revealing high-fidelity details. The data and code are at <https://github.com/yifita/3pu>.

1. Introduction

The success of neural super-resolution techniques in image space encourages the development of upsampling methods for 3D point sets. A recent plethora of deep learning super-resolution techniques have achieved significant improvement in single image super-resolution performance [9, 27, 32, 47]; in particular, multi-step methods have been shown to excel in their performance [11, 30, 62]. Dealing

with 3D point sets, however, is challenging since, unlike images, the data is unstructured and irregular [3, 17, 19, 34, 55]. Moreover, point sets are often a result of customer-level scanning devices, and they are typically sparse, noisy and incomplete. Thus, upsampling techniques are particularly important, and yet the adaption of image-space techniques to point sets is far from straightforward.

Neural point processing is pioneered by PointNet [41] and PointNet++ [42], where the problem of irregularity and the lack of structure is addressed by applying *shared* multilayer perceptrons (MLPs) for the feature transformation of individual points, as well as a symmetric function, e.g., max pooling, for global feature extraction. Recently, Yu *et al.* [59] introduced the first end-to-end point set upsampling network, PU-Net, where both the input and the output are the 3D coordinates of a point set. PU-Net extracts multi-scale features based on PointNet++ [42] and concatenates them to obtain aggregated multi-scale features on each input point. These features are expanded by replication, then transformed to an upsampled point set that is located and uniformly distributed on the underlying surface. Although multiscale features are gathered, the level of detail available in the input patch is fixed, and thus both high-level and low-level geometric structures are ignored. The method consequently struggles with input points representing large-scale or fine-scale structures, as shown in Figures 11 and 12.

In this paper, we present a patch-based progressive upsampling network for point sets. The concept is illustrated in Figures 1 and 2. The multi-step upsampling breaks a, say, 16 \times -upsampling network, into four 2 \times subnets, where

* Corresponding author: Hui Huang (hhzhiyan@gmail.com)

each subnet focuses on a different level of detail. To avoid exponential growth in points and enable end-to-end training for large upsampling ratios and dense outputs, all subnets are fully patch-based, and the input patch size is adaptive with respect to the present level of detail. Last but not least, we propose a series of architectural improvements, including novel dense connections for point-wise feature extraction, code assignment for feature expansion, as well as bilateral feature interpolation for inter-level feature propagation. These improvements contribute to further performance boost and significantly improved parameter efficiency.

We show that our model is robust under noise and sparse inputs. It compares favorably against existing state-of-the-art methods in all quantitative measures and, most importantly, restores fine-grained geometric details.

2. Related work

Optimization-based approaches. Early optimization-based point set upsampling methods resort to shape priors. Alexa *et al.* [2] insert new points at the vertices of the Voronoi diagram, which is computed on the moving least squares (MLS) surface, assuming the underlying surface is smooth. Aiming to preserve sharp edges, Huang *et al.* [23] employ an anisotropic *locally optimal projection* (LOP) operator [22, 36] to consolidate and push points away from the edges, followed by a progressive edge-aware upsampling procedure. Wu *et al.* [53] fill points in large areas of missing data by jointly optimizing both the surface and the inner points, using the extracted meso-skeleton to guide the surface point set resampling. These methods rely on the fitting of local geometry, e.g., normal estimation, and struggle with multiscale structure preservation.

Deep learning approaches. PointNet [41], along its multiscale variant PointNet++ [41], is one of the most prominent point-based networks. It has been successfully applied in point set segmentation [10, 40], generation [1, 13, 56], consolidation [14, 45, 58], deformation [57], completion [15, 60] and upsampling [58, 59, 61]. Zhang *et al.* [61] extend a PointNet-based point generation model [1] to point set upsampling. Extensive experiments show its generalization to different categories of shapes. However, note that [1] is trained on the entire object, which limits its application to low-resolution input. PU-Net [59], on the other hand, operates on patch level, thus handles high-resolution input, but the upsampling results lack fine-grained geometry structures. Its follow-up work, the EC-Net [58], improves restoration of sharp features by minimizing a point-to-edge distance, but it requires a rather expensive edge annotation for training. In contrast, we propose a multi-step, patch-based architecture to channel the attention of the network to both global and local features. Our method also differs from the PU-Net and EC-Net in feature extraction, expansion, and loss computation, as discussed in

Section 3.2 and 4.

Multiscale skip connections in deep learning. Modern deep convolutional neural networks (CNN) [29] process multiscale information using skip-connections between different layers, e.g. U-Net [44], ResNet [16] and DenseNet [20]. In image super-resolution, state-of-the-art methods such as LapSRN [30] and ProSR [51] gain substantial improvement by carefully designing layer connections with progressive learning schemes [25, 50], which usually contribute to faster convergence and better preservation of all levels of detail. Intuitively, such multiscale skip-connections are useful for point-based deep learning as well. A few recent works have exploited the power of multiscale representation [12, 24, 28, 37, 49] and skip-connection [8, 43] in 3D learning. In this paper, we focus on point cloud upsampling and propose intra-level and inter-level point-based skip-connections.

3. Method

Given an unordered set of 3D points, our network generates a denser point set that lies on the underlying surface. This problem is particularly challenging when the point set is relatively sparse, or when the underlying surface has complex geometric and topological structures. In this paper, we propose an end-to-end progressive learning technique for point set upsampling. Intuitively, we train a multi-step patch-based network to learn the information from different levels of detail. As shown in Figures 2 and 3, our model consists of a sequence of upsampling network units. Each unit has the same structure, but we employ it on different levels of detail. The information of all levels is shared via our intra-level and inter-level connections inside and between the units. By progressively training all network units end-to-end, we achieve significant improvements over previous works. We first present the global design of our network and then elaborate on the upsampling units.

3.1. Multi-step upsampling network

Multi-step supervision is common practice in neural image super-resolution [11, 30, 62]. In this section, we first discuss the difficulties in adapting multi-step learning to point set upsampling, which motivates the design of our multi-step *patch-based* supervision method. Next, we illustrate the end-to-end training procedure for a cascade of upsampling network units for large upsampling ratios and high-resolution outputs.

Multi-step patch-based receptive field. Ideally, a point set upsampling network should span the receptive field adaptively for various scales of details to learn geometric information from multiple scales. However, it is challenging to apply a multi-scope receptive field on a dense irregular point set due to practical constraints. In contrast to im-

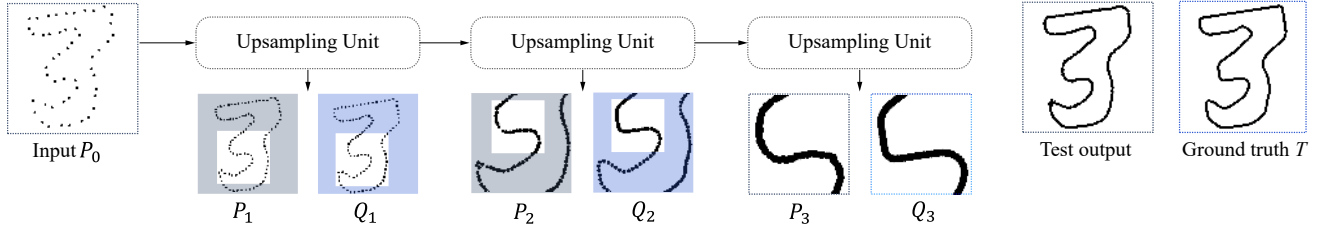


Figure 2: Overview of our multi-step patch-based point set upsampling network with 3 levels of detail. Given a sparse point set as input, our network predicts a high-resolution set of points that agree with the ground truth. Instead of training an $8\times$ -upsampling network, we break it into three $2\times$ steps. In each training step, our network randomly selects a local patch as input, upsamples the patch under the guidance of ground truth, and passes the prediction to the next step. During testing, we upsample multiple patches in each step independently, then merge the upsampled results to the next step.

ages, point sets do not have the regular structure, and the neighborhoods of points are not fixed sets. Neighborhood information must be *collected* by, e.g., k -nearest neighbors (k NN) search. This per-layer and per-point computation is rather expensive, prohibiting a naive implementation of a multi-step upsampling network to reach large upsampling ratios and dense outputs. Therefore, it is necessary to optimize the network architecture, such that it is scalable to a high-resolution point set.

Our key idea is to use a multi-step patch-based network, and the patch size should be adaptive to the scope of receptive fields at the present step. Note that in neural point processing, the scope of a receptive field is usually defined by the k NN size used in the feature extraction layers. Hence, if the neighborhood size is fixed, the receptive field becomes narrower as the point set grows denser. This observation suggests that it is unnecessary for a network to process all the points when the receptive field is relatively narrow. As shown in Figure 2, our network *recursively* upsamples a point set while at the same time reduces its spatial span. This multi-step patch-based supervision technique allows for a significant upsampling ratio.

Multi-step end-to-end training. Our network takes L steps to upsample a set of points by a factor of 2^L . For L levels of detail, we train a set of subnet units $\{U_1, U_2, \dots, U_L\}$. We train such a sequence of upsampling units by progressively activating the training of units; it has been used in many multiscale neural image processing works [25, 51].

More specifically, our entire training process has $2L - 1$ stages, i.e., every upsampling unit has two stages except the first one. We denote the currently targeted level of detail by \hat{L} . In the first stage of $U_{\hat{L}}$ we fix the network parameters of units U_1 to $U_{\hat{L}-1}$ and start the training of unit $U_{\hat{L}}$. In the second stage, we unleash the fixed units and train all the units simultaneously. This progressive training method is helpful because an immature unit can impose destructive gradient turbulence on the previous units [25].

We denote the ground truth model, prediction patch and reference patch with T , P and Q respectively and use \hat{L} and

ℓ to denote the targeted level of detail and an intermediate level, as illustrated in Figure 2 and 6. In practice, we recursively shrink the spatial scope by confining the input patch to a fixed number of points (N). For more technical detail about extracting such input patches on-the-fly and updating the reference patches accurately, please refer to Section 3.3.

3.2. Upsampling network unit

Let us now take a closer look at an upsampling network unit U_ℓ . It takes a patch from $P_{\ell-1}$ as input, extracts deep feature, expands the number of features, compresses the feature channels to d -dimensional coordinates P_ℓ . In the following, we explain each component in greater detail.

Feature extraction via intra-level dense connections.

We strive for extracting structure-aware features ($N \times C$) from an input point set ($N \times d$). In neural image processing, skip-connection is a powerful tool to leverage features extracted across different layers of the network [16, 20, 21, 35]. Following PointNet++ [42], most existing point-based networks extract multiple scales of information by hierarchically downsampling the input point sets [33, 59]. Skip-connections have been used to combine multiple levels of features. However, a costly point correspondence search must be applied prior to skip-connections, due to the varying point locations caused by the downsampling step.

We propose an architecture that facilitates *efficient* dense connections on point sets. Inspired by the dynamic graph convolution [46, 52], we define our local neighborhood in feature space. The point features are extracted from a local neighborhood that is computed dynamically via k NN search based on feature similarity. As a result, our network obtains long-range and nonlocal information without point set subsampling.

As shown in Figure 5, our feature extraction unit is composed of a sequence of dense blocks. In each dense block, we convert the input to a fixed number (C') of features, group the features using feature-based KNN, refine each grouped feature via a chain of densely connected MLPs, and finally compute an order-invariant point feature via max-pooling.

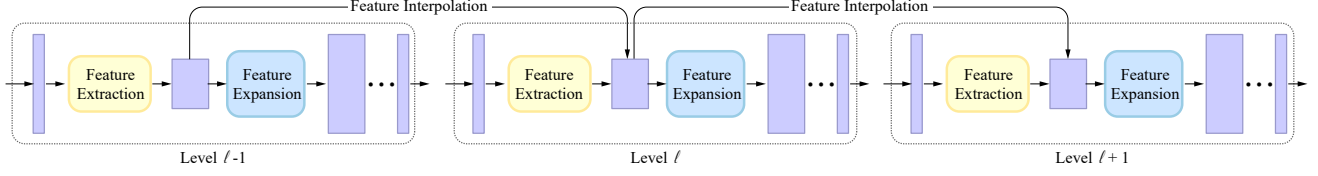


Figure 3: Illustration of three upsampling network units. Each unit has the same structure but applied on different levels.

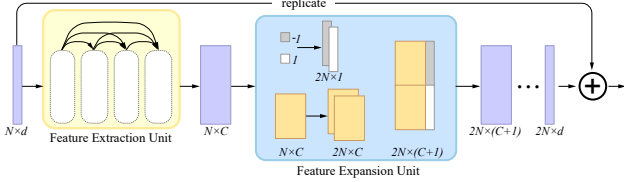


Figure 4: Illustration of one upsampling network unit.

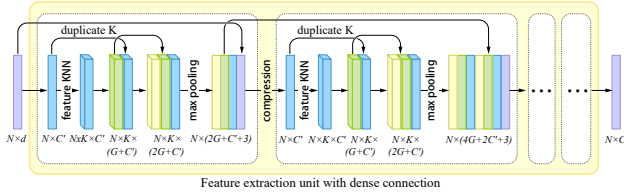


Figure 5: Illustration of the feature extraction unit with dense connections.

We introduce dense connections both within and between the dense blocks. Within the dense blocks, each MLP’s output, i.e., a fixed number (G) of features, is passed to *all* subsequent MLPs; between the blocks, the point features produced by each block are fed as input to *all* following blocks. All these skip-connections enable explicit information re-use, which improves the reconstruction accuracy while significantly reducing the model size, as demonstrated in Section 4. Overall, our $16\times$ -upsampling network with four $2\times$ -upsampling units has much fewer network parameters than a $4\times$ -upsampling PU-Net [59]: 304K vs. 825K.

Feature expansion via code assignment. In the feature expansion unit, we aim to transform the extracted features ($N \times C$) to an upsampled set of coordinates ($2N \times d$).

PU-Net [59] replicates the per-point features and then processes each replicant independently by an individual set of MLPs. This approach may lead to clustered points around the original points positions, which is alleviated by introducing a repulsion loss. Instead of training the network to disentangle the replicated features in-place, we explicitly offer the network the information about the position variation.

In conditional image generation models [39], a category-variable is usually concatenated to a latent code to generate images of different categories. Similarly, we assign a 1D code, with value -1 and 1 , to each of those duplicated features to transform them to different locations, as shown

in Figure 4. Next, we use a set of MLPs to compress the $2N \times (C + 1)$ features to $2N \times d$ residuals, which we add to the input coordinates to generate the output points.

Our experiments show that the proposed feature expansion method results in a well distributed point set without using an additional loss. Also, the number of network parameters is independent of the upsampling ratio, since all expanded features share the consecutive MLPs.

Our feature expansion method is also related to recent point cloud generative models FoldingNet [56] and AtlasNet [13], where the coordinates of a 2D point are attached to the learned features for point generation. Here, we show that the choice of an attached variable can be as simple as a 1D variable.

Inter-level skip connection via bilateral feature interpolation. We introduce inter-level skip-connections to enhance the communication between the upsampling units, which serves as bridges for features extracted with different scopes of the receptive fields, as shown in Figure 3.

To pass features from previous levels the current level, the key is a feature interpolation technique that constructs corresponding features from the previous upsampling unit, as the upsampling and patch extraction operations change the point correspondence. Specifically, we use bilateral interpolation. For the current level ℓ , we denote by p_i and f_i the coordinates of the i -th point and its features generated by the feature extraction unit respectively, and \mathcal{N}'_i denotes the spatial k -NN of p_i from level ℓ' . the interpolated feature for f_i can be written as:

$$\tilde{f}_i = \frac{\sum_{i' \in \mathcal{N}'_i} \theta(p_i, p_{i'}) \psi(f_i, f_{i'}) f_{i'}}{\sum_{i' \in \mathcal{N}'_i} \theta(p_i, p_{i'}) \psi(f_i, f_{i'})}, \quad (1)$$

with the joint weighting functions: $\theta(p_1, p_2) = e^{-\left(\frac{\|p_1 - p_2\|}{r}\right)^2}$, $\psi(f_1, f_2) = e^{-\left(\frac{\|f_1 - f_2\|}{h}\right)^2}$. The width parameters r and h are computed using average distance to the closest neighbor.

One way to implement the inter-level connection is to interpolate and concatenate \tilde{f}_i from *all* previous layers, i.e., use dense links the same as those within the feature extraction units. However, doing so would result in a very wide network, with ℓC features in level ℓ (typically $C = 216$), causing scalability issues and optimization difficulties [51]. Instead, we apply residual skip-connections, i.e.,

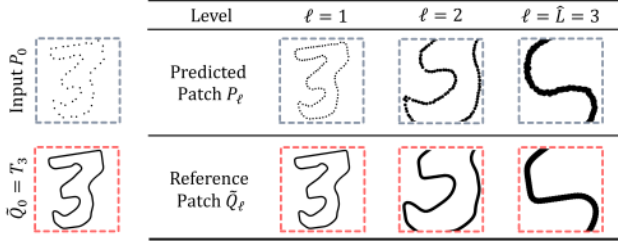


Figure 6: Extraction of patches for $\hat{L} = 3$ during training. In this example, since there are only a small number of input points in 2D data, the first level contains the whole input shape ($N = |P_0|$).

$f_i = \tilde{f}_i + f_i$. By applying such residual links per-level, contextual information from coarser scales can be propagated through the entire network and incorporated for the restoration of finer structures. We learn through experiments that both dense links and residual links contribute positively to the upsampling result, but the latter has better performance in terms of memory efficiency, training stability and reconstruction accuracy.

3.3. Implementation details

Iterative patch extraction. In each training step, the target resolution \hat{L} is fixed. $P_{\hat{L}}$ and $Q_{\hat{L}}$ denote the prediction and reference patch in \hat{L} , whereas $T_{\hat{L}}$ denotes the *entire* reference shape in this resolution. We compute $P_{\hat{L}}$ and $Q_{\hat{L}}$ iteratively from a series of intermediate predictions and references, denoted as P_ℓ and \tilde{Q}_ℓ where $\ell = 1 \dots \hat{L} - 1$.

More specifically, the input to level ℓ is obtained using k NN ($k = N$) around a random point $p_{\ell-1}^*$ in $P_{\ell-1}$. \tilde{Q}_ℓ should match the spatial extent of P_ℓ but has a higher resolution, hence it can be extracted by k NN search in $\tilde{Q}_{\ell-1}$ using the same query point $p_{\ell-1}^*$, whereas $k = 2^{\hat{L}-\ell+1}N$. Note that we normalize the patches to a unit cube to improve the computational stability. In Figure 6 we illustrate the described procedure for $\hat{L} = 3$.

For inference, the procedure differs from above in two points: 1. In each level, we extract H overlapping input patches to ensure coverage of the entire input point set, the query points are sampled with farthest sampling; 2. We obtain P_ℓ by first merging the H overlapping partial outputs and then resampling with farthest sampling such that $|P_\ell| = 2|P_{\ell-1}|$. The resampling leads to uniform point distribution despite overlapping regions.

Using a small N could theoretically restrict the contextual information, while a larger N could unnecessarily increase the input complexity thus training difficulty. In our experiments, the choice of the input patch size N is not that critical for the upsampling quality.

Loss function. We use Euclidean distance for patch extraction for its speed and flexibility. This implies that the patch pairs P_ℓ and Q_ℓ might have misalignment problems on their borders. We observe that the loss computed on

those unmatched points adds noise and outliers in the result. Thus, we propose a modified Chamfer distance:

$$\mathcal{L}(P, Q) = \frac{1}{|P|} \sum_{p \in P} \xi \left(\min_{q \in Q} \|p - q\|^2 \right) + \frac{1}{|Q|} \sum_{q \in Q} \xi \left(\min_{p \in P} \|p - q\|^2 \right), \quad (2)$$

where the function ξ filters outliers above a threshold δ :

$$\xi(d) = \begin{cases} d, & d \leq \delta \\ 0, & \text{otherwise} \end{cases}. \text{ We set } \delta \text{ to be a multiple of the average nearest neighbor distance so as to dynamically adjust to patches of different scales.}$$

4. Results

In this section, we compare our method quantitatively and qualitatively with state-of-the-art point upsampling methods, and evaluate various aspects of our model. Please refer to the supplementary for further implementation details and extended experiments.

The metrics used for evaluation are (i) Chamfer distance, (ii) Hausdorff distance [4] and (iii) point-to-surface distance computed against the ground truth mesh.

Training and testing data. We generate two datasets for our experiments: MNIST-CP, Sketchfab and ModelNet10 [54]. MNIST-CP consists of 50K and 10K training and testing examples of 2D *contour points* extracted from the MNIST dataset [31]. Given a set of 2D pixel points, we apply Delaunay triangulation [5], Loop surface subdivision [38], boundary edge extraction, and WLOP [22] to generate a uniformly distributed point set lying on the contour curve of the image. The number of points in input P and ground truth point sets T_1, T_2 and T_3 are 50, 100, 200 and 800, respectively. Sketchfab consists of 90 and 13 highly detailed 3D models downloaded from SketchFab [48] for training and testing, respectively. ModelNet10 is comprised of 10 categories, containing 3991 and 908 CAD meshes for training and testing, respectively. We use the Poisson-disk sampling [7] implemented in Meshlab [6] to sample input and ground truth point sets with the number of points ranging from 625 to 80000. Our data augmentation includes random rotation, scaling and point perturbation with gaussian noise.

Comparison. We compare our method on relatively sparse (625 points) and dense (5000 points) inputs with three state-of-the-art point set upsampling methods: EAR [23], PU-Net [59] and EC-Net [58]. The code of these methods is publicly available. For EAR, we set the parameter $\sigma_n = 35^\circ$ to favor sharp feature preservation. For PU-Net and EC-Net, we obtain $16\times$ results by iteratively applying their $4\times$ -upsampling model twice, as advised by the authors. As for comparison, we train a four-step $16\times$ model using our method, where the initial patch size falls into a similar level of detail as PU-Net. For all experiments, we add to the input Gaussian noise with 0.25% magnitude of the model dimensions.

Method	Sparse input			Dense input			# Param.
	CD	HD	P2F	CD	HD	P2F	
EAR	0.67	7.75	5.25	0.09	1.82	1.88	-
PU	0.72	9.24	6.82	0.41	5.45	3.39	814K
EC	0.91	13.4	6.42	0.24	4.21	2.64	823K
Ours	0.54	6.92	3.32	0.06	1.31	1.11	304K

Table 1: Quantitative comparison with state-of-the-art approaches for $16\times$ upsampling from 625 and 5000 input points tested on Sketchfab dataset.

	10^{-3}	bathub	bed	chair	desk	dresser	monitor	n. stand	sofa	table	toilet
CD	PU	1.01	1.12	0.82	1.22	1.55	1.19	1.77	1.13	0.69	1.39
	EC	1.43	1.81	1.8	1.30	1.43	2.04	1.88	1.79	1.00	1.72
	ours	0.70	0.77	0.90	0.96	1.13	0.83	1.37	0.67	0.58	1.02
HD	PU	10.77	12.39	10.38	13.29	14.08	14.01	16.21	11.66	9.7	14.74
	EC	15.71	23.17	18.65	16.12	16.37	30.48	20.29	19.97	12.42	18.58
	ours	7.76	9.36	9.70	9.19	11.33	9.90	13.52	8.37	5.87	10.95

Table 2: Quantitative comparison with state-of-the-art approaches on ModelNet10 dataset for $16\times$ upsampling from 625 input points.

Table 1 and 2 summarizes the quantitative comparison conducted using Sketchfab and ModelNet10. Note that because many models in ModelNet10 are not watertight, we omit the point-to-surface distance in Table 2. Examples of the upsampling results are provided in Figures 11 and 12 for visual comparison, where we apply surface reconstruction to the upsampled point sets using PCA normal estimation (neighborhood number = 25) [18] and screened Poisson reconstruction (depth = 9) [26]. As seen in Figures 11 and 12, EAR generates competitive results for denser inputs but struggles with sparse inputs. As shown in Table 1, the performance of PU-Net on sparse and dense inputs is similar, revealing its limitation for high levels of detail. For denser inputs, EC-Net produces clean and more well defined outputs than PU-Net, but also shows signs of over-sharpening. For sparse input though, EC-Net produces more artifacts, possibly because the geodesic KNN, which EC-Net is built upon, becomes unreliable under sparse inputs. In comparison, our method outperforms all these methods quantitatively by a large margin. Qualitatively, our results are less noisy and contain notably more details.

Ablation study. An ablation study quantitatively evaluates the contribution of each of our proposed components:

1. Multi-stage architecture: we train a $2\times$ -upsampling model for all levels of detail and test by iteratively applying the model 4 times.
2. End-to-end training: we train each upsampling unit separately.
3. Progressive training: instead of progressively activating the training of each upsampling unit as described in Section 3.1, we train all units simultaneously.
- 4-6. Dense feature extraction, expansion, and inter-level skip-connections: we either remove or replace each of these modules with their counterpart in PU-Net.

As Table 3 shows, all components contributes positively

Removed/Replaced component	CD 10^{-3}	HD 10^{-3}	P2F 10^{-3}	Param.
1. Multi-stage architecture	0.69	9.98	4.07	65K
2. End-to-end training	0.73	9.91	3.34	263K
3. Progressive end-to-end training	0.55	7.46	3.49	304K
4. Dense feature extraction	0.61	9.17	4.17	2855K
5. Feature expansion	0.73	9.83	5.30	1642K
6. Inter-level skip-connections	0.61	7.65	3.38	263K
Our full model	0.54	6.92	3.32	304K

Table 3: Ablation study with $16\times$ -upsampling factor tested on the Sketchfab dataset using 625 points as input. We evaluate the contribution of each proposed component quantitatively with Chamfer distance (CD), Hausdorff distance (HD) and mean point-to-surface distance (P2F), and also report the number of parameters in the rightmost column.




Figure 7: Study of patch-based progressive upsampling. From left to right: input with 50 points, (i) direct $16\times$ upsampling, (ii) iterative $2\times$ upsampling trained with augmented data, (iii) multi-stage network trained separately, (iv) progressively train all networks but omit the per-stage patch extraction technique introduced in Section 3.1, and finally (v) progressively train all networks with patch extraction.

to the full model. In particular, removing multi-stage architecture significantly increased the difficulty of the task, resulting in artifacts shown in Figure 8b. We observe similar artifacts when the upsampling units are trained separately (Figure 8c), as the networks cannot counteract the mistakes made in previous stages. The proposed dense feature extraction, feature expansion, and inter-level skip-connections considerably improve the upsampling results. Moreover, the feature extraction and expansion unit contribute to significant parameter reduction.

Study of patch-based progressive upsampling. We evaluate the effect of our core idea, patch-based progressive upsampling, in greater detail. For this purpose, we start from the architecture proposed by PU-Net and add the techniques introduced in Section 3.1 one by one. Specifically, we conduct the following experiments on MNIST-CP dataset: (i) train a PU-Net with direct $16\times$ upsampling, (ii) train one $2\times$ PU-Net using training examples sampled with all available patch densities and then apply it iteratively 4 times, (iii) train a network for each level of detail separately, (iv) progressively train all networks but omit the per-stage patch extraction technique introduced in Section 3.1, and finally (v) progressively train all networks with patch extraction.

The results are shown in Figure 7. Both direct upsampling and single-stage model ((i) and (ii)) are unable to reconstruct faithful geometry in curvy regions, suggesting that a multi-stage architecture is necessary for capturing high levels of detail. The multi-stage PU-Net (iii) notably




Figure 8: Visual comparison for ablation study. We perform $16\times$ -upsampling from 625 points (left). (a)-(d) show a point patch of the input and the results from the single-stage model, separately trained model and our full model.

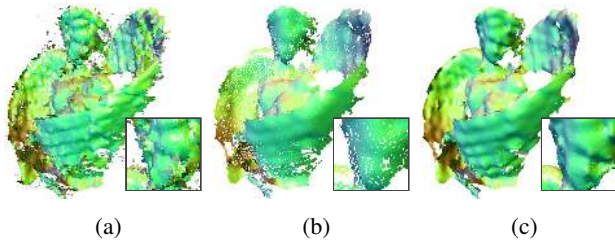


Figure 9: $16\times$ upsampling results using a real scan as input. Given a noisy input (a), we use WLOP [22] to obtain a consolidated point set (b), to which we apply our upsampling network (c).

improves the result but shows more artifacts compared with an end-to-end multi-stage model (iv), since the network has a chance to correct the mistakes introduced in earlier stages. Finally, applying adaptive patch extraction (v) further refines the local geometry, indicating that it helps the network to focus on local details by adapting the spatial span of input to the scope of receptive fields.

Stress test. To test the robustness to noise and sparsity, we subject an input point set to different noise levels ranging from 0% to 2%, and for sparsity we randomly remove 10% to 50% of the points from the input. The corresponding results from MNIST-CP datasets are shown in Figures 10a and 10b. Compared to PU-Net, our model is more robust against noise and sparsity.

Real world data. To test our model on real scans, we acquire input data using a hand-held 3D scanner Intel RealSense SR300. Albeit dense, such data is severely ridden with noise and outliers. Therefore, we first employ WLOP [22], a point set denoising tool known to be robust against noise and outliers, to consolidate and simplify the point set. We then apply our model to the resulting, denoised yet sparse point set and obtain a dense and clean output, as shown in Figure 9c.

5. Conclusion

In this work, we propose a progressive point set upsampling network that reveals detailed geometric structures from sparse and noisy inputs. We train our network step by step, where each step specializes in a certain level of detail. In particular, we direct the attention of our network to local

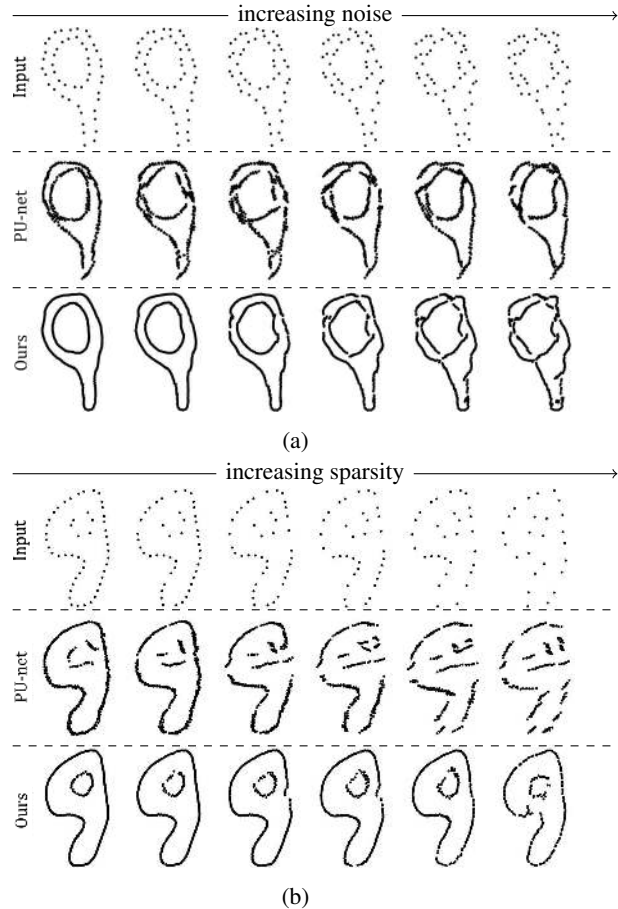


Figure 10: Stress test with increasing noise (a) and sparsity (b). The model is trained using 50 input points and Gaussian noise of 0.25% magnitude of the point set dimensions. In (a) we test with noise level of 0, 0.25%, 0.5%, 1%, 1.5% and 2%; in (b) we test with 50, 45, 40, 35, 30, and 25 input points.

geometric details by reducing the spatial span as the scope of the receptive field shrinks. Such adaptive patch-based architecture enables us to train on high-resolution point sets in an end-to-end fashion. Furthermore, we introduce dense connections for feature extraction, code assignment for efficient feature expansion, as well as bilateral feature interpolation for interlinks across the steps. Extensive experiments and studies demonstrate the superiority of our method compared with the state-of-the-art techniques.

Acknowledgement

We thank the anonymous reviewers for their constructive comments and the SketchFab community for sharing their 3D models. This work was supported in parts by SNF grant 200021_162958, ISF grant 2366/16, NSFC (61761146002), LHTD (20170003), and the National Engineering Laboratory for Big Data System Computing Technology.

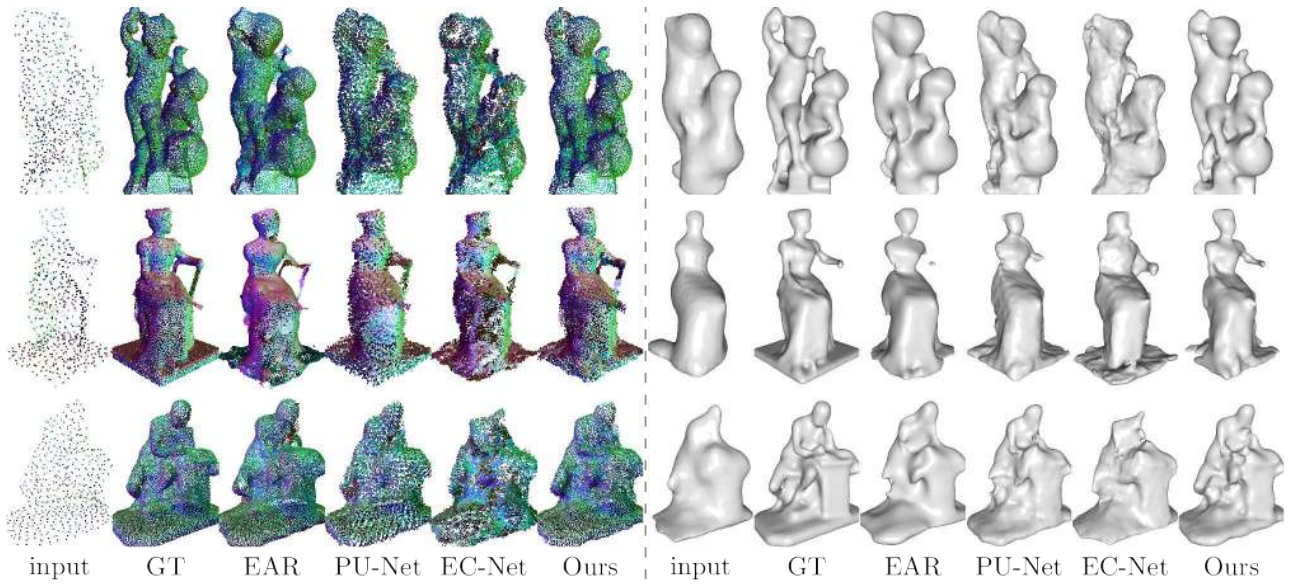


Figure 11: 16x upsampling results from 625 input points (left) and reconstructed mesh (right).

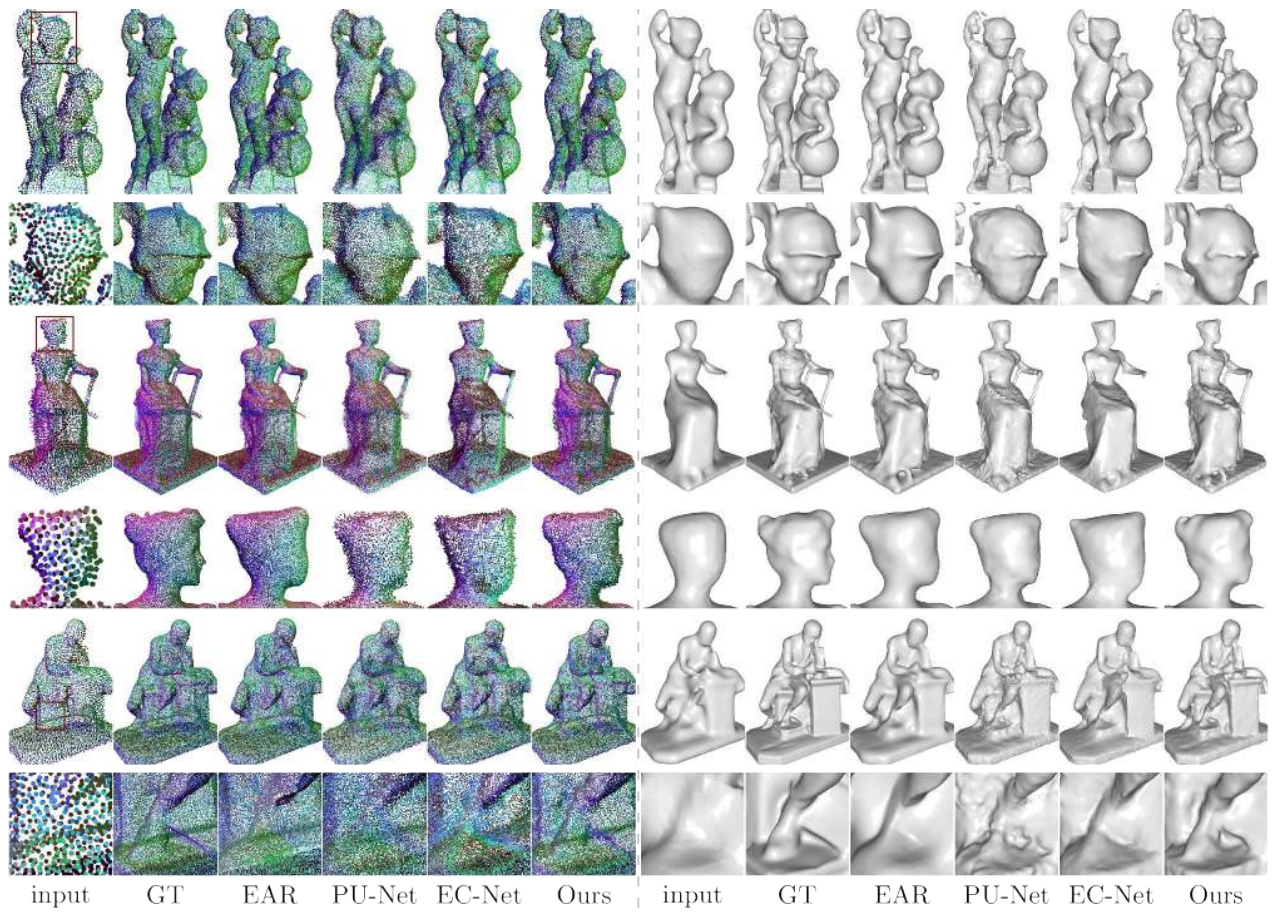


Figure 12: 16x upsampling results from 5000 input points (left) and reconstructed mesh (right).

References

- [1] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas. Learning representations and generative models for 3D point clouds. *Proc. Int. Conf. on Machine Learning*, 2018. 2
- [2] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. T. Silva. Computing and rendering point set surfaces. *IEEE Trans. Visualization & Computer Graphics*, 9(1):3–15, 2003. 2
- [3] M. Atzmon, H. Maron, and Y. Lipman. Point convolutional neural networks by extension operators. *ACM Trans. on Graphics (Proc. of SIGGRAPH)*, 2018. 1
- [4] M. Berger, J. A. Levine, L. G. Nonato, G. Taubin, and C. T. Silva. A benchmark for surface reconstruction. *ACM Trans. on Graphics*, 32(2):20, 2013. 5
- [5] J.-D. Boissonnat, O. Devillers, S. Pion, M. Teillaud, and M. Yvinec. Triangulations in CGAL. *Computational Geometry*, 22:5–19, 2002. 5
- [6] 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. 5
- [7] M. Corsini, P. Cignoni, and R. Scopigno. Efficient and flexible sampling with blue noise properties of triangular meshes. *IEEE Trans. Visualization & Computer Graphics*, 18(6):914–924, 2012. 5
- [8] H. Deng, T. Birdal, and S. Ilic. PPF-FoldNet: Unsupervised learning of rotation invariant 3D local descriptors. *arXiv preprint arXiv:1808.10322*, 2018. 2
- [9] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Analysis & Machine Intelligence*, 38(2):295–307, 2016. 1
- [10] F. Engelmann, T. Kontogianni, J. Schult, and B. Leibe. Know what your neighbors do: 3D semantic segmentation of point clouds. *arXiv preprint arXiv:1810.01151*, 2018. 2
- [11] Y. Fan, H. Shi, J. Yu, D. Liu, W. Han, H. Yu, Z. Wang, X. Wang, and T. S. Huang. Balanced two-stage residual networks for image super-resolution. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition Workshops*, pages 1157–1164. IEEE, 2017. 1, 2
- [12] M. Gadelha, R. Wang, and S. Maji. Multiresolution tree networks for 3D point cloud processing. *arXiv preprint arXiv:1807.03520*, 2018. 2
- [13] T. Groueix, M. Fisher, V. G. Kim, B. Russell, and M. Aubry. AtlasNet: A papier-mâché approach to learning 3D surface generation. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, 2018. 2, 4
- [14] P. Guerrero, Y. Kleiman, M. Ovsjanikov, and N. J. Mitra. PCPNet learning local shape properties from raw point clouds. *Computer Graphics Forum*, 37(2):75–85, 2018. 2
- [15] S. Gurumurthy and S. Agrawal. High fidelity semantic shape completion for point clouds using latent optimization. *arXiv preprint arXiv:1807.03407*, 2018. 2
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, pages 770–778, 2016. 2, 3
- [17] P. Hermosilla, T. Ritschel, P.-P. Vazquez, A. Vinacua, and T. Ropinski. Monte carlo convolution for learning on non-uniformly sampled point clouds. *ACM Trans. on Graphics (Proc. of SIGGRAPH Asia)*, 37(6), 2018. 1
- [18] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Proc. of SIGGRAPH*, pages 71–78, 1992. 6
- [19] B.-S. Hua, M.-K. Tran, and S.-K. Yeung. Pointwise convolutional neural networks. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, pages 984–993, 2018. 1
- [20] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, 2017. 2, 3
- [21] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Q. Weinberger. Deep networks with stochastic depth. In *Proc. Euro. Conf. on Computer Vision*, pages 646–661. Springer, 2016. 3
- [22] H. Huang, D. Li, H. Zhang, U. Ascher, and D. Cohen-Or. Consolidation of unorganized point clouds for surface reconstruction. *ACM Trans. on Graphics (Proc. of SIGGRAPH Asia)*, 28(5):176:1–176:7, 2009. 2, 5, 7
- [23] H. Huang, S. Wu, M. Gong, D. Cohen-Or, U. Ascher, and H. Zhang. Edge-aware point set resampling. *ACM Trans. on Graphics*, 32(1):9:1–9:12, 2013. 1, 2, 5
- [24] M. Jiang, Y. Wu, and C. Lu. PointSIFT: A SIFT-like network module for 3D point cloud semantic segmentation. *arXiv preprint arXiv:1807.00652*, 2018. 2
- [25] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *Proc. Int. Conf. on Learning Representations*, 2018. 2, 3
- [26] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Trans. on Graphics*, 32(1):29:1–29:13, 2013. 6
- [27] J. Kim, J. Kwon Lee, and K. Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, pages 1646–1654, 2016. 1
- [28] R. Klokov and V. Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3D point cloud models. In *Proc. Int. Conf. on Computer Vision*, pages 863–872. IEEE, 2017. 2
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *In Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105, 2012. 2
- [30] W.-S. Lai, J.-B. Huang, N. Ahuja, and M.-H. Yang. Deep laplacian pyramid networks for fast and accurate superresolution. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, 2017. 1, 2
- [31] Y. LeCun and C. Cortes. MNIST handwritten digit database. <http://yann.lecun.com/exdb/mnist/>, 2010. 5
- [32] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. P. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, 2017. 1
- [33] J. Li, B. M. Chen, and G. H. Lee. So-net: Self-organizing network for point cloud analysis. In *Proc. IEEE Conf. on*

- Computer Vision & Pattern Recognition*, pages 9397–9406, 2018. 3
- [34] Y. Li, R. Bu, M. Sun, and B. Chen. Pointcnn. *arXiv preprint arXiv:1801.07791*, 2018. 1
- [35] D. Lin, Y. Ji, D. Lischinski, D. Cohen-Or, and H. Huang. Multi-scale context intertwining for semantic segmentation. In *Proc. Euro. Conf. on Computer Vision*, pages 603–619, 2018. 3
- [36] Y. Lipman, D. Cohen-Or, D. Levin, and H. Tal-Ezer. Parameterization-free projection for geometry reconstruction. *ACM Trans. on Graphics (Proc. of SIGGRAPH)*, 26(3):22:1–22:6, 2007. 2
- [37] X. Liu, Z. Han, Y.-S. Liu, and M. Zwicker. Point2Sequence: Learning the shape representation of 3D point clouds with an attention-based sequence to sequence network. *arXiv preprint arXiv:1811.02565*, 2018. 2
- [38] C. Loop. Smooth subdivision surfaces based on triangles. *Master's thesis, University of Utah, Department of Mathematics*, 1987. 5
- [39] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 4
- [40] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas. Frustum pointnets for 3D object detection from rgb-d data. *arXiv preprint arXiv:1711.08488*, 2017. 2
- [41] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, 2017. 1, 2
- [42] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *In Advances in Neural Information Processing Systems (NIPS)*, pages 5099–5108, 2017. 1, 3
- [43] D. Rethage, J. Wald, J. Sturm, N. Navab, and F. Tombari. Fully-convolutional point networks for large-scale point clouds. *arXiv preprint arXiv:1808.06840*, 2018. 2
- [44] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-assisted Intervention*, pages 234–241. Springer, 2015. 2
- [45] R. Roveri, A. C. Öztireli, I. Pandele, and M. Gross. Point-ProNets: Consolidation of point clouds with convolutional neural networks. *Computer Graphics Forum*, 37(2):87–99, 2018. 2
- [46] Y. Shen, C. Feng, Y. Yang, and D. Tian. Mining point cloud local structures by kernel correlation and graph pooling. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, 2018. 3
- [47] 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 *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, pages 1874–1883, 2016. 1
- [48] Sketchfab. <https://sketchfab.com>. 5
- [49] P.-S. Wang, C.-Y. Sun, Y. Liu, and X. Tong. Adaptive O-CNN: A patch-based deep representation of 3D shapes. *ACM Trans. on Graphics (Proc. of SIGGRAPH Asia)*, 2018. 2
- [50] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro. High-resolution image synthesis and semantic manipulation with conditional GANs. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, 2018. 2
- [51] Y. Wang, F. Perazzi, B. McWilliams, A. Sorkine-Hornung, O. Sorkine-Hornung, and C. Schroers. A fully progressive approach to single-image super-resolution. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition Workshops*, June 2018. 2, 3, 4
- [52] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph cnn for learning on point clouds. *arXiv preprint arXiv:1801.07829*, 2018. 3
- [53] S. Wu, H. Huang, M. Gong, M. Zwicker, and D. Cohen-Or. Deep points consolidation. *ACM Trans. on Graphics*, 34(6):176, 2015. 2
- [54] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, pages 1912–1920, 2015. 5
- [55] Y. Xu, T. Fan, M. Xu, L. Zeng, and Y. Qiao. Spidercnn: Deep learning on point sets with parameterized convolutional filters. *Proc. Euro. Conf. on Computer Vision*, 2018. 1
- [56] Y. Yang, C. Feng, Y. Shen, and D. Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, volume 3, 2018. 2, 4
- [57] K. Yin, H. Huang, D. Cohen-Or, and H. Zhang. P2p-net: bidirectional point displacement net for shape transform. *ACM Trans. on Graphics (Proc. of SIGGRAPH)*, 37(4):152, 2018. 2
- [58] L. Yu, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng. Ecnnet: an edge-aware point set consolidation network. *Proc. Euro. Conf. on Computer Vision*, 2018. 1, 2, 5
- [59] L. Yu, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng. Pu-net: Point cloud upsampling network. In *Proc. IEEE Conf. on Computer Vision & Pattern Recognition*, pages 2790–2799, 2018. 1, 2, 3, 4, 5
- [60] W. Yuan, T. Khot, D. Held, C. Mertz, and M. Hebert. Pcn: Point completion network. In *Proc. Int. Conf. on 3D Vision*, pages 728–737. IEEE, 2018. 2
- [61] W. Zhang, H. Jiang, Z. Yang, S. Yamakawa, K. Shimada, and L. B. Kara. Data-driven upsampling of point clouds. *arXiv preprint arXiv:1807.02740*, 2018. 2
- [62] Y. Zhao, G. Li, W. Xie, W. Jia, H. Min, and X. Liu. Gun: Gradual upsampling network for single image super-resolution. *IEEE Access*, 6:39363–39374, 2018. 1, 2