

Point Cloud Completion by Skip-attention Network with Hierarchical Folding

Xin Wen¹, Tianyang Li¹, Zhizhong Han², Yu-Shen Liu^{1*}

¹School of Software, BNRist, Tsinghua University, Beijing, China

²Department of Computer Science, University of Maryland, College Park, USA

{x-wen16, lity16}@mails.tsinghua.edu.cn h312h@umd.edu liuyushen@tsinghua.edu.cn

Abstract

Point cloud completion aims to infer the complete geometries for missing regions of 3D objects from incomplete ones. Previous methods usually predict the complete point cloud based on the global shape representation extracted from the incomplete input. However, the global representation often suffers from the information loss of structure details on local regions of incomplete point cloud. To address this problem, we propose Skip-Attention Network (SA-Net) for 3D point cloud completion. Our main contributions lie in the following two-folds. First, we propose a skip-attention mechanism to effectively exploit the local structure details of incomplete point clouds during the inference of missing parts. The skip-attention mechanism selectively conveys geometric information from the local regions of incomplete point clouds for the generation of complete ones at different resolutions, where the skip-attention reveals the completion process in an interpretable way. Second, in order to fully utilize the selected geometric information encoded by skip-attention mechanism at different resolutions, we propose a novel structure-preserving decoder with hierarchical folding for complete shape generation. The hierarchical folding preserves the structure of complete point cloud generated in upper layer by progressively detailing the local regions, using the skip-attended geometry at the same resolution. We conduct comprehensive experiments on ShapeNet and KITTI datasets, which demonstrate that the proposed SA-Net outperforms the state-of-the-art point cloud completion methods.

1. Introduction

Recently, point cloud has received an extensive attention as a format of 3D objects, which can be easily accessed by 3D scanning devices and depth cameras. However, the raw point clouds produced by those devices are usually

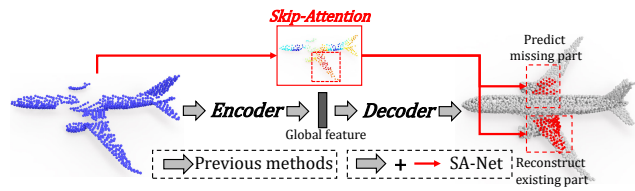


Figure 1. Illustration of the proposed skip-attention. Compared with the previous methods that simply rely on the global shape representation for completing point clouds, our skip-attention mechanism directly searches for informative local regions in input airplane shape, and selectively uses these regions for predicting the missing right wing or reconstructing the similar left wing (red).

sparse, noisy and mostly with serious missing regions due to the limited view angles or occlusion [47], which are difficult to be directly processed by the further shape analysis/rendering methods. Therefore, raw point cloud preprocessing becomes an important requirement for many real-world 3D computer vision applications. In this paper, we focus on the task of completing the missing regions of 3D shapes represented by point clouds.

The task of point cloud completion can be roughly decomposed into two targets [41, 47]. The first target is to preserve the geometric shape information of the original input point cloud, and the second target is to recover the missing regions according to the given inputs. In order to achieve such two targets, current studies usually followed the paradigm framework to learn a global shape representation from incomplete point clouds, which is further leveraged to estimate the missing geometric information [45, 47, 22]. However, the encoded global shape representation often suffers from the information loss of some structure details on local regions of incomplete point clouds, which should be fully preserved for further inferring the missing geometric information. As shown in Figure 1, to predict the complete wings of an airplane, the network should first preserve the existing left wing in the incomplete point cloud. And then, in order to infer the missing right wing, the network could refer to the existing left wing according to the pattern similarity between the regions of two similar wings.

*Corresponding author. This work was supported by National Key R&D Program of China (2018YFB0505400).

An intuitive idea to address this problem is to adopt the *skip-connection* mechanism like U-Net [35], which is widely used for local region reconstruction and reasoning in images. However, there are two problems for directly adopting skip-connection into point cloud completion. First, the previous skip-connection developed in [35] can not be directly applied to unordered inputs, since it concatenates the feature vectors according to the pixel order of 2D grids. Second, in the task of point cloud completion, not all the local region features under each level of resolutions will be of help for shape inferring and reconstruction. Equally revisiting them with skip-connection may introduce the information redundancy, and limit the feature learning ability of the entire network.

Therefore, in order to preserve the information of structure details while addressing the problem of skip-connection, we propose a novel deep neural network for point cloud completion, named *Skip-Attention Network* (SA-Net). The network is designed in an end-to-end framework, where an encoder-decoder architecture is specially designed for feature extraction and shape completion. The skip-attention refers to the attention based feature pipeline, which reveals completion process in an interpretable way. The skip-attention selectively conveys geometric information from the local regions of incomplete point clouds for the generation of the complete ones at different resolutions. The skip-attention enables the decoder to fully exploit and preserve the structure details on local regions. Compared with the skip-connection, the skip-attention can be generalized to unordered point clouds, since attention mechanism has no pre-requirements on the order of the input features. Moreover, our skip-attention provides an attentional choice for network to revisit the features under different resolutions, which allows the network to selectively incorporate the features encoded with desirable geometric information, and avoid the problem of information redundancy.

In order to fully utilize the selected geometric information from skip-attention at different resolutions, we further propose a structure-preserving decoder with *hierarchical folding* to generate complete point clouds. The hierarchical folding preserves the structure of point cloud generated in the upper layer, by progressively detailing the local regions using the skip-attended geometric information at the same resolution from the encoder. Specifically, the decoder has the same number of resolution levels as the encoder, with the skip-attention connecting each level of encoder to the corresponding level of decoder. In order to hierarchically fold the point clouds through levels, we propose to sample 2D grids with an increasing density from a 2D plane of fixed size. Compared with the decoders in existing point cloud completion methods [47, 41, 45], the proposed structure-preserving decoder can preserve the structure details on local regions under the whole resolution lev-

els, which enables the network to predict complete shape that maintains the global shape consistency while capturing more local region information. Our main contributions can be summarized as follows.

- We propose a novel Skip-Attention Network (SA-Net) for the point cloud completion task, which achieves state-of-the-art results. Moreover, the architecture of SA-Net can also be used for improving the performance of shape segmentation, and achieving the state-of-the-art results in unsupervised shape classification.
- We propose the skip-attention mechanism to fuse the informative local region features from encoder into the point features of decoder at different resolutions, which enables the network to infer the missing regions using more detailed geometry information from incomplete point clouds. In addition, skip-attention reveals the completion process in an interpretable way.
- We propose a structure-preserving decoder for high quality point cloud generation. It can progressively detail the point clouds at different resolutions with hierarchical folding, which hierarchically preserves the structure of complete shape at different resolutions.

2. Related Work

3D computer vision is an active research field in recent year [5, 11, 10, 12, 29, 13, 31], where the studies of 3D shape completion lead to many branches. For examples, geometry based [40, 2, 42, 23] methods exploit the geometric features of surface on the partial input to generate the missing part of 3D shapes, and alignment-based methods [37, 24, 32, 38] maintain a shape database and search for the similar patches to fill the incomplete regions of 3D shapes. Our method belongs to the deep learning based methods, which benefits a lot from the recent development of deep neural network in 3D computer vision [9, 28, 20, 16, 18, 15, 17, 14]. This branch can be further categorized according to the input form of 3D shapes.

Volumetric shape completion. 3D volumetric shape completion is a direction that is benefited a lot from the progress in 2D computer vision. Notable work like 3D-EPN [4] considered a progressive reconstruction of 3D volumetric shapes. And Han et. al [8] combined the inference of global structure with the local geometry refinement to directly generate the complete 3D volumetric shape of high resolution. More recently, the variational auto-encoder was introduced to learn a shape prior for inferring the latent representation of complete shapes [39]. Although fascinating improvements have been made in the research area of 3D volumetric data, the computational cost which is cubic to the resolution of input data makes it difficult to process fine-grained shapes.

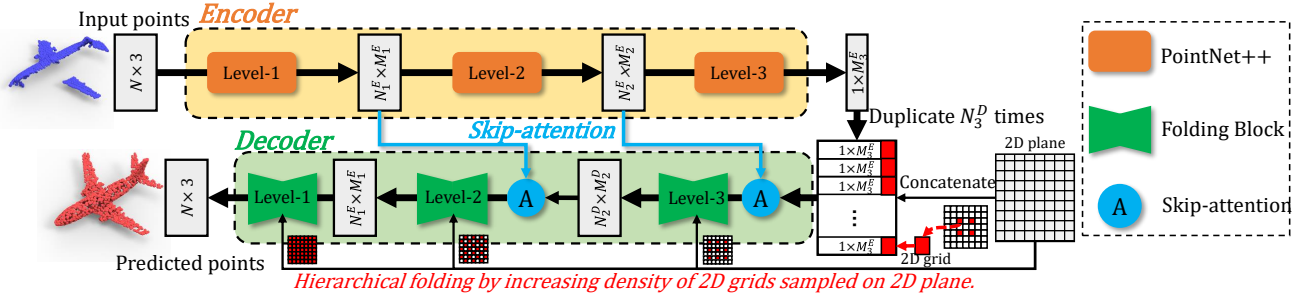


Figure 2. The overall architecture of SA-Net. SA-Net mainly consists of three modules: the encoder (yellow) aims to extract local region features from the input point clouds; the structure-preserving decoder (green) aims to reconstruct the complete point clouds and preserve the local region details; the skip-attention (sky blue) bridges the local region features in encoder and the point features in decoder.

Point cloud completion. The point cloud based 3D shape completion is a surging research area benefited from the pioneering work of PointNet [33] and PointNet++ [34]. As a compact representation of 3D shapes, point cloud can represent arbitrary detailed structure of 3D shape with a smaller storage cost compared to 3D volumetric data. Recent notable studies like PCN [47], FoldingNet [45] and AtlasNet [7] usually learn a global representation from partial point cloud, and generate the complete shape based on the learned global feature. Following the same practice, a tree-structured decoder was proposed in TopNet [41] for better structure-aware point cloud generation. By combining the reinforcement learning with the adversarial network, RL-GAN-Net [36] and Render4Completion [21] further improved the reality and consistency of the generated complete point cloud with the ground truth. However, most of these studies suffer from the information loss of structure details, as they predict the whole point cloud only from a single global shape representation.

3. The Architecture of SA-Net

Figure 2 shows the overall architecture of SA-Net, which consists of an encoder and a structure-preserving decoder. Between the encoder and the decoder, the skip-attention serves as the pipeline that connects the local region features (extracted from different resolutions in encoder) with the point features in the corresponding resolutions of decoder.

3.1. Encoder

Given the input point cloud of size $N=2,048$ with its 3-dimensional coordinates, the encoder of SA-Net aims to extract the features from the incomplete input point clouds. In SA-Net, we adopt the PointNet++ [34] framework as the backbone of our point cloud feature encoder. As shown in Figure 2, there are three levels of feature extraction, with the first level and the second level sampling the input point cloud into the size $N_1^E=512$ and $N_2^E=256$ (the superscript E denotes the encoder), and the last level grouping the input point cloud into a global representation. As a result, the encoder generates one global representation, and some lo-

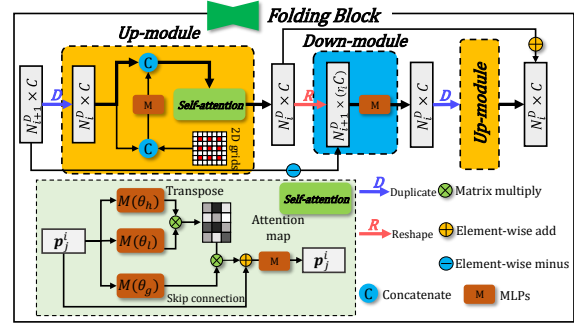


Figure 3. Illustration of the folding block, which consists of a down-module and two up-modules with self-attention inside. Folding block aims to lift the number of point features and refine the geometric information lying within these features.

cal region features extracted from different resolution levels for the input point clouds, respectively.

3.2. Structure-Preserving Decoder

Considering that the encoder extracts the local region features from different resolution levels, it is a natural practice for decoder to generate point features following the same way but with inverse resolution levels. This allows the skip-attention to establish a level-to-level connection between the extracted local region features in encoder and the generated point features in decoder. Inspired by this idea, we propose the structure-preserving decoder, which aims to progressively generate the complete point clouds and preserve the structure details of local regions under all resolution levels. Specifically, as shown in Figure 2, the structure-preserving decoder hierarchically folds the point clouds for three resolution levels, which is equal to the number of resolution levels in the encoder. Each resolution level of decoder consists of a skip-attention to convey the local region features from the same level of encoder, and a folding block to increase the number of point features.

3.3. Folding Block

Except for lifting the number of point features, the folding block also concerns the refinement of the expanded point features, which allows the decoder to produce more

consistent geometric details on the local regions of point clouds. Note that, such problem is usually ignored by previous methods, in which they either directly fold the entire point set based on the duplicated global representation [47, 45], or simply produce the point clouds through multi-layer perceptrons (MLPs) and reshape operations [41]. In SA-Net, we take the inspiration of the up-down-up framework from [26] to address this problem, which is adopted as the base of our folding block. Figure 3 shows the detailed structure of the folding block in the i -th level of decoder.

The up-module with hierarchical folding. As shown in the yellow part of Figure 3, for the input N_{i+1}^D (the superscript D denotes the decoder) point features from the previous level, the up-module first copies the point features by the time of up-sampling ratio $r_i = \frac{N_i^D}{N_{i+1}^D}$, and concatenates them with the 2D grids. Different from previous folding based decoders [26, 47, 45], which only have one resolution level for point cloud generation, the decoder in SA-Net progressively generates the point clouds for multiple resolution levels. In order to hierarchically fold point clouds through these levels, we propose to sample 2D grids with an increasing density from the 2D plane of fixed size. Specifically, for the N_i^D point features in the i -th level of decoder, the N_i^D 2D grids is evenly sampled from the 46×46 2D plane (the smallest number of square greater than 2,048), as illustrated in the up-module of Figure 3. These sampled 2D grids are then concatenated with the point features. After that, the point features with 2D grids are passed through MLPs and transformed into 3-dimensional latent codewords [45]. These 3-dimensional codewords are again concatenated with the point features in the i -th level of decoder.

In order to integrate the semantic and spatial relationships between these point features, we adopt a self-attention module with MLPs to establish the inner links between features, which aims to selectively fuse the similar features together through attention mechanism. This process is shown in the bottom half of Figure 3. Given the j -th point feature \mathbf{p}_j^i of i -th level in decoder, the skip-connection first calculates the attention scores $\{a_{j,k} | k = 1, 2, \dots, N_i^D\}$ between \mathbf{p}_j^i and all of the point features $\{\mathbf{p}_k^i | k = 1, 2, \dots, N_i^D\}$ in i -th level of decoder as

$$a_{j,k} = \frac{\exp(\mathbf{M}(\mathbf{p}_j^i | \theta_h)^T \cdot \mathbf{M}(\mathbf{p}_k^i | \theta_l))}{\sum_{n=1}^{N_i} \exp(\mathbf{M}(\mathbf{p}_j^i | \theta_h)^T \cdot \mathbf{M}(\mathbf{p}_n^i | \theta_l))} \quad (1)$$

where \mathbf{M} denotes the MLPs with parameters θ , and \mathbf{T} denotes the transposition operation. h and l indicate that two MLPs have different parameters. We take the weighted sum of point features $\{\mathbf{p}_k^i\}$ as the final context vector, and fuse it into the point feature \mathbf{p}_j^i as follow:

$$\mathbf{p}_j^i \leftarrow \mathbf{p}_j^i + \sum_{k=1}^{N_i^D} a_{j,k} \cdot \mathbf{M}(\mathbf{p}_k^i | \theta_g) \quad (2)$$

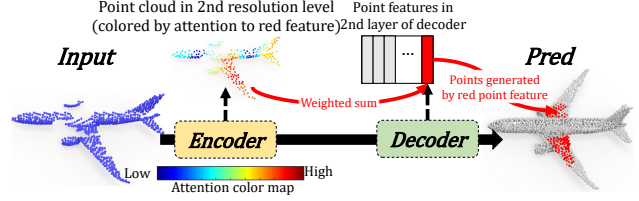


Figure 4. Illustration of the skip-attention. The skip-attention calculates the pattern similarity between local regions of complete point cloud (red points in *Pred*, which is generated by red feature) and the local regions of incomplete one. The similar local regions in the incomplete point clouds is selectively fused into the decoder with attentioned weighted sum.

The down-module. The point features expanded by the up-module actually occupy a small local region in the feature space, which can be aggregated as one local region feature through reshape and feature concatenation. Such aggregated local region feature can be regarded as a refined point feature of higher quality compared with the one in previous levels, since it contains not only the information from previous level of decoder, but also the detailed information produced by the current up-module. Then, followed by the MLPs and another up-module, the aggregated local region feature can be further used to reproduce the new point features with better structure details.

3.4. Skip-attention

The skip-attention serves as the pipeline to communicate the local region features extracted by encoder with the point features generated by decoder. It also interprets how the network completes shapes using information from incomplete ones. The skip-attention is designed for two purposes. First, when generating points that are located in the existing region of incomplete inputs, the skip-attention should fuse the feature of the same region from the encoder into the decoder, and guide the decoder to reconstruct more consistent structure details in such region. Second, when generating points that are located at the missing region of input, the skip-attention should search for referable similar regions in the original input point clouds, and guide the decoder to incorporate the shape of these similar regions as reference for inferring the shapes of missing regions. Both of the above purposes are achieved through an attention mechanism, as shown in Figure 4, where the semantic relativeness between point features in decoder and local region features in encoder are measured by attention scores, with the higher scores indicating the more significant pattern similarity (the wings of airplane). Then, the local region features are fused into point feature by weighted sum, and finally used for predicting related regions (also the wings of plane) in the complete point cloud.

There are different possible ways to calculate attentions for the skip-attention pipeline. In this paper, we do not explore the whole space but typically choose two straightforward

ward implementation, which work well in SA-Net. The first one in skip-attention is to directly adopt the learnable attention mechanism as described in the up-module. And the second one is to calculate the cosine similarity as the attention measurement between features. Compare with learnable attention, the unsmoothed (no softmax activation) cosine attention brings in more information from the previous encoder network, which can establish a strong connection between point features in decoder and local region features in encoder. On the other hand, the smoothed learnable attention can preserve more information from the original point features. For learnable attention, the attention score in the i -th resolution level is computed between the point feature \mathbf{p}_j^i from decoder and all of the local region features $\{\mathbf{r}_k^i | k = 1, 2, \dots, N_i^E\}$ from encoder, given as

$$a_{j,k}^L = \frac{\exp(M(\mathbf{p}_j^i | \theta_h^L)^T \cdot M(\mathbf{r}_k^i | \theta_l^L))}{\sum_{n=1}^{N_i^E} \exp(M(\mathbf{p}_j^i | \theta_h^L)^T \cdot M(\mathbf{r}_n^i | \theta_l^L))}, \quad (3)$$

where the superscript L denotes the word *learnable*. For cosine distance, the attention score is given as

$$a_{j,k}^C = \frac{(\mathbf{r}_k^i)^T \mathbf{p}_j^i}{\|\mathbf{r}_k^i\|_2 \|\mathbf{p}_j^i\|_2}, \quad (4)$$

where the superscript C denotes the word *cosine*. Same as the self attention in up-module, we fuse the weighted sum of local region features $\{\mathbf{r}_k^i\}$ into the point feature \mathbf{p}_j^i using element-wise addition, which is the same as Eq. (2). In ablation study (Sec. 4.2), we will quantitatively compare the performance of these two attentions.

3.5. Training

During training, the Chamfer distance (CD) \mathcal{L}_{CD} and the Earth Mover distance (EMD) \mathcal{L}_{EMD} are adopted as the optimization losses. The total loss for training is the weighted sum of the CD and EMD, defined as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{EMD} + \lambda \mathcal{L}_{CD}, \quad (5)$$

where λ is the weight parameter fixed to 10 for the experiments in our paper. The definition of \mathcal{L}_{CD} and \mathcal{L}_{EMD} will be detailed in **Supplementary**.

4. Experiments

By default, we use the cosine similarity based skip-attention for all experiments. In Sec. 4.2, we compare it with learnable attention. During evaluation, we mainly use Chamfer distance as the measurement to compare the predicted point clouds with the ground truth.

4.1. Evaluation of Completion Performance

Datasets. To evaluate the performance of SA-Net, we conduct experiments on two large scale datasets for point

cloud completion. For quantitative comparison, we follow [47] to evaluate our methods on ShapeNet dataset [3], and generate 8 partial point clouds for each object by back-projecting 2.5D depth images from 8 views into 3D. Unlike Render4Completion [21], we follow [41] to evaluate on the sparse input, which is more close to the real-world scenarios. We uniformly sample only 2,048 points on the mesh surfaces for both the complete and partial shapes. We also qualitatively evaluate SA-Net on KITTI dataset [6], since there is no ground truth for the incomplete shape of cars in KITTI.

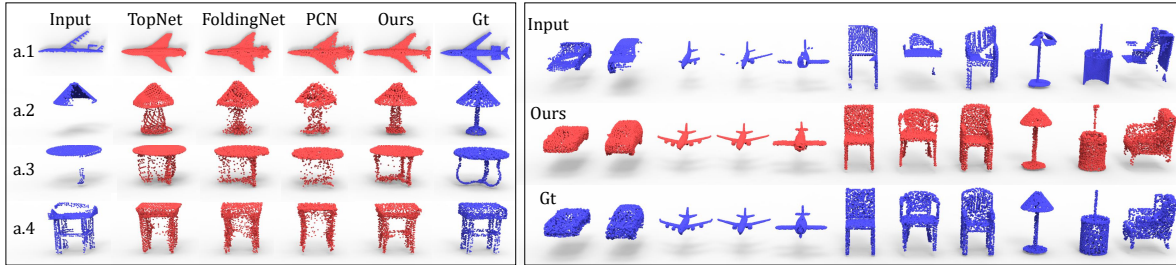
ShapeNet dataset. We use the per point Chamfer distance as the evaluation metric. In Table 1, SA-Net is compared with two point cloud completion methods PCN [47] and TopNet [41]. The reconstruction based unsupervised representation learning methods FoldingNet [45] and AtlasNet [7] are also included, since their basic encoder-decoder framework can also be generalized to point cloud completion task. The results of the above 4 methods are cited from [41]. The comparison shows that SA-Net outperforms the other methods on 6 out of 8 categories, and also achieves the best average Chamfer distance.

In Figure 5, we show the visualization results of point cloud completion using SA-Net and compare it with the other methods, from which we can find that SA-Net predicts more reasonable shape, while preserving more consistent geometric shapes for the existing parts. For example, in Figure 5(a.2) and 5(a.3), when predicting the missing lamp holders and table legs, the SA-Net generates more realistic shapes compared with the other three methods, and the points generated by SA-Net are arranged more tightly and shaped more close to the ground truth. In Figure 5(a.1) and 5(a.4), the SA-Net preserves the shapes of wings and the beam more consistently compared with the other three methods. The quantitative and qualitative improvements in shape completion task prove the effectiveness of skip-attention for introducing local region features, and the ability of structure-preserving decoder for utilizing the local region features to reconstruct completion point clouds. Moreover, in Table 2, we compare the number of trainable parameters in network of different methods, which shows that SA-Net has the least number of parameters, while achieving significantly better performance.

KITTI dataset. The KITTI dataset is collected from the real-world LiDAR scans, where the ground truth is missing for quantitative evaluation. Therefore, we qualitatively evaluate the performance of SA-Net by the visualization results. The complete cars are predicted using the parameters trained under car category in ShapeNet dataset for all methods in Figure 6. Note that in KITTI dataset, the point number of incomplete car has a large range of variation. In order to obtain a fixed point number of input, for the incomplete cars with more than 2048 points, we randomly choose

Table 1. Point cloud completion comparison on ShapeNet dataset in terms of per point Chamfer distance $\times 10^4$ (lower is better).

Methods	Average	Plane	Cabinet	Car	Chair	Lamp	Couch	Table	Watercraft
AtlasNet [7]	17.69	10.37	23.4	13.41	24.16	20.24	20.82	17.52	11.62
PCN [47]	14.72	8.09	18.32	10.53	19.33	18.52	16.44	16.34	10.21
FoldingNet [45]	16.48	11.18	20.15	13.25	21.48	18.19	19.09	17.8	10.69
TopNet [41]	9.72	5.5	12.02	8.9	12.56	9.54	12.2	9.57	7.51
SA-Net(Ours)	7.74	2.18	9.11	5.56	8.94	9.98	7.83	9.94	7.23



(a) Visualized completion comparison on ShapeNet.

(b) More completion results of SA-Net on ShapeNet.

Figure 5. Visualization of point cloud completion comparison on ShapeNet dataset. We compare SA-Net with other methods in (a), and in (b) we show more completion results of SA-Net.

Table 2. The number of trainable parameters in each method.

Methods	TopNet [41]	PCN [47]	FoldingNet [45]	SA-Net(Ours)
Params ($\times 10^6$)	9.97	5.29	2.40	1.67

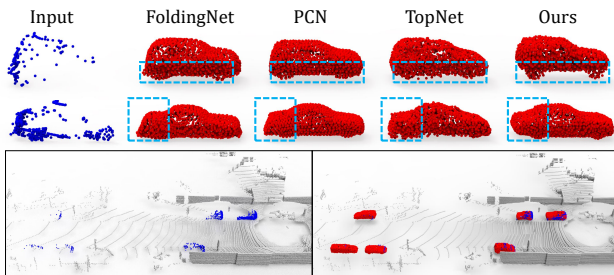


Figure 6. Visualization of the completion results on KITTI dataset.

2,048 points, otherwise, we randomly select points from the input to make up to the 2,048 points. The results are shown in Figure 6, from which we can find that our SA-Net predicts more structure details (car tiers) and shapes of higher quality (car trunks).

4.2. Ablation Study

In this subsection, we analysis the effect of important modules and hyper-parameters to SA-Net. All studies are typically conducted on the plane category for convenience.

Effect of attention. We developed three variations of SA-Net to verify the effectiveness of attention in SA-Net: (1) “No-skip” is the variation that removes the skip-attention from the SA-Net. (2) “Skip-L” is the variation that replaces the cosine attention in skip-attention by the learnable attention. (3) “Fold-C” is the variation that replaces the learnable attention by the cosine similarity in the self-attention of folding block. All three variations have the same structure as SA-Net except for the removed/replaced module. The results are shown in Table 3, in which the

original SA-Net achieves the best performance. The experimental results prove the effectiveness of attention used in SA-Net. The performance drop for replacing the attention in skip-attention (Skip-L) and self-attention (Fold-C) can be dedicated to the different design purposes for the two modules. The skip-attention aims to incorporate the local region features, and the unsmoothed cosine similarity allows more information to be fused into decoder. In contrast, the self-attention aims to learn a discriminative point features instead of simply merging the neighborhood features, therefore, smoothed weights (by softmax) in self-attention is more desirable for the network to preserve the original information of the point features. We specially note that, since removing the decoder of multi-resolution levels will also change the linkages of skip-attention, in Sec 4.3, we will instead evaluate the effectiveness of decoder on the task of unsupervised shape classification.

Table 3. The effect of each module to SA-Net (plane category).

Methods	No-skip	Skip-L	Fold-C	SA-Net
CD ($\times 10^4$)	2.31	2.25	2.34	2.18

Effect of optimization loss. To evaluate the effect of EMD loss and CD loss to SA-Net, we developed two variations: (1) “SA-Net-EMD” is the variation of SA-Net that is only trained using EMD loss; (2) “SA-Net-CD” is the variation that is only trained with CD loss. The comparison results are shown in Table 4, which proves that both EMD and CD contribute to the performance of SA-Net.

Table 4. The effect of each optimization loss (plane category).

Methods	SA-Net-EMD	SA-Net-CD	SA-Net
CD ($\times 10^4$)	2.39	2.23	2.18
EMD ($\times 10^2$)	3.06	4.58	3.02

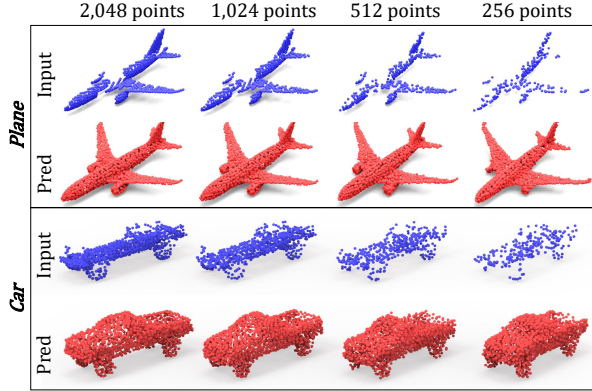


Figure 7. Visualization of completion results on different resolutions of input.

Effect of input point number. We analyze the robustness of SA-Net on various resolutions of inputs, especially for the performance on sparse input. In this experiment, we fix the number of output point clouds to 2,048, and evaluate the performance of SA-Net on the input point clouds with resolutions ranging from 256 to 2,048. For the point size less than 2,048, we use the same strategy in KITTI dataset to randomly select points from input, and lift the number of points up to 2,048. The model performance in terms of per point CD is reported in Table 5. In Figure 7, we visualize the completion quality under different point number of incomplete point clouds, in which the SA-Net shows a robust performance on all input resolutions.

Table 5. The effect of input point number (plane category).

#Points	2048	1024	512	256
CD ($\times 10^4$)	2.18	2.28	2.45	3.31

Visualization of skip-attention. In Figure 8, we visualize the attention in the second resolution level of decoder, which is to predict a complete plane. We compare the skip-attention learned for generating the empennage and part of the two wings. The points generated by the same point feature are colored by red in the left half of Figure 8(a) and 8(b), and the corresponding attention scores that point feature assigned to the local regions of incomplete point cloud are visualized in the right half. As shown in Figure 8(a), when generating points that belong to the empennage, the skip-attention searches for relative local regions (which is also the empennage) in the input point clouds for prediction. In Figure 8(b), when predicting the points of wings (where the right wing are missing), the skip-attention selects the region of left wing (by assigning higher attention) in incomplete point cloud for predicting the shape of both wings. Similar pattern is also observed on other categories as shown in Figure 8.

Visualization of hierarchical folding. In Figure 9, we visualize the hierarchical folding in decoder. We track the folding process of a specific vector colored by blue, and

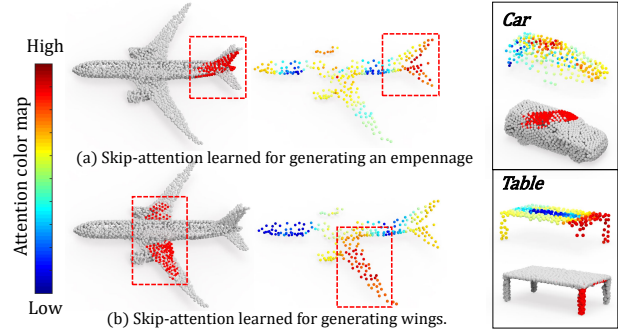


Figure 8. Visualization of the attention learned in skip-attention.

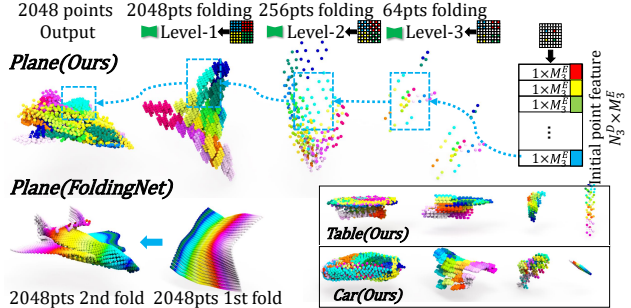
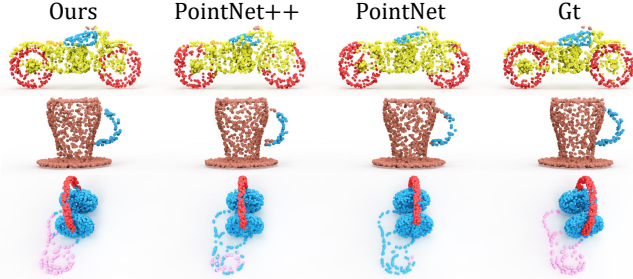


Figure 9. Visualization of the hierarchical folding in each level of decoder. We track the folding and point number expansion process of a specific initial vector colored by blue, and illustratively show the 2D grids sampling process.

denote the points derived from this blue vector with blue rectangular in each level. From a local perspective, we observe that each initial point feature successfully learns to generate a specific region on the plane. And in the case of blue initial point feature, it generates the left wing of the plane. On the other hand, from a global perspective, we can observe that the folding process of SA-Net does not restrictively follow the 2D manifold assumption like FoldingNet [45]. As pointed out by [41], enforcing learning from the 2D manifold structure may not be optimal for training, because the space of possible solutions is constrained. Therefore, the subtle deviation from 2D manifold, which is observed in SA-Net, is more flexible for learning to generate variant shapes and preserve better structure details. Both of the observations prove the effectiveness of hierarchical folding. In addition, we also visualize the folding process under car and table categories in Figure 9.

4.3. Model Analysis on Applications

Skip-attention for semantic segmentation. To further verify the effectiveness of skip-attention proposed in Sec 3.4, we conduct the semantic segmentation experiment on the ShapeNet dataset [46], where the dataset splittings follow the previous method of PointNet++ [34]. The segmentation variation (SA-Net-seg) of SA-Net uses exactly the same architecture as PointNet++, except for the skip-attention connecting the local region features in en-



(a) Visualized segmentation comparison on ShapeNet.



(b) More segmentation results of SA-Net on ShapeNet.

Figure 10. Segmentation visualization on ShapeNet. We compare SA-Net with baseline PointNet and PointNet++ in (a). In (b), we show more segmentation results of SA-Net. Note that there is no correspondence between colors and labels across object categories in (b).

coder with the features in interpolation layers. The comparison in terms of part-averaged intersection over union (pIoU, %) and mean per-class pIoU (mpIoU, %) [27] is shown in Table 6, from which we can find that SA-Net-seg drastically improves the segmentation performance compared with the baseline method of PointNet++. Specifically, the skip-attention improves the performance of backbone PointNet++ by 0.6% in terms of mIoUs. In Figure 10(a), we visualize the segmentation results and compared SA-Net-seg with the baseline PointNet and PointNet++, from which we can find that SA-Net-seg yields more precise prediction of semantic labels. Especially, the SA-Net-seg significantly improves the segmentation accuracy on the tier of the motorcycle, where the body and the tier are heavily overlapped with each other. Such improvement results from the local region features conveyed by skip-attention from the encoder, which helps the interpolation layers make more discriminative prediction in the local regions. Figure 10(b) gives more segmentation results.

Table 6. Semantic segmentation results (%) on ShapeNet.

Methods	pIoU	mpIoU
PointNet [33]	83.7	80.4
PointNet++ [34]	85.1	81.9
SO-Net [25]	84.9	81.0
DGCNN[43]	85.1	82.3
PointCNN [27]	86.1	84.6
SA-Net-seg(Ours)	85.7	83.0

Structure-preserving decoder for unsupervised representation learning in shape classification. In order to verify the effectiveness of our structure-preserving decoder, we further conduct unsupervised shape classification experiments on ModelNet40 [44]. The training and test settings on ModelNet40 also follow the PointNet++ [34]. In this experiment, we use a classification variation (SA-Net-clc) of SA-Net, in which we remove the skip-attention from SA-Net. The reason is that we use the global representation for predicting class label by a support vector machine (SVM), and remove the skip-attention can enhance the information embedded in the global representation, since it forces the decoder to decode a whole point cloud only based on the

Table 7. The classification comparison under ModelNet40.

Methods	Supervised	Accuracy(%)
PointNet[33]	Yes	89.2
PointNet++ [34]	Yes	90.7
PointCNN[27]	Yes	92.2
DGCNN[43]	Yes	92.2
SO-Net[25]	Yes	90.9
LGAN[1]	No	85.7
LGAN[1](MN40)	No	87.3
FoldingNet[45]	No	88.4
FoldingNet[45](MN40)	No	84.4
MAP-VAE[19]	No	90.2
L2G[30]	No	90.6
SA-Net-clc(Ours)	No	90.6

single global representation. The encoder and decoder are trained by reconstructing itself. In Table 7, we compare the classification performance of SA-Net-clc with counterpart methods, where all results are obtained under 1,024 points input without normal vector. From Table 7 we can find that our SA-Net-clc achieves the best performance in the unsupervised learning methods. The result of SA-Net-clc is also comparable with the supervised methods. Especially, we note that the classification accuracy of our SA-Net-clc is only 0.1% lower than the supervised PointNet++, which is exactly the same backbone used as our encoder.

5. Conclusion

We propose a novel Skip-Attention Network (SA-Net) for point cloud completion. Through the proposed skip-attention, SA-Net can effectively utilize the features of local regions in input point clouds for completion task. In order to exploit the local regions at different resolutions, the structure-preserving decoder is further proposed to progressively generate point clouds, and incorporate local region features at different resolutions. The completion experiments on ShapeNet and KITTI prove the effectiveness of SA-Net. The segmentation and classification experiments on ShapeNet and ModelNet40 further demonstrate the effectiveness of skip-attention and structure-preserving decoder, respectively.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3D point clouds. In *International Conference on Machine Learning*, pages 40–49, 2018.
- [2] Matthew Berger, Andrea Tagliasacchi, Lee Seversky, Pierre Alliez, Joshua Levine, Andrei Sharf, and Claudio Silva. State of the art in surface reconstruction from point clouds. In *Proceedings of the Conference of the European Association for Computer Graphics*, volume 1, pages 161–185, 2014.
- [3] Angel X Chang, Thomas Funkhouser, Leonidas J Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3D model repository. *arXiv:1512.03012*, 2015.
- [4] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3D-encoder-predictor CNNs and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5868–5877, 2017.
- [5] Ge Gao, Yu-Shen Liu, Meng Wang, Ming Gu, and Jun-Hai Yong. A query expansion method for retrieving online bim resources based on industry foundation classes. *Automation in construction*, 56:14–25, 2015.
- [6] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision Meets Robotics: The KITTI dataset. *International Journal of Robotics Research (IJRR)*, 2013.
- [7] Thibault Groueix, Matthew Fisher, Vladimir Kim, Bryan Russell, and Mathieu Aubry. AtlasNet: A papier-mâché approach to learning 3D surface generation. In *CVPR 2018*, 2018.
- [8] Xiaoguang Han, Zhen Li, Haibin Huang, Evangelos Kalogerakis, and Yizhou Yu. High-resolution shape completion using deep neural networks for global structure and local geometry inference. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 85–93, 2017.
- [9] Zhizhong Han, Xinhai Liu, Yu-Shen Liu, and Matthias Zwicker. Parts4Feature: Learning 3D global features from generally semantic parts in multiple views. In *International Joint Conference on Artificial Intelligence*, 2019.
- [10] Zhizhong Han, Zhenbao Liu, Junwei Han, Chi-Man Vong, Shuhui Bu, and C.L.Philip Chen. Mesh convolutional restricted boltzmann machines for unsupervised learning of features with structure preservation on 3D meshes. *IEEE Transactions on Neural Network and Learning Systems*, 28(10):2268 – 2281, 2017.
- [11] Zhizhong Han, Zhenbao Liu, Junwei Han, Chi-Man Vong, Shuhui Bu, and C.L.P. Chen. Unsupervised learning of 3D local features from raw voxels based on a novel permutation voxelization strategy. *IEEE Transactions on Cybernetics*, 49(2):481–494, 2019.
- [12] Zhizhong Han, Zhenbao Liu, Junwei Han, Chi-Man Vong, Shuhui Bu, and Xuelong Li. Unsupervised 3D local feature learning by circle convolutional restricted boltzmann machine. *IEEE Transactions on Image Processing*, 25(11):5331–5344, 2016.
- [13] Zhizhong Han, Zhenbao Liu, Chi-Man Vong, Yu-Shen Liu, Shuhui Bu, Junwei Han, and CL Philip Chen. BoSCC: Bag of spatial context correlations for spatially enhanced 3D shape representation. *IEEE Transactions on Image Processing*, 26(8):3707–3720, 2017.
- [14] Zhizhong Han, Zhenbao Liu, Chi-Man Vong, Yu-Shen Liu, Shuhui Bu, Junwei Han, and CL Philip Chen. Deep Spatiality: Unsupervised learning of spatially-enhanced global and local 3D features by deep neural network with coupled softmax. *IEEE Transactions on Image Processing*, 27(6):3049–3063, 2018.
- [15] Zhizhong Han, Honglei Lu, Zhenbao Liu, Chi-Man Vong, Yu-Shen Liu, Matthias Zwicker, Junwei Han, and C.L. Philip Chen. 3D2SeqViews: Aggregating sequential views for 3D global feature learning by CNN with hierarchical attention aggregation. *IEEE Transactions on Image Processing*, 28(8):3986–3999, 2019.
- [16] Zhizhong Han, Mingyang Shang, Yu-Shen Liu, and Matthias Zwicker. View inter-prediction GAN: Unsupervised representation learning for 3D shapes by learning global shape memories to support local view predictions. In *33rd AAAI Conference on Artificial Intelligence*, 2019.
- [17] Zhizhong Han, Mingyang Shang, Zhenbao Liu, Chi-Man Vong, Yu-Shen Liu, Matthias Zwicker, Junwei Han, and CL Philip Chen. SeqViews2SeqLabels: Learning 3D global features via aggregating sequential views by RNN with attention. *IEEE Transactions on Image Processing*, 28(2):658–672, 2018.
- [18] Zhizhong Han, Mingyang Shang, Xiyang Wang, Yu-Shen Liu, and Matthias Zwicker. Y2Seq2Seq: Cross-modal representation learning for 3D shape and text by joint reconstruction and prediction of view and word sequences. *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [19] Zhizhong Han, Xiyang Wang, Yu-Shen Liu, and Matthias Zwicker. Multi-Angle Point Cloud-VAE: Unsupervised feature learning for 3D point clouds from multiple angles by joint self-reconstruction and half-to-half prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10442–10451, 2019.
- [20] Zhizhong Han, Xiyang Wang, Chi-Man Vong, Yu-Shen Liu, Matthias Zwicker, and CL Chen. 3DViewGraph: Learning global features for 3D shapes from a graph of unordered views with attention. In *International Joint Conference on Artificial Intelligence*, 2019.
- [21] Tao Hu, Zhizhong Han, Abhinav Shrivastava, and Matthias Zwicker. Render4Completion: Synthesizing multi-view depth maps for 3D shape completion. 2019.
- [22] Tao Hu, Zhizhong Han, and Matthias Zwicker. 3D shape completion with multi-view consistent inference. In *34rd AAAI Conference on Artificial Intelligence*, 2020.
- [23] Wei Hu, Zeqing Fu, and Zongming Guo. Local frequency interpretation and non-local self-similarity on graph for point cloud inpainting. *IEEE Transactions on Image Processing*, 28(8):4087–4100, 2019.
- [24] Evangelos Kalogerakis, Siddhartha Chaudhuri, Daphne Koller, and Vladlen Koltun. A probabilistic model for

- component-based shape synthesis. *ACM Transactions on Graphics*, 31(4):55, 2012.
- [25] Jiaxin Li, Ben M Chen, and Gim Hee Lee. SO-Net: Self-organizing network for point cloud analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9397–9406, 2018.
- [26] Ruihui Li, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. PU-GAN: A point cloud upsampling adversarial network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7203–7212, 2019.
- [27] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. PointCNN: Convolution on x-transformed points. In *Advances in Neural Information Processing Systems*, pages 820–830, 2018.
- [28] Xinhai Liu, Zhizhong Han, Fangzhou Hong, Yu-Shen Liu, and Matthias Zwicker. LRC-Net: Learning discriminative features on point clouds by encoding local region contexts. In *The 14th International Conference on Geometric Modeling and Processing*, 2020.
- [29] Xinhai Liu, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Point2Sequence: Learning the shape representation of 3D point clouds with an attention-based sequence to sequence network. In *33rd AAAI Conference on Artificial Intelligence*, 2019.
- [30] Xinhai Liu, Zhizhong Han, Xin Wen, Yu-Shen Liu, and Matthias Zwicker. L2G Auto-Encoder: Understanding point clouds by local-to-global reconstruction with hierarchical self-attention. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 989–997. ACM, 2019.
- [31] Yu-Shen Liu, Karthik Ramani, and Min Liu. Computing the inner distances of volumetric models for articulated shape description with a visibility graph. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2538–2544, 2011.
- [32] Andelo Martinovic and Luc Van Gool. Bayesian grammar learning for inverse procedural modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 201–208. IEEE, 2013.
- [33] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [34] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *In Advances in Neural Information Processing Systems*, pages 5099–5108, 2017.
- [35] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [36] Muhammad Sarmad, Hyunjoo Jenny Lee, and Young Min Kim. RL-GAN-Net: A reinforcement learning agent controlled gan network for real-time point cloud shape completion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5898–5907, 2019.
- [37] Tianjia Shao, Weiwei Xu, Kun Zhou, Jingdong Wang, Dongping Li, and Baining Guo. An interactive approach to semantic modeling of indoor scenes with an rgbd camera. *ACM Transactions on Graphics*, 31(6):136, 2012.
- [38] Chao-Hui Shen, Hongbo Fu, Kang Chen, and Shi-Min Hu. Structure recovery by part assembly. *ACM Transactions on Graphics*, 31(6):180, 2012.
- [39] David Stutz and Andreas Geiger. Learning 3D shape completion from laser scan data with weak supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1955–1964, 2018.
- [40] Minhyuk Sung, Vladimir G Kim, Roland Angst, and Leonidas Guibas. Data-driven structural priors for shape completion. *ACM Transactions on Graphics (TOG)*, 34(6):175, 2015.
- [41] Lyne P Tchapmi, Vineet Kosaraju, Hamid Reza Tofighi, Ian Reid, and Silvio Savarese. TopNet: Structural point cloud decoder. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 383–392, 2019.
- [42] Duc Thanh Nguyen, Binh-Son Hua, Khoi Tran, Quang-Hieu Pham, and Sai-Kit Yeung. A field model for repairing 3D shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5676–5684, 2016.
- [43] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5):1–12, 2019.
- [44] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A deep representation for volumetric shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1912–1920, 2015.
- [45] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. FoldingNet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 206–215, 2018.
- [46] Li Yi, Vladimir G Kim, Duygu Ceylan, Ichao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas J Guibas. A scalable active framework for region annotation in 3D shape collections. In *International Conference on Computer Graphics and Interactive Techniques*, volume 35, page 210, 2016.
- [47] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. PCN: Point completion network. In *2018 International Conference on 3D Vision (3DV)*, pages 728–737. IEEE, 2018.