

Point cloud completion via structured feature maps using a feedback network

ZeJia Su
Shenzhen University
Shenzhen, China

zejiasu.36@gmail.com

Haibin Huang
Kuaishou Technology
Beijing, China

jackiehuanghaibin@gmail.com

Chongyang Ma
Kuaishou Technology
Beijing, China

chongyangm@gmail.com

Hui Huang
Shenzhen University
Shenzhen, China

hhzhiyan@gmail.com

Ruizhen Hu
Shenzhen University
Shenzhen, China

ruizhen.hu@gmail.com

Abstract

In this paper, we tackle the challenging problem of point cloud completion from the perspective of feature learning. Our key observation is that to recover the underlying structures as well as surface details, given partial input, a fundamental component is a good feature representation that can capture both global structure and local geometric details. We accordingly first propose FSNet, a feature structuring module that can adaptively aggregate point-wise features into a 2D structured feature map by learning multiple latent patterns from local regions. We then integrate FSNet into a coarse-to-fine pipeline for point cloud completion. Specifically, a 2D convolutional neural network is adopted to decode feature maps from FSNet into a coarse and complete point cloud. Next, a point cloud upsampling network is used to generate a dense point cloud from the partial input and the coarse intermediate output. To efficiently exploit local structures and enhance point distribution uniformity, we propose IFNet, a point upsampling module with a self-correction mechanism that can progressively refine details of the generated dense point cloud. We have conducted qualitative and quantitative experiments on ShapeNet, MVP, and KITTI datasets, which demonstrate that our method outperforms state-of-the-art point cloud completion approaches.

1. Introduction

In this paper, we study the problem of point cloud completion, i.e. recovering a full point cloud given a partial observation. It is an important component of many real-world applications, such as 3D data scanning [2], acquisition [28], robot navigation [15] and so on.

Point cloud completion is a challenging task as it needs

to recover both missing topological structures and geometric details from incomplete input. Traditional methods [3, 19] used hand-crafted features like surface smoothness or symmetry priors. Such empirical human-designed features suffer from performance degradation under changing illumination and in the presence of severe occlusion. Recently, deep neural network (DNN) based methods have been introduced into this task and have achieved promising improvements. Early works [4, 27] voxelized the 3D point cloud and applied a 3D convolutional neural network (CNN) to the volumetric data to complete the shapes. Due to the large computational cost of 3D CNNs, such methods generally output shapes with limited details. Pioneered by PointNet [22] and PointNet++ [23], other methods [47, 29] directly complete 3D point clouds. Compared to using voxels, learning on 3D points is more scalable and efficient, but also leads to the problem of feature learning due to the inherent irregularity and sparseness of point clouds.

A common practice in previous methods is to apply MLPs to point clouds in an encoder-decoder manner, where per-point features are aggregated to a global feature vector via a max-pooling operation and then decoded into a complete point cloud. Due to the use of pooling operations, these methods suffer from unavoidable information loss which results in unsatisfactory shape structures and blurred details in the missing regions. More recent methods have been proposed to improve feature learning, including use of shape priors [21], a skip-attention mechanism [36], and so on [49, 42]. However, most of these methods adopt feature vectors as the global representation and suffer from non-uniform distributions of the dense prediction results, due to the irregular and discrete nature of the point cloud, especially when the input shapes have large missing parts. As shown in Figure 1, the distributions of the point clouds predicted by previous methods are uneven *e.g.* in regions between the known and predicted parts.

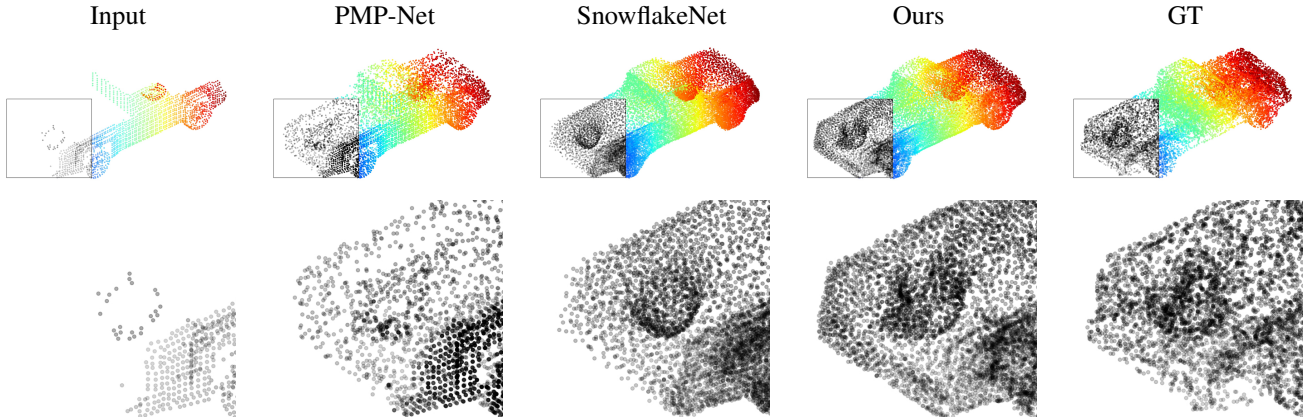


Figure 1. Our approach can generate complete point clouds from partial input with finer details and more uniform surface points compared to state-of-the-art methods.

To tackle these challenges, we propose a new point cloud completion method with two novel modules: FSNet for structured feature learning, which captures well both global structure and local geometric details of the input point cloud, and IFNet, for progressive detail generation and uniformity enhancement. Given a partial input, point-wise features are first extracted by an encoder and FSNet is then applied to form a 2D structured feature map. Instead of a max-pooling operation, FSNet relies on multi-head attention [30] to aggregate the features, where a learnable set and the input features are used as queries and key-value pairs, respectively, and the output matrices constitute the structured feature map. By exploiting latent patterns among local regions, FSNet can learn a global representation with rich semantic information which is distinguishable at the instance level. The 2D structured feature map then allows us to utilize a 2D convolutional neural network as a decoder to predict the complete structures and generate a coarse point cloud as the input to the upsampling stage. Motivated by the success of feedback mechanisms in the image super-resolution task [9], we propose IFNet with a self-correction strategy to upsample the coarse point cloud in a progressive manner. A sparse encoding module is applied to extract sparse features from the coarse point cloud and input, and followed by an upsampling module to generate initial dense features. The final dense point cloud is then iteratively refined through IFNet, where details are progressively added and uniformity is gradually improved by feeding back the projection error to the initial dense features.

We have conducted experiments on ShapeNet, MVP and KITTI datasets to evaluate our method. The results show that our approach can outperform state-of-the-art methods.

To summarize, our contributions are as follows:

- FSNet, a novel feature aggregation network which adaptively organizes the unordered features into a 2D structured feature map, which can retain more infor-

mation and represent more fine-grained global features than a feature vector,

- IFNet, an iterative feedback network to upsample and refine the coarse completion via a multi-step self-correction procedure, which can further exploit local details and facilitate consolidation and uniformity of the point cloud, and
- an evaluation of different algorithms on both synthetic and real-world datasets, which shows that our approach outperforms state-of-the-art methods in terms of both completion quality and surface uniformity.

2. Related Work

2.1. 3D shape completion

3D shape completion has drawn increasing attention in recent years. Early geometric methods [3, 19] complete objects by leveraging predefined geometric features, such as surface smoothness or symmetry priors. Since these methods rely on hand-crafted features, they are only valid under particular circumstances. Recently, convolutional neural networks (CNNs) have been widely used for processing regularly arranged data due to their ability to learn features. Some methods [4, 27] introduce 3D voxel grids or distance fields as the representation for 3D data, and utilize a 3D CNN to complete objects. However, use of a 3D CNN leads to large computational and memory costs that are cubic in the resolution of the volumetric data, while reducing the resolution limits the processing of fine-grained shapes.

Recent works tend to adopt 3D point clouds as the representations of 3D objects due to their convenience and flexibility. Pioneered by Pointnet [22], several methods use MLPs for point cloud completion under an encoder-decoder framework, where per-point features are aggregated to a




Figure 2. Pipeline.

global feature vector (GFV) by the max-pooling operation. PCN [47] and MSN [18] complete the point cloud in a coarse-to-fine fashion. TopNet [29] introduces a hierarchical tree-structure network that takes the geometric structure into consideration. RL-GAN-Net [26] and Render4Completion [10] focus on adversarial learning to improve the realism and consistency of the generated shape. However, these methods suffer from loss of structural details, as they predict the point cloud only from a single global vector.

To better preserve shape structures and complete surface details, SA-Net [36] and ASHF-Net [51] introduce a skip-attention mechanism to further revisit the low-level features from the encoder. SoftpoolNet [34] proposes a soft pooling module to replace the max-pooling operator, to keep more information by considering multiple features. NSFA [49] aggregates different features to represent the known and missing parts separately. GRNet [42] and VE-PCN [31] introduce 3D voxels as the intermediate representation to help the network to infer the complete shape. PF-Net [11], DeCo [1], and PoinTr [46] only generate the missing part of the object to preserve the spatial arrangements of the original part. VRCNet [21] and ASFM-Net [39] improve the global features by narrowing the distribution difference between the incomplete and complete point clouds. CRN [32] and SnowflakeNet [40] generate the complete point cloud in a progressive manner. Other notable work such as PMP-Net [37] formulates completion as a point cloud deformation process, where point-wise paths are predicted to move each point of the incomplete input to complete the point cloud. SpareNet [41] proposes a style-based point generator with adversarial rendering for point cloud completion. Cycle4Completion [35] improves completion quality by establishing a geometric correspondence between complete shapes and incomplete ones.

2.2. Point cloud upsampling

Point cloud upsampling aims to generate a uniform dense point cloud to represent the underlying surface of the object within local patches. PU-Net [45] uses PointNet++ [23] to extract multi-scale features and expands the features through multi-branch MLPs. With additional edge and surface annotations, EC-Net [44] improves the edge quality of PU-Net by formulating an edge-aware joint loss to learn the geometry of edges. MPU [43] proposes a

progressive network to upsample points in multiple stages, which requires supervision of intermediate outputs. PUGAN [17] presents a generative adversarial network (GAN) to learn the distribution of dense point clouds. PUGeoNet [25] proposes a geometric-centric network by learning local parameterization and normal direction for each point, which needs additional normal annotations. Recently, PU-GCN [24] designed a novel feature expansion network called NodeShuffle, which utilizes graph convolutional networks (GCNs) to encode local information.

3. Method

We now formally introduce our proposed method. As shown in Figure 2, we use a two-stage pipeline which works in a coarse-to-fine manner and consists of a coarse structure completion stage and a point cloud upsampling stage. More specifically, taking a partial point cloud $P^i \in \mathbb{R}^{N^i \times 3}$ as input, a novel feature structuring module FSNet firstly aggregates point-wise features encoded from P^i into a 2D structured feature map SFM and feeds it to a point cloud decoding network to produce a coarse point cloud $P^c \in \mathbb{R}^{N^c \times 3}$, where N^i and N^c denote the numbers of input points and points after coarse completion, respectively. Then, a sparse encoding module and a novel iterative feedback network IFNet are applied sequentially to upsample and refine the coarse completion P^c , followed by an offset regression module to obtain a dense point cloud $P^d \in \mathbb{R}^{N^d \times 3}$ with fine-grained details, where N^d denotes the number of points after dense completion.

3.1. Coarse Point Cloud Generation

3.1.1 Background

The goal of the first stage is to generate a coarse and complete point cloud. The main challenge is to recover a coarse shape that has correct global structures while maintaining sufficient local geometric details with respect to the input. To this end, we propose a feature structuring network, FSNet, that can adaptively aggregate point-wise features into a 2D feature map and then decode it into a coarse point cloud, as shown in Figure 3 (a).

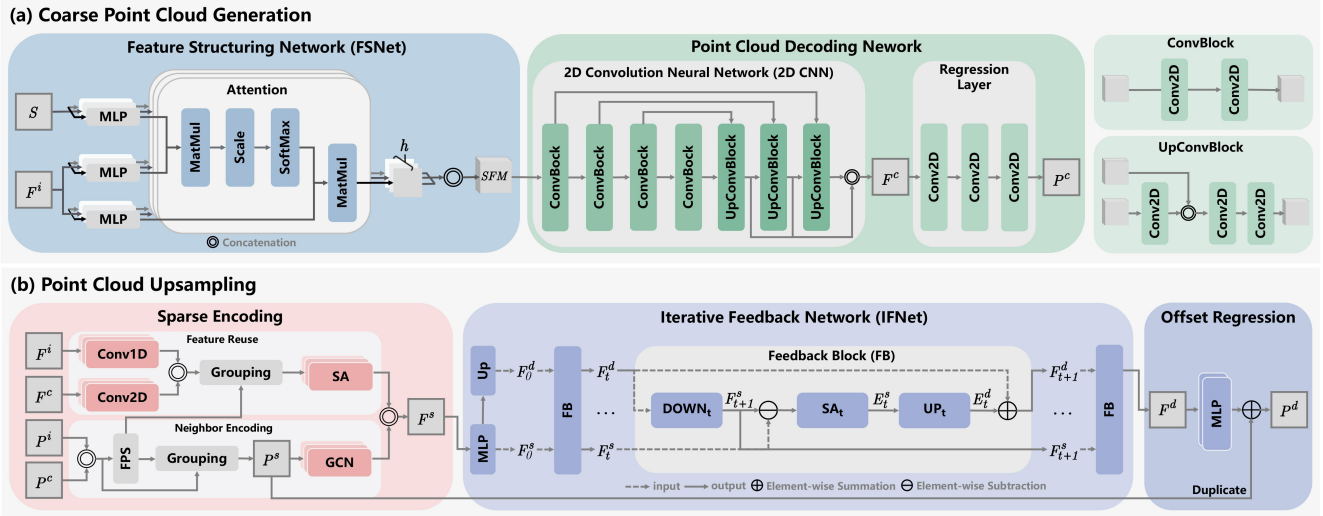


Figure 3. Network architectures of (a) our coarse point cloud generation module and (b) our point cloud upsampling module.



Figure 4. Normalized attention weights of some representative channels in the SFM.

3.1.2 Feature structuring network

To preserve permutation invariance of the input points, max pooling is widely used by previous methods, but it often leads to unavoidable information loss of local details. Inspired by Set Transformer [16], we propose FSNet, which utilizes multi-head attention [30] instead of max pooling to aggregate point-wise features in a data-driven way, while ignoring the order of the input points.

Specifically, given point-wise features $F^i \in \mathbb{R}^{N^i \times C^i}$ extracted from the input point cloud, we explore the latent patterns in local regions via an attention mechanism. We first define a learnable set with k vectors $S \in \mathbb{R}^{k \times C^i}$ and the F^i as the queries and key-value pairs. Then, we linearly project the queries, keys and values h times through different MLPs to dimension d . We further perform an attention function [30] on the projected features, yielding h matrices with shape $k \times d$, where each matrix is generated as the weighed sum of the projected input features and the attention weights which determine the geometric region corresponding to the channel extracted from the input point cloud. These matrices are then concatenated to generate a 2D structured feature map $SFM \in \mathbb{R}^{h \times k \times d}$ with shape $k \times d$ and h channels. Note that the output feature map is independent of the order of the input points, and naturally takes all vectors of input features into consideration thanks to the attention mechanism.

To better understand the learned SFM, we visualize the normalized attention weights of some representative channels using heatmaps in Figure 4. We observe that the distribution of attention weights shows different patterns across channels; some reveal semantic correlation, while others are geometrically complementary, indicating the structured feature map contains different combinations of local information from the input point cloud.

3.1.3 Point cloud decoding network

With the learned SFM, the next step is to generate a coarse but complete point cloud that recovers the overall structure of the target object. The key challenge here is to infer the features of missing regions and recover the underlying structure for diverse topologies. Our key observation is that since SFM has a regular 2D format and contains rich local information about the input, we may model the point cloud structure via 2D convolution operations. Thus, we build a decoding network that learns to generate a point cloud from the SFM. It is composed of two parts: a 2D CNN with UNet structure and a regression layer. In more detail, given the SFM as input, the 2D CNN produces coarse features $F^c \in \mathbb{R}^{C^c \times k \times d}$, and then the regression layer transforms F^c to an intermediate output P^c with shape $3 \times (k/2) \times (d/2)$, which is reshaped to form the corresponding coarse point cloud $P^c \in \mathbb{R}^{N^c \times 3}$.

Note that as the structure information is implicitly modeled in the 2D CNN, local patches in P^c actually correspond to local neighboring point regions, which represent the local structure of the object. Figure 5 shows some example local point regions corresponding to local patches in P^c , with increasing sizes from the top to bottom rows, where we use sliding windows of different sizes to extract the lo-




Figure 5. Representative structures learned by the decoding network. Local patches in the 2D structured point cloud \bar{P}^c correspond to local point regions, which grow with increasing patch size from the top row to the bottom row.

cal patches. We can see that point regions corresponding to smaller patches represent low-level structures of the object, and point regions grow with increasing local patch size in \bar{P}^c to represent high-level structures of the object. This shows that the CNN-based decoding network can precisely model the structure of point cloud using the given SFM.

3.2. Point Cloud Upsampling

3.2.1 Approach

After generating the coarse point cloud P^c , a point cloud upsampling stage follows to produce a dense point cloud P^d with fine-grained details and uniform distribution. As Figure 3 (b) shows, our proposed point cloud upsampling stage has three steps: sparse encoding, feature expansion, and offset regression.

3.2.2 Sparse encoding

The first step is a sparse encoding module which serves as sparse feature preparation and provides both local and contextual information to the following steps. It consists of two branches: a neighbor encoding path to preserve geometric structures, and a feature reuse path to explore structural details in local regions by revisiting the input features F^i and coarse features F^c . For the neighbor encoding path, we first combine P^c with P^i , followed by farthest point sampling and a grouping operation, which produces a sparse and relatively uniform point cloud $P^s \in \mathbb{R}^{N^s \times 3}$, where N^s denotes the number of sparse sample points. We then follow the GCN structure in DGCNN [33] with adaption of EdgeConv [33] as the GCN layer, which sequentially aggregates features from the local neighbors for each point. For the feature reuse path, F^i and F^c are passed to a sequence of 1D and 2D convolution layers separately. The output features are concatenated, followed by a grouping operation

and a sequence of self-attention units [48] to enhance feature integration. At the end, features from the two paths are concatenated as the sparse features $F^s \in \mathbb{R}^{N^s \times C^s}$ and fed into the feature expansion step.

3.2.3 Iterative feedback network

Inspired by DBPN [9] for image super resolution tasks, we introduce a feedback mechanism into feature expansion by constructing an iterative feedback network (IFNet) that can effectively expand the sparse features through multi-step refinement. The network first reduces the dimension of F^s from C^s to c using an MLPs and expands it through an upsampling unit, producing the initial sparse features $F_0^s \in \mathbb{R}^{N^s \times c}$ and dense point features $F_0^d \in \mathbb{R}^{rN^s \times c}$, where r denotes the upsampling ratio. It then feeds the initial features into a sequence of feedback blocks (FBs) to perform self-correction.

More specifically, a feedback block comprises an upsampling unit, a downsampling unit and a self-attention unit, defined as:

$$\text{scale down : } F_{t+1}^s = \text{DOWN}_t(F_t^d, r) \quad (1)$$

$$\text{difference : } E_t^s = \text{SA}_t(F_{t+1}^s - F_t^s) \quad (2)$$

$$\text{scale difference up : } E_t^d = \text{UP}_t(E_t^s, r) \quad (3)$$

$$\text{corrected features : } F_{t+1}^d = F_t^d + E_t^d \quad (4)$$

where SA denotes the self-attention unit [48], and UP(F, r) and Down(F, r) denote upsampling and downsampling operations on features F with a ratio r . Detailed network architectures are shown in the supplementary materials.

Each feedback block takes sparse features $F_t^s \in \mathbb{R}^{N^s \times c}$ and dense features $F_t^d \in \mathbb{R}^{rN^s \times c}$ as inputs, and maps the dense features F_t^d to new sparse features $F_{t+1}^s \in \mathbb{R}^{N^s \times c}$. The difference $E_t^s \in \mathbb{R}^{N^s \times c}$ between known sparse features F_t^s and the reconstructed sparse features F_{t+1}^s is computed, followed by a self-attention unit to produce a more discriminative error by capturing long-range contextual information. Afterwards, the difference is mapped to intermediate dense features $E_t^d \in \mathbb{R}^{rN^s \times c}$, and the corrected dense features $F_{t+1}^d \in \mathbb{R}^{rN^s \times c}$ are obtained by adding E_t^d to F_t^d . Intuitively, each block performs a self-correction procedure by feeding back the projection error to the initial dense features, which enables the networks to produce dense features that help to produce a dense point cloud with fine-grained details and uniform distribution.

3.2.4 Offset regression

In the last process, we generate the dense point cloud P^d based on the produced dense features $F^d \in \mathbb{R}^{N^d \times c}$ by predicting offsets between the coordinates of points in the sparse and dense point clouds, where the residual offsets

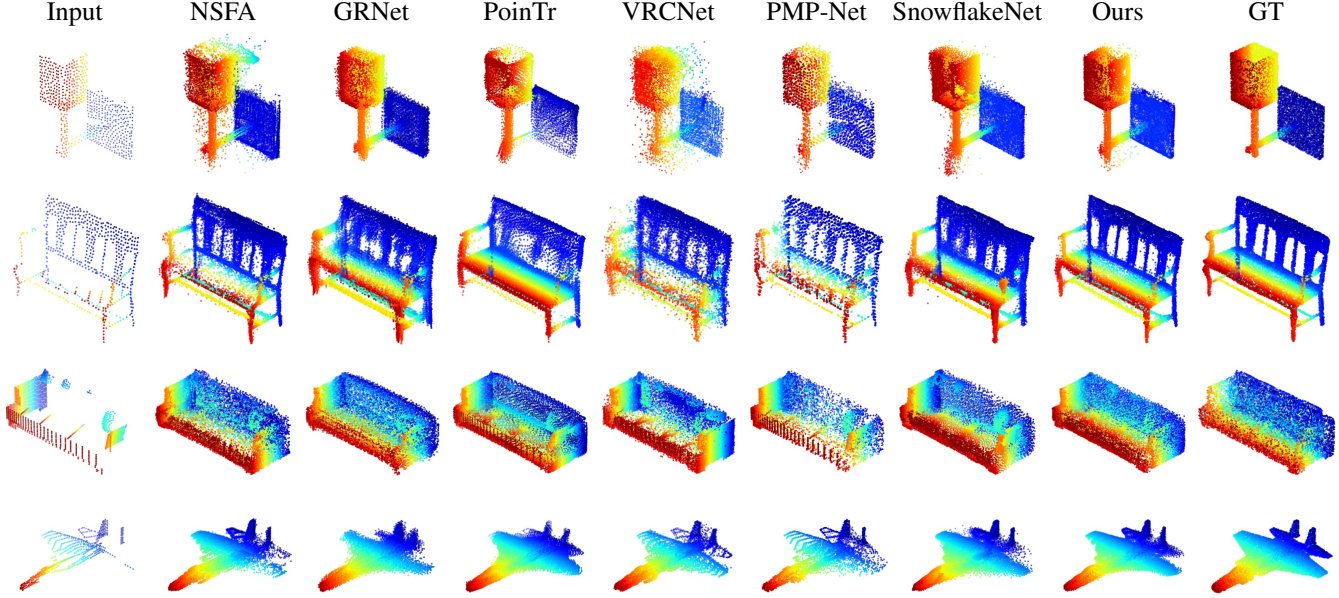


Figure 6. Completion results on the ShapeNet dataset.

are regressed through two layers of MLPs and the sparse point cloud P^s is replicated r times before residual summation. Note that the number of dense completion points $N^d = rN^s$.

3.3. Loss function

We choose Chamfer distance (CD) as the reconstruction loss due to its efficiency, following PSGN [5] and TopNet [29]. The Chamfer distance between two point sets X and Y is defined as:

$$L_{CD}(X, Y) = L_{X, Y} + L_{Y, X}, \quad (5)$$

where

$$L_{X, Y} = \frac{1}{|X|} \sum_{x \in X} \min_{y \in Y} \|x - y\|_2^2. \quad (6)$$

Since we generate the complete point cloud in a coarse-to-fine fashion, we jointly optimize the coarse point cloud P^c and dense point cloud P^d via the CD loss. The overall training loss is defined as:

$$L = L_{CD}(P^s, P^g) + L_{CD}(P^d, P^g) \quad (7)$$

where $P^g \in \mathbb{R}^{N^g \times 3}$ denotes the ground truth point cloud, and N^g denotes the number of ground truth points.

4. Experiments

4.1. Datasets

We evaluated our method using three popular benchmarks, including ShapeNet [38], MVP [21], and KITTI [6]:

4.1.1 ShapeNet

The ShapeNet dataset for point cloud completion is derived from PCN [47]; 30,974 samples are selected in 8 categories. The ground truth point clouds containing 16,384 points are uniformly sampled from mesh surfaces. The partial point clouds are generated by back-projecting 2.5D depth maps from 8 random views into 3D. For a fair comparison, we use the same training / validation / test splits as PCN.

4.1.2 MVP

The MVP dataset consists of 16 categories of partial and complete point clouds generated from CAD-models selected from the ShapeNet dataset [38]. There are 62400 and 41600 shape pairs in the training and testing sets, respectively. Unlike other datasets, the complete point clouds in MVP have different resolutions, including 2048, 4096, 8192 and 16384, so can be used to evaluate completion quality at different resolutions. We use the same training / test splits as in VRCNet [21] to evaluate our method.

4.1.3 KITTI

The KITTI dataset is composed of a sequence of real-world LiDAR scans, also derived from PCN [47]. For each frame, a car point cloud is extracted within the car object bounding boxes, resulting in 2,401 partial point clouds. There is no ground truth for this dataset.

Method	Plane	Cabinet	Car	Chair	Lamp	Sofa	Table	Vessel	Avg.
PCN [47]	1.40	4.45	2.45	4.84	6.24	5.13	3.57	4.06	4.02
TopNet [29]	2.15	5.62	3.51	6.35	7.50	6.95	4.78	4.36	5.15
MSN [18]	1.54	7.25	4.71	4.54	6.48	5.89	3.80	3.85	4.76
NSFA [49]	1.75	5.31	3.43	5.01	4.73	6.41	4.00	3.56	4.27
PF-Net [11]	1.55	4.43	3.12	3.96	4.21	5.87	3.35	3.89	3.80
CRN [32]	1.46	4.21	2.97	3.24	5.16	5.01	3.99	3.96	3.75
GRNet [42]	1.53	3.62	2.75	2.95	2.65	3.61	2.55	2.12	2.72
PoinTr [46]	0.99	4.80	2.52	3.68	3.07	6.53	3.10	2.02	3.34
VRCNet [21]	1.53	4.66	2.66	4.62	5.50	5.70	4.39	3.58	4.08
PMP-Net [37]	1.22	4.18	2.85	3.51	2.14	4.22	2.89	1.88	2.86
SnowflakeNet [40]	0.86	3.40	2.36	2.68	2.06	4.46	2.16	1.74	2.47
ASHF-Net [51]	1.40	3.49	2.32	2.82	2.52	3.48	2.42	1.99	2.55
Ours	0.80	3.53	2.13	2.48	1.79	3.64	1.98	1.81	2.27

Table 1. Shape completion results (CD loss $\times 10^4$) on the ShapeNet dataset.

Points	2,048		4,096		8,192		16,384	
	CD	F1	CD	F1	CD	F1	CD	F1
PCN [47]	9.77	0.320	7.96	0.458	6.99	0.563	6.02	0.638
TopNet [29]	10.11	0.308	8.20	0.440	7.00	0.533	6.36	0.601
MSN [18]	7.90	0.432	6.17	0.585	5.42	0.659	4.90	0.710
CRN [32]	7.25	0.434	5.83	0.569	4.90	0.680	4.30	0.740
ECG [20]	6.64	0.476	5.41	0.585	4.18	0.690	3.58	0.753
VRCNet [21]	5.96	0.499	4.70	0.636	3.64	0.727	3.12	0.791
PMPNet [37]	6.33	0.479	4.63	0.584	3.52	0.680	2.79	0.753
SnowflakeNet [40]	6.06	0.500	4.80	0.615	3.49	0.739	2.75	0.805
Ours	5.53	0.503	4.20	0.648	3.19	0.760	2.33	0.810

Table 2. Shape completion results (CD loss $\times 10^4$ and F-score@1%) with various resolutions on the MVP dataset.

4.2. Implementation Details

We use DGCNN [33] as the encoder to extract features from the input point cloud. Both d and k are set to 64 for generating a coarse completion P^c with $N^c = 1024$ points. The number of channels h for SFM is set to 32, determined by experiment. We typically sample $N^s = 1024$ points to obtain the sparse point cloud P^s , which is sufficient to cover the overall structures of the objects. We may generate dense output P^d with various resolutions, including $N^d = 2048, 4096, 8192, 16384$, where the upsampling ratio r is set to 2, 4, 8, and 16, respectively.

Our networks are implemented using PyTorch and optimized using an Adam optimizer [13] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The initial learning rate is 10^{-3} and is multiplied by 0.7 per 10 epochs. We train the networks with a batch size of 32 on four NVIDIA TITAN Xp GPUs.

4.3. Comparison with State-of-the-Art

We have compared our method to several state-of-the-art point cloud completion approaches, both quantitatively

and visually.

4.3.1 Results on ShapeNet

A quantitative comparison of our method to other methods on the ShapeNet dataset is shown in Table 1: our method outperforms all competitive methods in terms of average CD across all categories. Qualitative results are shown in Figure 6: our method can precisely predict the missing parts of the input while other methods tend to output blurred point clouds in the missing region. Moreover, our method can produce more uniformly-distributed dense point clouds with less noise than other methods, benefiting from IFNet with its self-correction procedure.

4.3.2 Results on MVP

Since our method can generate complete point cloud with various resolutions by modifying the upsampling rate r , we compare our method with existing methods that support multi-resolution completion. Except for CD loss, we also

Methods	FD ($\times 10^{-3}$)	MMD ($\times 10^{-3}$)	Consistency ($\times 10^{-3}$)	Uniformity		
				0.4%	0.8%	1.2%
PCN [47]	2.235	1.366	1.557	3.662	7.710	10.823
TopNet [29]	5.354	0.636	0.568	1.353	1.219	0.950
MSN [18]	0.434	2.259	1.951	0.822	0.523	0.383
NSFA [49]	1.281	0.891	0.491	0.992	0.767	0.552
PF-Net [11]	1.137	0.792	0.436	0.881	0.682	0.491
CRN [32]	1.023	0.872	0.431	0.870	0.673	0.485
GRNet [42]	0.816	0.568	0.313	0.632	0.489	0.352
VRCNet [21]	2.586	0.378	0.259	0.751	0.582	0.416
SnowflakeNet [40]	0.220	0.664	0.557	0.979	0.730	0.499
ASHF-Net [51]	0.773	0.541	0.298	0.602	0.466	0.335
Ours	0.586	0.336	0.314	0.506	0.313	0.148

Table 3. Shape completion results on the KITTI dataset: fidelity distance (FD), minimal matching distance (MMD), consistency, and uniformity.

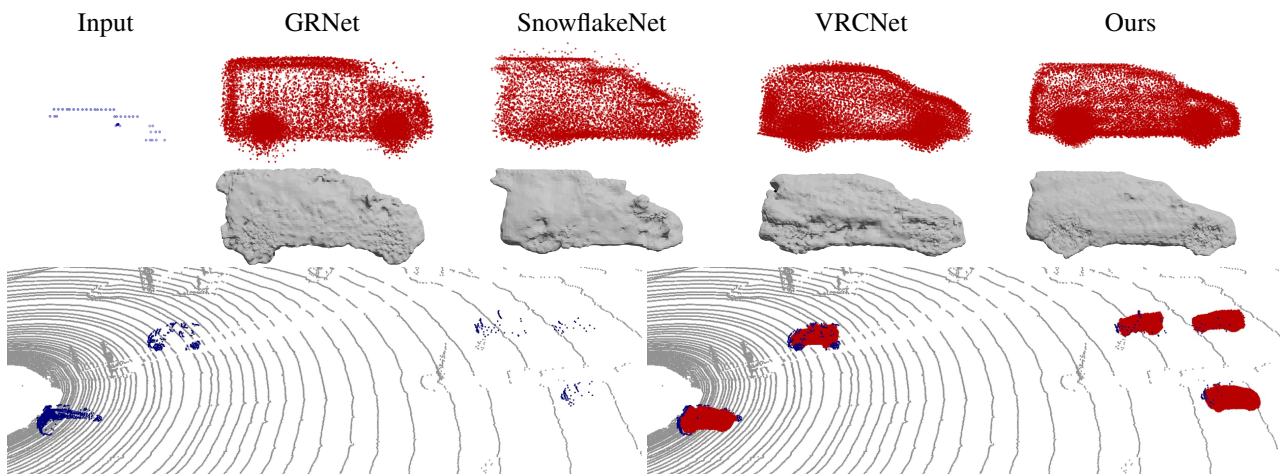


Figure 7. Qualitative point cloud completion and surface reconstruction results on the KITTI dataset.

use F-Score [14] to evaluate the distance between objects as in [21]. As shown in Table 2, our method outperforms all the other methods in terms of both CD and F-Score@1%.

4.3.3 Results on KITTI

Since there are no complete ground truth point clouds for KITTI, we follow the protocol of ASHF-Net [51] to evaluate the performance, where Fidelity, Minimal Matching Distance (MMD), Consistency and Uniformity are used as evaluation metrics. Table 3 shows the advantages of the point clouds generated by our method over the other methods in terms of distribution uniformity, which demonstrates that the proposed IFNet is beneficial to produce uniformly distributed point cloud through the iterative self-correction procedure. Besides, our approach achieves the lowest MMD among all the methods. Other than the quantitative comparison, we also visualize point cloud comple-

tion results and surface reconstruction results in Figure 7, where the mesh is created by Poisson Surface Reconstruction [12]. We can observe that other methods produce more artifacts in the reconstructed surfaces due to the generation of unevenly distributed point set, while our method produce smoother surface and clearer structures in the completion results.

4.4. Ablation Study and Parameters

We demonstrate the effectiveness of the proposed networks through an ablation study, and consider the choice of parameters. We conducted all experiments on the ShapeNet dataset and all experimental settings including the network architectures were the same as described in 4.2, except for the analysis.




Figure 8. Visual comparison of coarse point clouds generated by SFM and GFV.

4.4.1 Feature structuring network

We compared the performance of point cloud completion to quantitatively validate the effect of the FSNet; the structured feature map SFM was generated with various numbers of channels h . Additionally, we replace FSNet by a max-pooling operation to produce a global feature vector (GFV) as the global representation instead of SFM, and the coarse point cloud was generated by the module in SnowflakeNet [40]. As Table 4 shows, the completion quality of the SFM-based method for both coarse and dense point clouds is better than for the GFV-based method. A visual comparison of the coarse point clouds generated by different methods is provided in Figure 8, and we see that our method provides complete shapes with clearer structures. We also note that CD decreases to the lowest value as h increases to 32, which demonstrates the effectiveness of FSNet. However, CD increases when h rises to 64, which indicates that although using SFM with more channels can represent richer latent patterns, it may lead to information redundancy.

Method	CD	
	Coarse	Dense
GFV	9.77	2.64
Ours ($h = 8$)	7.68	2.40
Ours ($h = 16$)	7.62	2.35
Ours ($h = 32$)	7.26	2.27
Ours ($h = 64$)	7.64	2.30

Table 4. Shape completion (CD loss $\times 10^4$) on the ShapeNet dataset for a GFV-based method and SFM-based method, with various numbers of channels.

We also provide a visualization of the learned SFM projection for various testing samples from the ShapeNet dataset in Figure 9. Compared to a recent GFV based method [40], the distribution of SFM projections is more uniform at instance level across different categories, demonstrating that SFM has learned sufficient local information to distinguish different shapes. Note that we use the same encoder as [40] to conduct the comparison.

Method	Plane	Cabinet	Car	Chair	Lamp	Sofa	Table	Vessel	Avg.
MB [45]	0.99	3.80	2.30	3.64	3.12	4.26	2.77	2.61	2.94
DP [17]	1.00	3.88	2.30	3.46	3.07	4.24	2.75	2.60	2.92
NS [24]	0.96	3.79	2.28	3.23	2.95	4.17	2.68	2.51	2.82
Ours ($T = 1$)	0.80	3.64	2.21	3.05	2.23	3.87	2.31	1.88	2.50
Ours ($T = 3$)	0.79	3.45	2.24	2.86	1.99	3.98	2.27	2.03	2.45
Ours ($T = 5$)	0.82	3.52	2.18	2.97	1.85	3.65	2.12	1.92	2.38
Ours ($T = 7$)	0.84	3.47	2.18	2.63	1.84	3.73	2.04	1.90	2.33
Ours ($T = 9$)	0.80	3.53	2.13	2.48	1.79	3.64	1.98	1.81	2.27

Table 5. Shape completion results (CD loss $\times 10^4$) on the ShapeNet dataset with different feature expansion networks.

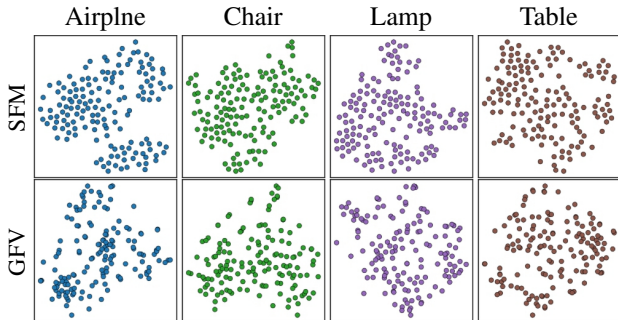


Figure 9. Compared to GFV used in [40], the feature embedding of the learned *SFM* is more uniformly distributed.

4.4.2 Iterative feedback network

Feature expansion is an important component of recent point cloud upsampling methods. However, existing feature expansion methods, including multi-branch MLPs (MB) [45], duplication-based expansion (DP) [17] and NodeShuffle (NS) [24] generally perform expansion in one or two steps, which limits their abilities in the point cloud completion task which has incomplete and non-uniform inputs. To demonstrate the effectiveness of IFNet, we conducted a comparison in which IFNet in our pipeline is either replaced by other feature expansion methods or has various numbers T of feedback blocks. Quantitative results are shown in Table 5, which demonstrates the advantage of IFNet over other methods in the point cloud completion task; the completion quality improves as the number of feedback blocks increases. Qualitative completion results and close-ups are illustrated in Figure 10, showing that the dense point clouds generated by our method have fewer outliers and less noise, demonstrating that the proposed IFNet facilitates consolidation and uniformity of output point clouds.

For more comprehensive understanding of the progressive refinement process in IFNet, we visualize various intermediate completion results in Figure 11, where the offsets between the coordinates of points in P^s and the dense

point cloud in step t are regressed from the corresponding dense features F_t^d . It shows that as the dense features are progressively corrected via the feedback mechanism, local details and surface uniformity of the dense point cloud are gradually improved. Note that only the dense features from the last feedback block are fed to the offset regression module during both training and test stages, and there is no need to supervise the intermediate results as in previous methods [40, 11].

5. Conclusions

In this study, we revisit the problem of point cloud completion and propose two novel networks which can be integrated into a coarse-to-fine pipeline. By replacing the max pooling operation with FSNet, we are able to efficiently aggregate both global and local information from partial observations. Moreover, we introduce IFNet into the upsampling stage, which can work in a self-correcting manner, helping to progressively refine local details of the output shape. Experiments on multiple datasets indicate that our method achieves state-of-the-art performance.

However, our method still has certain limitations. FSNet incurs an extra computational cost compared to max-pooling, and it may lead to information redundancy as noted in the parameter study. Furthermore, IFNet takes multiple steps to expand the sparse features, which is also relatively time-consuming. It is worth exploring ways to address these limitations and the possibility of applying our method to other tasks like point cloud consolidation and upsampling. Moreover, since we use multi-head attention in FSNet and self-attention in IFNet, it would be interesting to explore other more sophisticated attention mechanisms [8], e.g. transformer-based modules that have proved helpful in 3D point cloud processing [7, 50], to see if the results can be further improved.

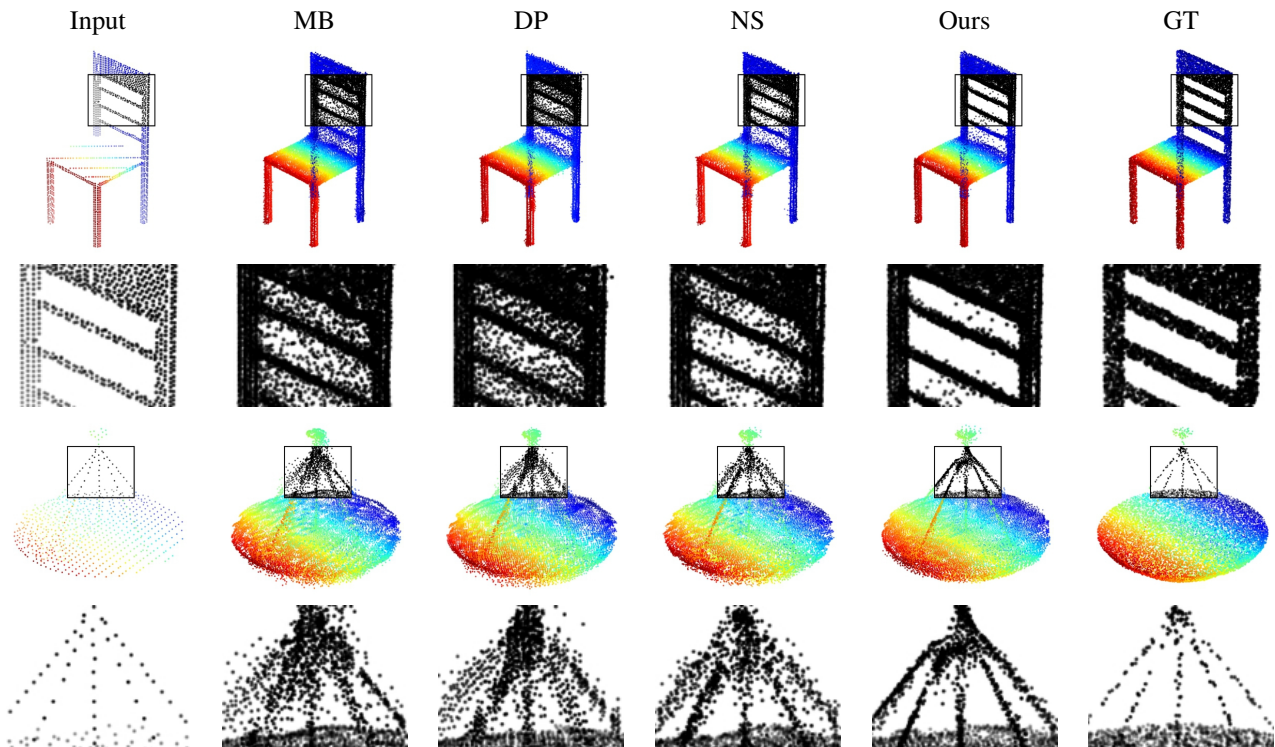


Figure 10. Example shape completion results using the ShapeNet dataset using different feature expansion networks.




Figure 11. Example shape completion results using the ShapeNet dataset regressed from the intermediate features in IFNet.

References

- [1] A. Alliegro, D. Valsesia, G. Fracastoro, E. Magli, and T. Tommasi. Denoise and contrast for category agnostic shape completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4629–4638, 2021. 3
- [2] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1534–1543, 2016. 1
- [3] M. Berger, A. Tagliasacchi, L. Seversky, P. Alliez, J. Levine, A. Sharf, and C. Silva. State of the art in surface reconstruction from point clouds. In *Eurographics 2014-State of the Art Reports*, volume 1(1), pages 161–185, 2014. 1, 2
- [4] A. Dai, C. R. Qi, and M. 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. 1, 2
- [5] H. Fan, H. Su, and L. J. Guibas. A point set generation net-

- work for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. 6
- [6] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 6
- [7] M.-H. Guo, J.-X. Cai, Z.-N. Liu, T.-J. Mu, R. R. Martin, and S.-M. Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7(2):187–199, 2021. 10
- [8] M.-H. Guo, T.-X. Xu, J.-J. Liu, Z.-N. Liu, P.-T. Jiang, T.-J. Mu, S.-H. Zhang, R. R. Martin, M.-M. Cheng, and S.-M. Hu. Attention mechanisms in computer vision: A survey. *arXiv preprint arXiv:2111.07624*, 2021. 10
- [9] M. Haris, G. Shakhnarovich, and N. Ukita. Deep back-projection networks for super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1664–1673, 2018. 2, 5
- [10] T. Hu, Z. Han, A. Shrivastava, and M. Zwicker. Render4completion: Synthesizing multi-view depth maps for 3d shape completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3
- [11] Z. Huang, Y. Yu, J. Xu, F. Ni, and X. Le. Pf-net: Point fractal network for 3d point cloud completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7662–7670, 2020. 3, 7, 8, 10
- [12] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)*, 32(3):1–13, 2013. 8
- [13] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [14] A. Knapitsch, J. Park, Q.-Y. Zhou, and V. Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017. 8
- [15] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726–1743, 2013. 1
- [16] J. Lee, Y. Lee, J. Kim, A. Kosiorok, S. Choi, and Y. W. Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International Conference on Machine Learning*, pages 3744–3753. PMLR, 2019. 4
- [17] R. Li, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng. Pagan: a point cloud upsampling adversarial network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7203–7212, 2019. 3, 10
- [18] M. Liu, L. Sheng, S. Yang, J. Shao, and S.-M. Hu. Morphing and sampling network for dense point cloud completion. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34(07), pages 11596–11603, 2020. 3, 7, 8
- [19] N. J. Mitra, M. Pauly, M. Wand, and D. Ceylan. Symmetry in 3d geometry: Extraction and applications. In *Computer Graphics Forum*, volume 32(6), pages 1–23. Wiley Online Library, 2013. 1, 2
- [20] L. Pan. Ecg: Edge-aware point cloud completion with graph convolution. *IEEE Robotics and Automation Letters*, 5(3):4392–4398, 2020. 7
- [21] L. Pan, X. Chen, Z. Cai, J. Zhang, H. Zhao, S. Yi, and Z. Liu. Variational relational point completion network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8524–8533, 2021. 1, 3, 6, 7, 8
- [22] C. R. Qi, H. Su, K. Mo, and L. 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*, pages 652–660, 2017. 1, 2
- [23] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. 1, 3
- [24] G. Qian, A. Abualshour, G. Li, A. Thabet, and B. Ghanem. Pu-gcn: Point cloud upsampling using graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11683–11692, 2021. 3, 10
- [25] Y. Qian, J. Hou, S. Kwong, and Y. He. Pugeo-net: A geometry-centric network for 3d point cloud upsampling. In *European Conference on Computer Vision*, pages 752–769. Springer, 2020. 3
- [26] M. Sarmad, H. J. Lee, and Y. M. Kim. Rl-gan-net: A reinforcement learning agent controlled gan network for real-time point cloud shape completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5898–5907, 2019. 3
- [27] D. Stutz and A. 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. 1, 2
- [28] M. Tarini, H. P. Lensch, M. Goesele, and H.-P. Seidel. 3d acquisition of mirroring objects using striped patterns. *Graphical Models*, 67(4):233–259, 2005. 1
- [29] L. P. Tchappmi, V. Kosaraju, H. Rezatofghi, I. Reid, and S. Savarese. Topnet: Structural point cloud decoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 383–392, 2019. 1, 3, 6, 7, 8
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2, 4
- [31] X. Wang, M. H. Ang, and G. H. Lee. Voxel-based network for shape completion by leveraging edge generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13189–13198, 2021. 3
- [32] X. Wang, M. H. Ang Jr, and G. H. Lee. Cascaded refinement network for point cloud completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 790–799, 2020. 3, 7, 8
- [33] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019. 5, 7
- [34] Y. Wang, D. J. Tan, N. Navab, and F. Tombari. Softpool-net: Shape descriptor for point cloud completion and classification. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 70–85. Springer, 2020. 3

- [35] X. Wen, Z. Han, Y.-P. Cao, P. Wan, W. Zheng, and Y.-S. Liu. Cycle4completion: Unpaired point cloud completion using cycle transformation with missing region coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13080–13089, 2021. [3](#)
- [36] X. Wen, T. Li, Z. Han, and Y.-S. Liu. Point cloud completion by skip-attention network with hierarchical folding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1939–1948, 2020. [1](#), [3](#)
- [37] X. Wen, P. Xiang, Z. Han, Y.-P. Cao, P. Wan, W. Zheng, and Y.-S. Liu. Pmp-net: Point cloud completion by learning multi-step point moving paths. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7443–7452, 2021. [3](#), [7](#)
- [38] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. 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. [6](#)
- [39] Y. Xia, Y. Xia, W. Li, R. Song, K. Cao, and U. Stilla. Asfm-net: Asymmetrical siamese feature matching network for point completion. *arXiv preprint arXiv:2104.09587*, 2021. [3](#)
- [40] P. Xiang, X. Wen, Y.-S. Liu, Y.-P. Cao, P. Wan, W. Zheng, and Z. Han. Snowflakenet: Point cloud completion by snowflake point deconvolution with skip-transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5499–5509, 2021. [3](#), [7](#), [8](#), [9](#), [10](#)
- [41] C. Xie, C. Wang, B. Zhang, H. Yang, D. Chen, and F. Wen. Style-based point generator with adversarial rendering for point cloud completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4619–4628, 2021. [3](#)
- [42] H. Xie, H. Yao, S. Zhou, J. Mao, S. Zhang, and W. Sun. Gr-net: Gridding residual network for dense point cloud completion. In *European Conference on Computer Vision*, pages 365–381. Springer, 2020. [1](#), [3](#), [7](#), [8](#)
- [43] W. Yifan, S. Wu, H. Huang, D. Cohen-Or, and O. Sorkine-Hornung. Patch-based progressive 3d point set upsampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5958–5967, 2019. [3](#)
- [44] L. Yu, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng. Ecn-net: an edge-aware point set consolidation network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 386–402, 2018. [3](#)
- [45] L. Yu, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng. Punct-net: Point cloud upsampling network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2790–2799, 2018. [3](#), [10](#)
- [46] X. Yu, Y. Rao, Z. Wang, Z. Liu, J. Lu, and J. Zhou. PointR: Diverse point cloud completion with geometry-aware transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12498–12507, 2021. [3](#), [7](#)
- [47] W. Yuan, T. Khot, D. Held, C. Mertz, and M. Hebert. Pcn: Point completion network. In *2018 International Conference on 3D Vision (3DV)*, pages 728–737. IEEE, 2018. [1](#), [3](#), [6](#), [7](#), [8](#)
- [48] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena. Self-attention generative adversarial networks. In *International conference on machine learning*, pages 7354–7363. PMLR, 2019. [5](#)
- [49] W. Zhang, Q. Yan, and C. Xiao. Detail preserved point cloud completion via separated feature aggregation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 512–528. Springer, 2020. [1](#), [3](#), [7](#), [8](#)
- [50] H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16259–16268, 2021. [10](#)
- [51] D. Zong, S. Sun, and J. Zhao. Ashf-net: Adaptive sampling and hierarchical folding network for robust point cloud completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35(4), pages 3625–3632, 2021. [3](#), [7](#), [8](#)