

Preserving Topological and Geometric Embeddings for Point Cloud Recovery

Kaiyue Zhou^{1,2}, Zelong Tan^{1,2}, Hongxiao Wang³, Ya-Li Li^{1,2}, Shengjin Wang^{1,2*}

¹Department of Electronic Engineering, Tsinghua University

²Beijing National Research Center for Information Science and Technology (BNRist)

³Academy for Multidisciplinary Studies, Capital Normal University

Abstract

Recovering point clouds involves the sequential process of sampling and restoration, yet existing methods struggle to effectively leverage both topological and geometric attributes. To address this, we propose an end-to-end architecture named **TopGeoFormer**, which maintains these critical properties throughout the sampling and restoration phases. First, we revisit traditional feature extraction techniques to yield topological embedding using a continuous mapping of relative relationships between neighboring points, and integrate it in both phases for preserving the structure of the original space. Second, we propose the **InterTwinning Attention** to fully merge topological and geometric embeddings, which queries shape with local awareness in both phases to form a learnable 3D shape context facilitated with point-wise, point-shape-wise, and intra-shape features. Third, we introduce a full geometry loss and a topological constraint loss to optimize the embeddings in both Euclidean and topological spaces. The geometry loss uses inconsistent matching between coarse-to-fine generations and targets for reconstructing better geometric details, and the constraint loss limits embedding variances for better approximation of the topological space. In experiments, we comprehensively analyze the circumstances using the conventional and learning-based sampling/upsampling/recovery algorithms. The quantitative and qualitative results demonstrate that our method significantly outperforms existing sampling and recovery methods.

Code — <https://github.com/ky-zhou/TopGeoFormer>

1 Introduction

Current 3D scanning technologies are essential for capturing detailed representations of objects, environments, or even entire landscapes in the form of dense point clouds. They generally achieve high accuracy and resolution, especially for small objects. Here, we define **recovery** as the process of both downsampling and subsequent restoration of a point cloud. To integrate data from various remote devices efficiently, point cloud recovery becomes a fundamental methodology in various applications, ranging from remote robot maneuver, augmented reality interaction, to medical surgery interaction. As an emerging field, it is different from point cloud completion, because the focus shifts

*Corresponding author: wsgsj@tsinghua.edu.cn
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

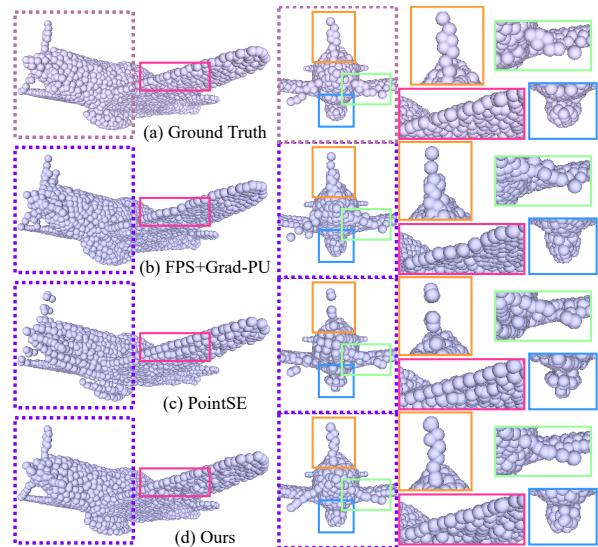


Figure 1: Overview of an aircraft and front views of recovered details. More details such as continuity and curvature (propeller, fuel tank, and wing) are preserved by our method.

from dealing with imperfection to efficiently sampling and restoring the point cloud while minimizing information loss. Moreover, point cloud completion requires transmitting the entire input directly to the decoding phase, which inherently conflicts with the purpose defined in recovery. In contrast to compression (Schwarz et al. 2018) focusing solely on transmission efficiency, recovery aims to jointly optimize both geometric fidelity and transmission efficiency.

The quality of recovered points depends on the procedure of sampling and restoration. An intuitive paradigm generally utilizes a traditional downsampling algorithm, farthest point sampling (FPS), cascaded by a followed upsampling method for this task. However, such a two-stage pipeline may lead to the loss of fine details due to the nature of FPS, which aims to uniformly sample points on a surface. On the contrary, the sampling process is essential to the overall reconstruction accuracy, since point cloud sampling must maintain the integrity of spatial relationships between points while balancing geometric fidelity with data volume. This is more challenging than image downsampling techniques due to the

orderless nature of point clouds.

To compensate for this tradeoff, existing point cloud recovery methods tend to learn displacements based on FPS-sampled points to capture the geometric pattern. Point Set Self-Embedding (Li et al. 2022) (PointSE) was first proposed to explore this domain, using an end-to-end paradigm for sampling and restoring objects represented by point clouds. However, it directly formulates the topological features using many k-nearest-neighbors (kNN) operations without fully leveraging that rich information with the shape, resulting extremely more expensive computation and less effective results. Meanwhile, the Chamfer distance (CD) loss is applied on recovered point clouds with their ground truth counterparts in a consistent manner, producing predictions with fewer local geometric details.

In this paper, we focus on the extremely under-explored research field, i.e., point cloud recovery, by proposing *TopGeoFormer* for preserving both topological and geometric attributes. First, we formulate a Down-Preservation module generating a 3D shape context to capture intra-shape features. Instead of directly mapping the 1D global code to the Euclidean space, we devise the InterTwining Attention (ITA) that queries a general learnable 2D shape code with point-wise features guided by topological embeddings to contextually encode local geometric patterns. Second, we propose an Up-Preservation phase, which integrates ITA, Up-Preserving Attention (UPA), and residual multi-layer perceptrons (MLPs), aiming to restore dense points. Third, to learn explicit local details, we propose to optimize the network based on the inconsistent mappings between the restored point clouds at various resolutions and the ground truth. Comparing to consistently matching the ground truth resolution with generations, this strategy effectively refines the local details of the coarse point clouds produced by intermediate layers. Lastly, we propose a topological constraint loss to guide the optimization process, which limits embedding variance in the latent space and slightly enhances recovery performance. By such designs, we achieve superior performance compared to state-of-the-arts by a large margin in terms of both sampling and recovery tasks. Qualitative results indicate that sampled and restored objects exhibit significantly more original and realistic local geometry. We summarize our contributions as follows:

- We revisit traditional feature extraction to fully explore the utility of topological embedding, which is subtraction relations between a local centroid and its neighbors. Such embedding is essential to guide intra-shape features for precise sampling and restoration.
- To merge topological and geometric embeddings, we propose intertwining attention by interacting local and point-wise features with the rough global shape code, which can be transformed to intra-shape features to further enrich the current shape information for closely resembling the original structure.
- Our losses are specifically designed for the proposed paradigm with respect to both Euclidean and topological spaces for preservation of fine-details. The full geometry loss always utilizes ground truth to guide coarse-to-fine

predictions, and the topological constraint loss ensures the encoded embeddings to remain coherent with their original local structures.

2 Related Work

Point Cloud Sampling. Point cloud sampling (downsampling) methods (Li, Chen, and Lee 2018; Dovrat, Lang, and Avidan 2019; Lang, Manor, and Avidan 2020; Wu et al. 2023) are generally integrated with understanding and reconstruction downstreams. Our method can be categorized as reconstruction-oriented sampling that preserves shape details without any semantic priors (Liu et al. 2024a).

Point Cloud Upsampling. Point cloud upsampling falls into three main types: shuffle-based (Yu et al. 2018; Qian et al. 2021; Luo et al. 2021; Ye et al. 2021; Qiu, Anwar, and Barnes 2022), tile-based (Wang et al. 2019b; Li et al. 2019, 2021; Rong et al. 2024), and others based on implicit functions (Ma et al. 2021; Feng et al. 2022; Zhao et al. 2022; He et al. 2023), diffusion models (Qu et al. 2024), or other modalities (Qian et al. 2020; Zhao, Hui, and Xie 2021; Liu et al. 2024b; Yang et al. 2024). This task generally restores uniform points and fills small holes, which differentiates from recovery that focuses more on restoring details.

Point Cloud Completion. Point cloud completion is analogous to recovery by merging restoration and upsampling into an individual network. However, serving as a rough guidance, the partial input is inappropriate for recovery where the input that is the entire point cloud should remain unknown during recovery. Without other priors rather than coordinates, PCN (Yuan et al. 2018) was first proposed to enlighten many classic methods (Tchapmi et al. 2019; Wen et al. 2021; Xiang et al. 2021; Cui et al. 2023). Recent methods are developed using transformers (Yu et al. 2021; Zhou et al. 2022; Li et al. 2023; Chen et al. 2023; Zhu et al. 2023) or semantic information (Xu et al. 2023a; Xiang et al. 2023; Xia et al. 2023; Xu et al. 2023b; Zhou et al. 2024).

Point Cloud Recovery. Pioneering the realm of recovery, PointSE presents an end-to-end network to sample and restore point clouds simultaneously (Li et al. 2022). As the only method in the community, it proposes that such a paradigm is well suited for long-range communication scenarios, where processing or transmitting the entire point cloud can be impractical due to data volume constraints.

3 Methodology of TopGeoFormer

In this work, we address the task of point cloud recovery while preserving the target’s topological and geometric properties. As the overview shown in Fig. 2, our network has an end-to-end architecture, comprising a down-preservation (DP) phase and two up-preservation (UP) phases. Given a point set $P = \{p_i | i = 1, \dots, N\} \in \mathbb{R}^{N \times 3}$, we aim to learn a mapping function f that transforms $p \in P$ continuously into a sampled point cloud R_0 representing the coarse shape that closely resembles the original shape P , and then into a restored point cloud R . Such transformation process can preserve the topological and geometric properties if the mapping function f is continuous and the local Euclidean neighborhoods exhibit sufficient overlap across

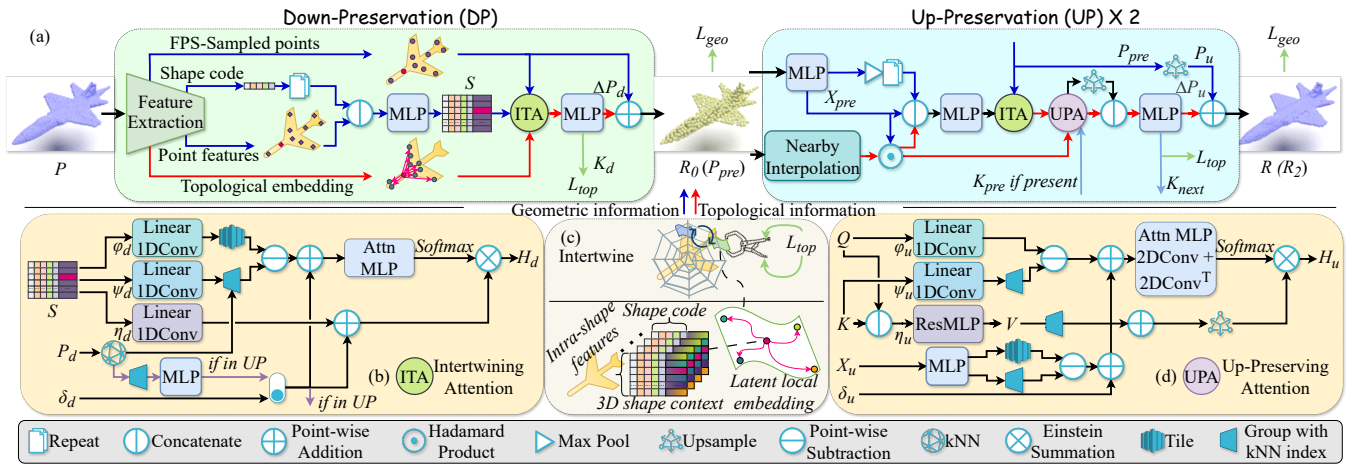


Figure 2: Architecture of our *TopGeoFormer* (a) consists of a Down-Preservation (DP) phase and two cascaded Up-Preservation (UP) phases, to preserve both topological and geometric attributes. (b) The learnable 2D shape code S is (c) intertwined with latent embeddings to reflect tangent-space structures. (d) Assisted with ITA, UPA preserves similar features for restoration.

submanifolds. Therefore, we aim to improve recovery capacity in terms of shape details based on **topological** and **geometric** attributes. We name our approach as the Topology-Geometry Preserving Transformer (*TopGeoFormer*), even though it is inherently a local method (Ma and Fu 2012).

3.1 Down-Preservation

Accurate downsampling plays a crucial role in precise object recovery, as the shape code s (typically encoded by PointNet-like (Qi et al. 2017a) structures (Qi et al. 2017b; Wang et al. 2019a)) provides essential guidance for determining a general destination where a seed should move to. Generally, s encodes partial topological and geometric information of the manifold M . However, directly mapping a shape code s to formulate seeds may introduce additional noise (Xiang et al. 2021). Moreover, simply concatenating s with the non-differentially sampled points still limits the comprehensive awareness on local details, since gradients cannot flow back fluently to update local features.

Extracting Topological and Geometric Embeddings. To preserve more local information, instead of only mapping the shape code to a coarse point cloud, we revisit PointNet++ (Qi et al. 2017b) and DGCNN (Wang et al. 2019a) to additionally reuse the topological information formulated during shape code extraction. Fig. 2(a) shows the structure of our DP phase. Given P , we extract the shape code $s \in \mathbb{R}^{1 \times C_3}$ through a feature extraction module consisting of three set abstraction (SA) layers. Each SA layer is composed of FPS and local grouping based on k -nearest-neighbors (kNN), generating FPS-sampled points $P_d \in \mathbb{R}^{N_d \times 3}$ as the geometric embedding and their corresponding features $X_d \in \mathbb{R}^{N_d \times C_1}$ as

$$X_d = MLP(Cat(P_d, \delta_d)), \delta_d = MLP(\mathcal{D}), \quad (1)$$

where $\mathcal{N}_P(p_i)$ denotes the k neighbors of $p_i \in P$, the subtraction operation $\mathcal{D}_i = \{p_i - p_j | p_j \in \mathcal{N}_P(p_i)\}$ in Euclidean space yields a vector that represents both the magnitude and

direction between points, forming the foundation for topological embedding. The subscript d in Eq. 1 denotes the first SA layer, representing the set that non-differentially sampled points belong to. We map $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^{N_d} \in \mathbb{R}^{N_d \times k \times 3}$ to a higher-dimensional latent space $\mathbb{R}^{N_d \times k \times m}$ (where $m > 3$) using an MLP as the continuous map. The MLP can capture complex and fine-grained topological structures, such that X_d preserves the key topological embeddings (e.g., connectivity) of the original space. Finally, a max-pooling operator aggregates features in X_d along the neighborhood dimension, achieving topological invariant while capturing prominent topological features and resisting to small local perturbations.

During feature extraction, we set the downsampling ratios for the SA layers to 4, 16, and N , respectively, and retain P_d and X_d from the first SA for subsequent steps. Preserving local nearest neighbors ensures that the underlying manifold of a point cloud is respected when encoding the data into a different (latent) space during dimensionality reduction. Therefore, we concatenate X_d with s and feed it into a residual MLP (*ResMLP*) to formulate a 2D point-shape fusing code $S \in \mathbb{R}^{N_d \times C_1}$ that comprehensively mixes the local awareness of each point with the overall shape:

$$S = ResMLP(Cat(X_d, Repeat(s))), \quad (2)$$

where s is duplicated to fit the size of X_d . By merging shape information, S also makes the network permutation-invariant, an essential property for point cloud recovery, given that points may be transmitted asynchronously from different remote sources. Moreover, S serves as a compact latent encoding of the manifold M in the DP phase, which is designed to preserve the topological properties (such as continuity, connectedness) of the original point set P . Thus, S encapsulates the structure of the manifold M onto which the point set P or its subsets P_d are mapped. Essentially, S encodes both the geometric and topological structure of M , allowing the reconstruction of the shape in Euclidean

space via tangent space approximations that preserve the local structure of the manifold.

Intertwining Attention. In order to establish a point-wise, point-shape, and intra-shape context in the manifold, we propose to feed S into a topology-geometry intertwining attention (ITA) module for mapping the latent embeddings guided by tangent space approximations back to the Euclidean space. This process is illustrated in Fig. 2(b). In ITA, S is first transformed by three linear MLPs, denoted as φ_d , ψ_d , and η_d , respectively. Simultaneously, the topological embedding δ_d is used to guide the 2D shape code S that describes a general shape (such as a plane sketch), formulating a 3D shape context that encodes local patterns intertwined with a group of 2D shape codes. The benefit of this operation is twofold: (i) it reduces redundant neighbor computations; (ii) it preserves exhaustive local pattern with respect to P . Recall that X_d merely represents a rough measure of connectivity due to dimensionality reduction. Here, we aim to learn a more comprehensive manifold by interacting all intra-shape features (representing a same shape) with local embeddings δ_d to enrich the shape information. The attention vector H_d is formulated as:

$$h_i = \sum_{p_j \in \mathcal{N}(p_i)} \rho(\gamma(\varphi_d(S) - \psi_d(S) + \delta_d)) \otimes (\eta_d(S) + \delta_d), \quad (3)$$

$$H_d = \{h_i\}_{i=1}^{N_d} \in \mathbb{R}^{N_d \times C},$$

where N_d is the number of downsampled points, ρ is the *Softmax* activation function, γ is an MLP layer with two linear layers and a *ReLU* nonlinear activation function, C is the dimension of feature, and all transformations of S are broadcasted or grouped to match the dimension of $\delta_d \in \mathbb{R}^{N_d \times k \times C}$. Finally, we feed these intertwining attention vectors H_d into another *ResMLP* to regress the displacement vectors $\Delta P_d \in \mathbb{R}^{N_d \times 3}$:

$$\Delta P_d = \tanh(K_d), K_d = \text{ResMLP}(H_d), \quad (4)$$

which are added back to the non-differentially sampled points P_d , eventually endowing differentiability to these points as $R_0 \in \mathbb{R}^{N_d \times 3}$:

$$R_0 = P_d + \Delta P_d. \quad (5)$$

Our sampling (DP) is designed to locally preserve both topological and geometric embeddings. Since the features of all (sampled) points are differentiable, the encoded features comprehensively reflect the topological structure of a point cloud throughout sampling and restoration phases.

3.2 Up-Preservation

To generate geometric details while precisely preserving topological properties, we propose to progressively restore the point cloud in a coarse-to-fine manner. We devise this preservation procedure to be two cascaded UP phases. For simplicity, we primarily describe an individual UP in Fig. 2 in this section. We denote the input to the UP phase as P_{pre} for general expression, where $P_{pre} = R_0$ in the first UP phase. Inspired by decoders using transposed convolution (Xiang et al. 2021; Zhou et al. 2022), we adopt their self-attention layer to predict displacements for upsampled

points. Furthermore, we update the upsampling features by interpolating the generated points, and cooperate the proposed ITA as a plug-in component with the Up-Preserving Attention (UPA) in each UP module.

Interpolating features between sparsely distributed seeds and generated points (Zhou et al. 2022) is impractical, since seed features are prohibited from transmitting to the restoration stage. Instead, we dynamically update the points to represent the trend for $p_i \in P_{pre}$ to be upsampled in a small local area. We thus calculate the weights to adjust point-wise features as the geometric embedding X_{pre} , which formulates the upsampling feature with topological embedding as:

$$X_u = \frac{\mathbf{1}^\top (\mathcal{D}^* \odot X_{pre})}{\mathbf{1}^\top \mathcal{D}^*}, \mathbf{1}^\top \mathcal{D}^* = \sum_{j \in \mathcal{N}_{P_{pre}}(p_i)} \frac{1}{p_i - p_j}, \quad (6)$$

where $|\mathcal{N}_{P_{pre}}(p_i)| = 3$ and $p_i \in P_{pre}$. Parallely, P_{pre} is encoded by a point-wise MLP to form X_{pre} . Both embeddings are then used to enrich the upsampling information in the cascaded intertwining and up-preserving attentions.

We first use ITA to enhance the query $Q = \text{ResMLP}(\text{Cat}(X_u, \text{Max}(X_{pre}), X_{pre}))$ for the following UPA, and $K = Q$ if not present in the current phase. We hereby substitute S with $\text{Max}(X_{pre})$ as the updated 2D shape code. δ_u and its corresponding kNN indices are calculated, since any kNN cache or local pattern is prohibited from transferring between sampling and restoration phases in the recovery task. Still, we can reuse such kNN cache in the subsequent UPA to avoid duplicated computations.

Up-Preserving Attention. We leverage UPA as illustrated in Fig. 2(d) to learn the attention vector for upsampling while preserving the shape information. Let φ_u and ψ_u denote linear 1D convolution (*1DConv*) layers for transforming Q and K , and η_u a *ResMLP* for concatenating Q and K to formulate the value vector V . Without loss of generality, notations such as Q , K , V , etc., omit the index of the current UP phase. Meanwhile, we use an MLP to map X_u to the same dimension of V . With the additional upsampling feature X_u , we then use an MLP γ_u consisting of a linear 2D layer, a *ReLU* nonlinear layer, and a 2D transposed convolution layer (*2DConv^T*) to merge topological, upsampling, and intra-shape features, formulating the attention vector H_u in each UP. Akin to Eq. 4 and 5, the displacements are predicted using H_u . In a coarse-to-fine manner, the obtained R_l can be further used as P_{pre} in the next UP phase. Finally, the restored point cloud R is obtained after two UP phases.

3.3 Optimization

Conventional coarse-to-fine methods for point cloud reconstruction (Xiang et al. 2021; Zhou et al. 2022; Zhu et al. 2023) typically leverage progressive CD losses for optimization. However, the target point cloud is matched to the same size as the predicted point cloud using FPS, leading to less information preserved during the coarse predictions. To address this, we propose to optimize the coordinates at various resolutions using the original-resolution ground truth as the target. Hence, our full reconstruction loss is formulated as:

$$\mathcal{L}_{geo} = \sum_{l=0}^2 \mathcal{L}_{CD}(R_l, P), \quad (7)$$

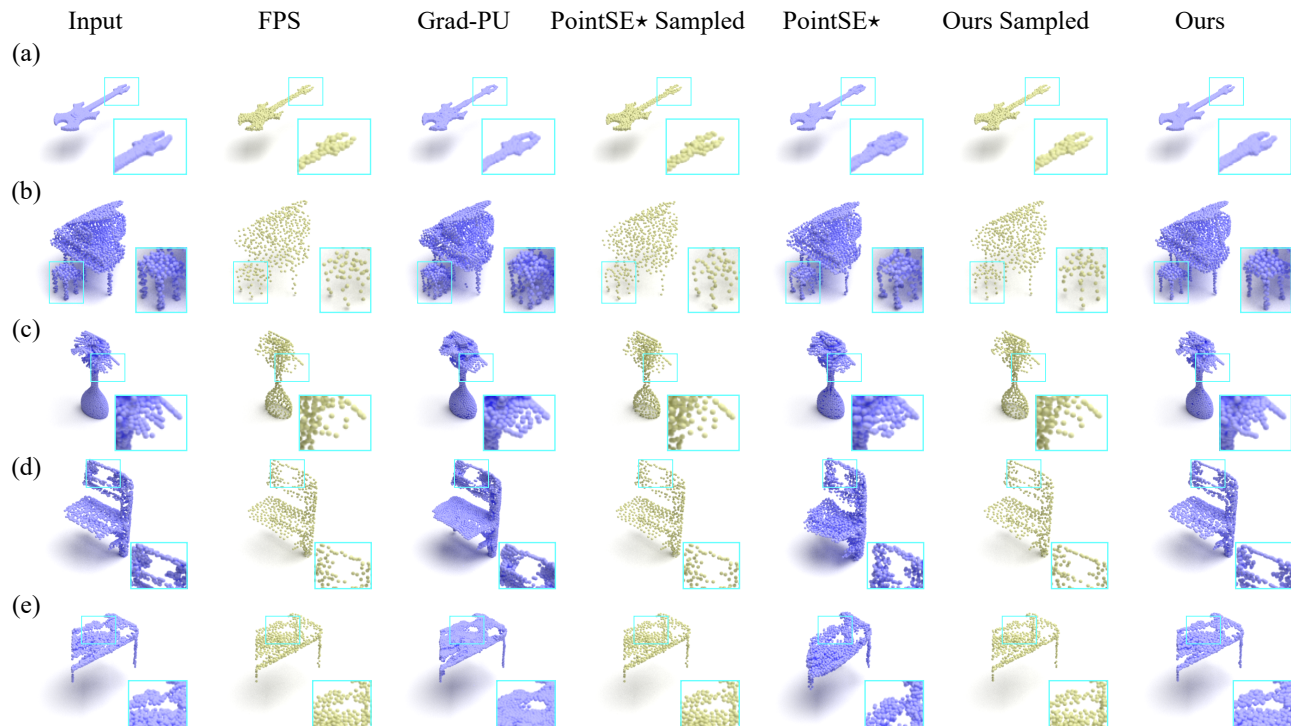


Figure 3: Qualitative visualizations. (a-c) guitar, piano, and plant in the random set. (d-e) chair and table in the partial set. Both our sampling (DP) and recovery methods preserve plausible details and overall shape.

where R_1 is generated by the first UP phase, $R_2 = R$, and \mathcal{L}_{CD} represents the Chamfer distance (CD), formulated as:

$$\mathcal{L}_{CD} = \frac{1}{|P|} \sum_{x \in P} \min_{y \in R} \|x - y\|_2 + \frac{1}{|R|} \sum_{y \in R} \min_{x \in P} \|y - x\|_2, \quad (8)$$

where $\|\cdot\|_2$ denotes the Euclidean L2-norm of CD and $|\cdot|$ is the magnitude of a matrix.

According to manifold learning (Belkin and Niyogi 2003), maintaining small variance in the encoded embeddings implies that the local linear approximations of the manifold (i.e., the tangent spaces) are well-preserved during encoding. This means that local neighborhoods in the original space, which lie on the manifold, remain coherent in the encoded space (Ma and Fu 2012). Therefore, we propose to constrain the encoded embeddings to have small variance as illustrated in Fig 2(c), such that embeddings that differ substantially are limited to overwhelm others. Formally, we denote the topological embedding constraint loss as follows:

$$\mathcal{L}_{top} = \sum_{l=0}^2 \sum_{i=1}^{N_l} \text{mean}(K_l^2), \quad (9)$$

where $N_0 = |R_0| = \frac{N}{4}$, $N_1 = |R_1| = \frac{N}{2}$, $N_2 = |R_2| = N$, $K_0 = K_d$. Eq. 9 encourages the topological embeddings to preserve local information mapped into the manifold.

Finally, the overall loss function can be formulated as:

$$\mathcal{L} = \lambda \mathcal{L}_{geo} + \mathcal{L}_{top}, \quad (10)$$

where λ is empirically set to 1000. Eq. 10 corresponds to optimizations on both geometric and topological levels, re-

distributing both sampled and recovered points with respect to preserving local shape details.

4 Experiment

Following point cloud reconstruction works (Li et al. 2022, 2019; Qian et al. 2021; Li et al. 2021; He et al. 2023), we report metrics: Chamfer distance (CD), Hausdorff distance (HD) using PyTorch3D (Ravi et al. 2020), and Earth mover’s distance (EMD) using (Liu et al. 2020). For quantitative experiments, we use: ModelNet40 (Wu et al. 2015), and ScanObjectNN (Uy et al. 2019) to generate three sets: uniform set, random set, and partial set. For generalization studies and qualitative experiments, we use the collected 10 scenes (KITTI-10) from KITTI (Geiger et al. 2013) and 30 scenes (ScanNet-30) from ScanNet (Dai et al. 2017). Our network is trained on an Nvidia RTX3090 GPU, with a batch size of 32 for 120 epochs. We use the Adam optimization with initial learning rate 0.005, which is decayed by 0.5 every 30 epochs. During training, point sets are augmented with random mirroring, scaling, and rotation, respectively. Unless otherwise stated, all experiments are conducted with our unified *PyTorch* implementations for fair comparison.

4.1 Point Cloud Recovery

Object-Level. Table 1 shows the comparisons between our method and other recovery methods, where results in the first 3 rows are directly reported from PointSE (separated by a horizontal line). As the only work in point cloud recovery, we re-implement PointSE in *PyTorch* (denoted as \star).

Method	Uniform			Random			Partial		
	CD	HD	EMD	CD	HD	EMD	CD	HD	EMD
PU-GCN	1.13	1.11	6.57	1.30	1.19	7.48	0.72	0.83	6.82
Dis-PU	0.99	1.29	6.24	1.14	1.41	7.32	0.63	0.99	6.41
PointSE	0.76	0.74	4.51	0.90	0.85	6.45	0.47	0.49	4.63
RepKPU	1.00	0.36	3.09	0.99	0.35	3.07	0.51	0.24	2.63
SampleNet	1.25	1.01	7.41	1.54	1.72	9.13	0.86	1.39	10.8
APES	0.83	0.88	3.36	0.85	0.95	3.33	0.46	1.22	2.92
PointSE*	0.85	0.35	2.93	0.86	0.34	2.97	0.38	0.15	2.57
Ours	0.63	0.29	2.82	0.62	0.29	2.82	0.26	0.18	2.39

Table 1: Quantitative comparison on recovery in terms of CD (10^3), HD (10^2), and EMD (10^2), between our methods and state-of-the-art recovery and upsampling methods.

Method	KITTI-10			ScanNet-30		
	CD	HD	EMD	CD	HD	EMD
PointSE*	1.39	0.81	1.48	1.65	0.65	1.20
Ours	0.73	0.61	1.45	1.26	0.58	1.20

Table 2: CD (10^5), HD (10^4), and EMD (10^4) comparison on randomly selected scenes from KITTI and ScanNet.

For *PyTorch*-based evaluations, we compute the metrics for RepKPU (Rong et al. 2024), SampleNet (Lang, Manor, and Avidan 2020), APES (Wu et al. 2023), PointSE*, and our method. It is worth noting that directly adapting completion methods (Xiang et al. 2021; Zhou et al. 2022; Zhu et al. 2023) into the recovery task generally performs worse than baselines in Table 1, since completion emphasizes more on the decoder with additional guidance from input. Compared to the second best, our method reduces CD on the three sets by 17%, 27%, and 32%, respectively, while also achieving superior performance in terms of almost all metrics. The upsampling methods, PU-GCN, Dis-PU, and RepKPU, use FPS-sampled points as input. The result highlights the importance of the sampling ability for shape recovery. Moreover, sampling methods, i.e., SampleNet and APES, generally perform worse than PointSE and ours, even though they are cascaded with the same decoder (Xiang et al. 2021) for adjusting from understanding-oriented into a reconstruction-oriented paradigm. Note that, the decoder from SnowflakeNet requires a global shape code for rough seed generation, while ours does not need to transmit any high-dimensional code into the decoding phase. This further strengthens the transmission’s efficiency from remote ends.

Meanwhile, we visualize the sampling and corresponding recovery results in Fig. 3. Both our sampled and recovered points better preserve the original details and overall distribution. For instance, in (a-c), our method preserves the guitar’s headstock, the chair beside piano, and the leaves of the plant, respectively; and in (d-e), our method maintains either more continuous edges or more accurate holes.

Scene-Level. On large scenes, we test the model trained on the random set using the collected scenes from KITTI-10 and ScanNet-30. We first normalize a scene and then divide it into patches with 2048 points. After recovery, the output patches are merged back together using their respective normalization parameters. Quantitative and qualitative results are demonstrated in Table 2 and Fig. 4, where our method consistently outperforms PointSE*, particularly with respect

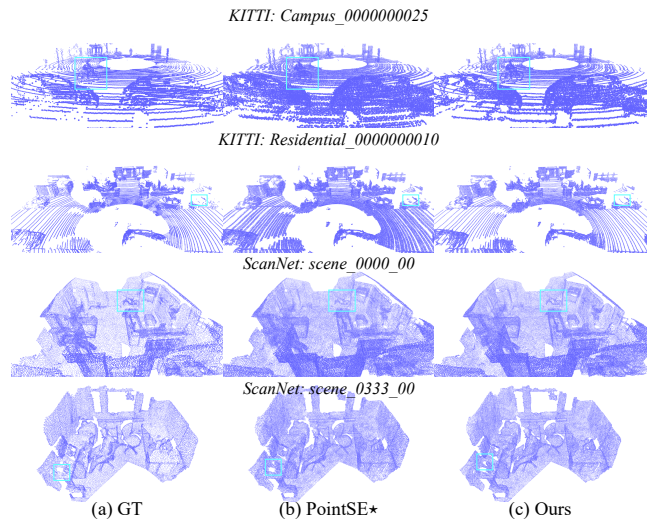


Figure 4: Recovery comparison on large scenes.

Method	4096	2048	1024	512
SampleNet	1.23/2.19	1.32/3.71	2.46/5.19	3.88/8.76
FPS	0.39/0.80	0.71/1.64	1.34/3.25	2.63/6.42
APES	0.45/1.22	0.85/2.41	1.64/4.69	3.25/9.17
Ours	0.36/0.75	0.69/1.51	1.31/2.99	2.54/5.88

Table 3: CD (10^3) errors of sampling methods at different resolutions using recovered/sampled points against GT.

to outline appearance.

4.2 Point Cloud Sampling

To demonstrate the superiority of our DP sampling module, we test on the random set of ModelNet40 at various resolutions: 4096, 2048, 1024, and 512, with a sampling ratio set to 4. For fair comparison, we use the same decoder and loss (i.e., without our \mathcal{L}_{geo} and \mathcal{L}_{top}) of SnowflakeNet for all sampling methods. Moreover, sampling-only networks (i.e., SampleNet, our DP) are trained independently with GT as target if differentiable, and otherwise we directly compute CD of sampled points and GT from recovery networks (i.e., FPS, APES). In APES, we adjust the sampling ratios of its cascaded encoding layers be to 1 and 4 (ours 4 and 16), as we find that the original point number in the first layer yields the best results. As shown in Table 3, among all sampling methods, our sampling significantly outperforms others across all resolution levels in both recovery and individual sampling tasks by 7.69%/6.25%, 2.82%/7.93%, 2.23%/8%, 3.42%/8.41%, respectively. This observation confirms not only the advantages of our designs to preserve topological and geometric properties, such as point-wise, point-shape and intra-shape correspondences, but also the significance of the sampling process in recovery. Furthermore, we visualize the sampled results at different resolutions of GT in Fig. 5, where the downsampling ratio is 4. Compared to APES and FPS, our sampling method preserves the clear track fairings on both wings of the aircraft, validating the preservation capacity of our method.

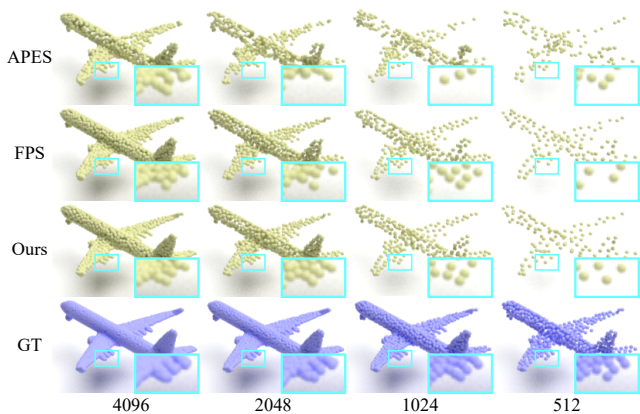


Figure 5: Sampling precisions. Track fairings are preserved.




Figure 6: Predicted displacements. Darker blue represents larger offset. Regions with more details are generally more concentrated.

Displacement Prediction. Our network aims to learn displacements to preserve fine details, and thus we visualize the magnitude of those point movements in Fig. 6. Different objects (aircraft and person) along with various instances are demonstrated to highlight consistency and symmetry in the results. For example, engines, wheels, and thin structures like wings in aircraft, as well as hands, feet, and tools in person are mostly concentrated by the method. These regions typically have smaller scales and higher geometric variance, which further demonstrates the preservation capacity of our sampling method.

4.3 Complexity Analysis

To demonstrate the efficiency of our method, Table 4 compares architecture parameters and FLOPS (G). Testing times include evaluations for individual metrics (CD, HD, and EMD). Notably, SeedFormer is only referenced for model size comparison, as it requires additional seed features during decoding. Evaluation times are indicated in parentheses. PointSE* has more feature extraction modules, consisting of multiple FPS sampling and kNN grouping operations. While they do not contain any parameters, they extremely increase the computational cost. Thus, our training time is 25% faster than PointSE*. We also include our small (S) model to show that performance is still better than PointSE* with fewer parameters. In summary, by eliminating redun-

Method	Param.	Train	Test	FLOPS	CD	HD	EMD
RepKPU	1.02M	-	19s(11s)	9G	0.99	0.35	3.07
PointSE*	0.94M	4h	30s(25s)	41G	0.86	0.34	2.97
APES	2.27M	5h	35s(29s)	48G	0.85	0.95	3.33
SeedFormer	2.27M	-	-	-	-	-	-
Ours (S)	0.46M	1.6h	20s(12s)	4G	0.66	0.29	2.78
Ours	2.03M	3h	24s(14s)	14G	0.62	0.29	2.82

Table 4: Architecture complexity comparison.

Sampling		Restoration		Loss		Metric			
P_d	δ_d	ITA	SPD	UP	\mathcal{L}_{geo}	\mathcal{L}_{top}	CD	HD	EMD
✓	-	✓	✓	-	✓	-	0.68	3.00	2.68
✓	✓	-	✓	-	✓	-	0.65	2.90	2.75
✓	✓	✓	✓	-	-	-	0.69	3.06	2.81
✓	✓	✓	✓	-	✓	-	0.64	2.91	2.77
✓	✓	✓	✓	-	✓	✓	0.63	2.91	2.83
✓	✓	✓	-	✓	✓	-	0.64	2.89	2.75
✓	✓	✓	-	✓	✓	✓	0.62	2.89	2.82

Table 5: Ablation study on the random set of ModelNet40 in terms of CD (10^3), HD (10^3), and EMD (10^2).

nant grouping calculations and leveraging the proposed ITA module (only 0.08M parameters), our model achieves the highest efficiency in both training and testing phases.

4.4 Ablation Study

We elaborate the ablation study in Table 5, where we evaluate the contributions of various components in both the sampling and recovery phases. In the sampling phase, we highlight the importance of the geometric and topological embeddings, as well as the proposed ITA. In the restoration phase, we compare the performance of different decoding modules, including SPD (Xiang et al. 2021) and our proposed UP. Note that, neither seed features (Zhou et al. 2022) nor global code (Xiang et al. 2021) is transmitted to our UP-based restoration phase. At last, we assess the significance of the proposed losses, \mathcal{L}_{geo} and \mathcal{L}_{top} . Specifically, the result of 2048 points without our proposed losses is also provided in Table 3. With these designs, our overall architecture achieves the lowest CD and HD but slightly higher EMD. This is because our predicted displacement prioritizes preserving shape and boundaries over enforcing uniformity, which is empirically a crucial factor for EMD.

5 Conclusion and Discussion

Unlike existing sampling and upsampling methods that focus on individual tasks, we propose an end-to-end network, *TopGeoFormer*, for point cloud recovery. By encoding both topological and geometric attributes, our method achieves outstanding performance in both recovery and sampling tasks. Extensive experiments demonstrate that our method deliberately rearranges points for precise shape coverage, leading the state-of-the-art. Our model can be deployed in remote systems for energy-efficient processing and environmental monitoring. We expect this study will serve as a foundational contribution to future recovery research.

Limitation: *TopGeoFormer* is designed to favor non-uniform inputs and assumes none or low levels of noise. Moreover, since it is trained solely on complete objects, it is not suited for predicting largely missing regions.

Acknowledgments

This work is supported by the research fund under Grant No. 20242910035 from the Tsinghua University-Jiangsu CRRC Digital Technology Co.,Ltd. Joint Research Center for Data Driven Intelligence of Industry.

References

- Belkin, M.; and Niyogi, P. 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6): 1373–1396.
- Chen, Z.; Long, F.; Qiu, Z.; Yao, T.; Zhou, W.; Luo, J.; and Mei, T. 2023. Anchorformer: Point cloud completion from discriminative nodes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13581–13590.
- Cui, R.; Qiu, S.; Anwar, S.; Liu, J.; Xing, C.; Zhang, J.; and Barnes, N. 2023. P2c: Self-supervised point cloud completion from single partial clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14351–14360.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE.
- Dovrat, O.; Lang, I.; and Avidan, S. 2019. Learning to Sample. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2760–2769.
- Feng, W.; Li, J.; Cai, H.; Luo, X.; and Zhang, J. 2022. Neural points: Point cloud representation with neural fields for arbitrary upsampling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18633–18642.
- Geiger, A.; Lenz, P.; Stiller, C.; and Urtasun, R. 2013. Vision meets Robotics: The KITTI Dataset. *International Journal of Robotics Research (IJRR)*.
- He, Y.; Tang, D.; Zhang, Y.; Xue, X.; and Fu, Y. 2023. Gradpu: Arbitrary-scale point cloud upsampling via gradient descent with learned distance functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5354–5363.
- Lang, I.; Manor, A.; and Avidan, S. 2020. Samplenet: Differentiable point cloud sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7578–7588.
- Li, J.; Chen, B. M.; and Lee, G. H. 2018. So-net: Self-organizing network for point cloud analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 9397–9406.
- Li, R.; Li, X.; Fu, C.-W.; Cohen-Or, D.; and Heng, P.-A. 2019. Pu-gan: a point cloud upsampling adversarial network. In *Proceedings of the IEEE/CVF international conference on computer vision*, 7203–7212.
- Li, R.; Li, X.; Heng, P.-A.; and Fu, C.-W. 2021. Point cloud upsampling via disentangled refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 344–353.
- Li, R.; Li, X.; Wong, T.-T.; and Fu, C.-W. 2022. Point Set Self-Embedding. *IEEE Transactions on Visualization and Computer Graphics*.
- Li, S.; Gao, P.; Tan, X.; and Wei, M. 2023. Proxy-former: Proxy alignment assisted point cloud completion with missing part sensitive transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9466–9475.
- Liu, J.; Li, J.; Wang, K.; Guo, H.; Yang, J.; Peng, J.; Xu, K.; Liu, X.; and Guo, J. 2024a. LTA-PCS: learnable task-agnostic point cloud sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28035–28045.
- Liu, M.; Sheng, L.; Yang, S.; Shao, J.; and Hu, S.-M. 2020. Morphing and sampling network for dense point cloud completion. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 11596–11603.
- Liu, Y.; Chen, R.; Li, Y.; Li, Y.; and Tan, X. 2024b. SPU-PMD: Self-Supervised Point Cloud Upsampling via Progressive Mesh Deformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5188–5197.
- Luo, L.; Tang, L.; Zhou, W.; Wang, S.; and Yang, Z.-X. 2021. Pu-eva: An edge-vector based approximation solution for flexible-scale point cloud upsampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16208–16217.
- Ma, B.; Han, Z.; Liu, Y.-S.; and Zwicker, M. 2021. Neural-Pull: Learning Signed Distance Functions from Point Clouds by Learning to Pull Space onto Surfaces. In *International Conference on Machine Learning (ICML)*.
- Ma, Y.; and Fu, Y. 2012. *Manifold learning theory and applications*, volume 434. CRC press Boca Raton.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in neural information processing systems*, 5099–5108.
- Qian, G.; Abualshour, A.; Li, G.; Thabet, A.; and Ghanem, B. 2021. Pu-gcn: Point cloud upsampling using graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11683–11692.
- Qian, Y.; Hou, J.; Kwong, S.; and He, Y. 2020. PUGeoNet: A geometry-centric network for 3D point cloud upsampling. In *European conference on computer vision*, 752–769. Springer.
- Qiu, S.; Anwar, S.; and Barnes, N. 2022. Pu-transformer: Point cloud upsampling transformer. In *Proceedings of the Asian conference on computer vision*, 2475–2493.
- Qu, W.; Shao, Y.; Meng, L.; Huang, X.; and Xiao, L. 2024. A conditional denoising diffusion probabilistic model for point

- cloud upsampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20786–20795.
- Ravi, N.; Reizenstein, J.; Novotny, D.; Gordon, T.; Lo, W.-Y.; Johnson, J.; and Gkioxari, G. 2020. Accelerating 3D Deep Learning with PyTorch3D. *arXiv:2007.08501*.
- Rong, Y.; Zhou, H.; Xia, K.; Mei, C.; Wang, J.; and Lu, T. 2024. RepKPU: Point cloud upsampling with kernel point representation and deformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21050–21060.
- Schwarz, S.; Preda, M.; Baroncini, V.; Budagavi, M.; Cesar, P.; Chou, P. A.; Cohen, R. A.; Krivokuća, M.; Lasserre, S.; Li, Z.; et al. 2018. Emerging MPEG standards for point cloud compression. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 9(1): 133–148.
- Tchapmi, L. P.; Kosaraju, V.; Rezatofighi, H.; Reid, I.; and Savarese, S. 2019. Topnet: Structural point cloud decoder. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 383–392.
- Uy, M. A.; Pham, Q.-H.; Hua, B.-S.; Nguyen, T.; and Yeung, S.-K. 2019. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1588–1597.
- Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S. E.; Bronstein, M. M.; and Solomon, J. M. 2019a. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (tog)*, 38(5): 1–12.
- Wang, Y.; Wu, S.; Huang, H.; Cohen-Or, D.; and Sorkine-Hornung, O. 2019b. Patch-based progressive 3d point set upsampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5958–5967.
- Wen, X.; Xiang, P.; Han, Z.; Cao, Y.-P.; Wan, P.; Zheng, W.; and Liu, Y.-S. 2021. 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*, 7443–7452.
- Wu, C.; Zheng, J.; Pfrommer, J.; and Beyerer, J. 2023. Attention-based point cloud edge sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5333–5343.
- Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1912–1920.
- Xia, Z.; Liu, Y.; Li, X.; Zhu, X.; Ma, Y.; Li, Y.; Hou, Y.; and Qiao, Y. 2023. Scpnet: Semantic scene completion on point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17642–17651.
- Xiang, P.; Wen, X.; Liu, Y.-S.; Cao, Y.-P.; Wan, P.; Zheng, W.; and Han, Z. 2021. Snowflakenet: Point cloud completion by snowflake point deconvolution with skip-transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5499–5509.
- Xiang, P.; Wen, X.; Liu, Y.-S.; Zhang, H.; Fang, Y.; and Han, Z. 2023. Retro-fpn: Retrospective feature pyramid network for point cloud semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 17826–17838.
- Xu, J.; Li, X.; Tang, Y.; Yu, Q.; Hao, Y.; Hu, L.; and Chen, M. 2023a. Casfusionnet: A cascaded network for point cloud semantic scene completion by dense feature fusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 3018–3026.
- Xu, M.; Wang, Y.; Liu, Y.; He, T.; and Qiao, Y. 2023b. CP3: Unifying point cloud completion by pretrain-prompt-predict paradigm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8): 9583–9594.
- Yang, B.; Pfreundschuh, P.; Siegart, R.; Hutter, M.; Moghadam, P.; and Patil, V. 2024. Tulip: Transformer for upsampling of lidar point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15354–15364.
- Ye, S.; Chen, D.; Han, S.; Wan, Z.; and Liao, J. 2021. Meta-PU: An arbitrary-scale upsampling network for point cloud. *IEEE transactions on visualization and computer graphics*, 28(9): 3206–3218.
- Yu, L.; Li, X.; Fu, C.-W.; Cohen-Or, D.; and Heng, P.-A. 2018. Pu-net: Point cloud upsampling network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2790–2799.
- Yu, X.; Rao, Y.; Wang, Z.; Liu, Z.; Lu, J.; and Zhou, J. 2021. PointR: Diverse point cloud completion with geometry-aware transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 12498–12507.
- Yuan, W.; Khot, T.; Held, D.; Mertz, C.; and Hebert, M. 2018. Pcn: Point completion network. In *2018 International Conference on 3D Vision (3DV)*, 728–737. IEEE.
- Zhao, W.; Liu, X.; Zhong, Z.; Jiang, J.; Gao, W.; Li, G.; and Ji, X. 2022. Self-supervised arbitrary-scale point clouds upsampling via implicit neural representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1999–2007.
- Zhao, Y.; Hui, L.; and Xie, J. 2021. Sspu-net: Self-supervised point cloud upsampling via differentiable rendering. In *Proceedings of the 29th ACM international conference on multimedia*, 2214–2223.
- Zhou, H.; Cao, Y.; Chu, W.; Zhu, J.; Lu, T.; Tai, Y.; and Wang, C. 2022. Seedformer: Patch seeds based point cloud completion with upsample transformer. In *European conference on computer vision*, 416–432. Springer.
- Zhou, K.; Dong, M.; Zhi, P.; and Wang, S. 2024. Cascaded Network with Hierarchical Self-Distillation for Sparse Point Cloud Classification. In *2024 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Zhu, Z.; Chen, H.; He, X.; Wang, W.; Qin, J.; and Wei, M. 2023. Svdformer: Complementing point cloud via self-view augmentation and self-structure dual-generator. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14508–14518.