2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

PointSFDA: Source-free Domain Adaptation for Point Cloud Completion

Xing He, Zhe Zhu, Liangling Nan, Wenshuo Peng, Honghua Chen, Mingqiang Wei

Abstract— **Point cloud completion is critical for autonomous driving and robotic perception, yet deep learning models often experience severe performance degradation under the domain gap between synthetic training and real-world data. While unsupervised domain adaptation (UDA) has been explored to mitigate this issue, its reliance on access to source datasets limits practical applicability, as source data are often proprietary or restricted. We pioneer source-free domain adaptation (SFDA) for point cloud completion, which adapts a pre-trained source model to an unlabeled target domain without requiring source data access. To this end, we propose PointSFDA, a framework that combines global knowledge transfer with target-specific local adaptation. Specifically, we design (i) a Coarse-to-Fine Point Cloud Distillation module to extract domain-invariant global geometric priors from the source model, and (ii) a Partial-Mask Consistency Training strategy to enforce prediction consistency across masking augmentations, enabling self-supervised learning of local target-domain geometry. Experiments on real-world datasets (KITTI, ScanNet) and synthetic benchmarks (ModelNet40, 3D-FUTURE) demonstrate that PointSFDA achieves significant improvements over state-of-the-art methods in cross-domain shape completion, establishing a practical and scalable solution for robotics applications. Our code is available at `https://github.com/Starak-x/PointSFDA`.**

I. INTRODUCTION

Point clouds serve as a widely used 3D data representation in autonomous driving and robotics. However, inherent challenges such as self-occlusion, light reflection, and limited sensor resolution often lead to incomplete and noisy raw point cloud data. Therefore, reconstructing complete and accurate shapes from partial observations is essential for downstream tasks.

Recent advances in supervised point cloud completion [1], [2], [3], [4], [5], [6], [7] have achieved strong results on synthetic benchmarks, where partial inputs are generated via virtual scanning. However, these methods degrade substantially in real-world scenarios due to the absence of ground-truth data. To mitigate this limitation, unsupervised methods based on unpaired data [8], [9], [10] have been proposed, but their outputs are often biased toward synthetic distributions, leaving the domain gap unresolved. Self-supervised approaches [11], [12], [13] instead exploit real-world inputs without paired ground truth, yet struggle to capture global structures and fine-grained geometry due to the lack of explicit supervision. These limitations highlight the importance of explicitly addressing distribution shifts between synthetic and real domains.

Unsupervised Domain Adaptation (UDA) emerges as a natural solution. As illustrated in Fig. 1, by aligning source (synthetic) and target (real) distributions during training, UDA methods [14], [15] demonstrate improved generalization. Yet, conventional UDA assumes simultaneous access to both source and target data, which is impractical in certain critical circumstances like autonomous driving where source data are often proprietary or inaccessible [16].

To address these challenges, we pioneer the study of source-free domain adaptation (SFDA) for point cloud completion. As illustrated in Fig. 1, SFDA circumvents the impractical source-data dependency of UDA by adapting the pre-trained model using only unlabeled real-world target scans. While SFDA has been explored for image classification [17], [18], [19], semantic segmentation [20], [16], [21], and object detection [22], [23], existing methods are inherently designed for semantic prediction tasks relying on pseudo-labeling or source-style sample synthesis. These strategies are ill-suited for point cloud completion, a dense geometric reconstruction task characterized by unstructured and unordered nature of point cloud. Effective adaptation in this context requires transferring geometric knowledge rather than semantic labels or features.

In this work, we propose PointSFDA, the first source-free domain adaptation framework for point cloud completion. We observe that shapes across domains often share similar global structures but differ in local point distributions. Based on this observation, our approach transfers global geometric knowledge from the pre-trained source model to the target domain while promoting self-supervised learning of local geometry directly from the target data. Specifically, we introduce Coarse-to-Fine Point Cloud Distillation, which directly distills geometric knowledge through point clouds rather than via feature distillation. Furthermore, we develop a Partial-Mask Consistency Training strategy that enforces point-wise predictive consistency under varying mask augmentations to learn the local geometry of the target domain in a self-supervised manner.

To evaluate the effectiveness of our method, we conduct experiments in both real-world datasets, including

This work was supported by the National Natural Science Foundation of China (No. T2322012, No. 62572240). (Corresponding author: Mingqiang Wei.)

X. He, Z. Zhu and M. Wei are with the School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, Nanjing, China (emails: hexing@nuaa.edu.cn; zhuzhe0619@nuaa.edu.cn; mingqiang.wei@gmail.com).

L. Nan is with the Urban Data Science Section, Delft University of Technology, Delft, Netherlands (email: liangliang.nan@tudelft.nl).

W. Peng is with the Department of Electronic Engineering, Tsinghua University, Beijing, China (email: gin2pws@gmail.com).

H. Chen is with the School of Data Science, Lingnan University, Hong kong SAR (email: honghuachen@LN.edu.hk).

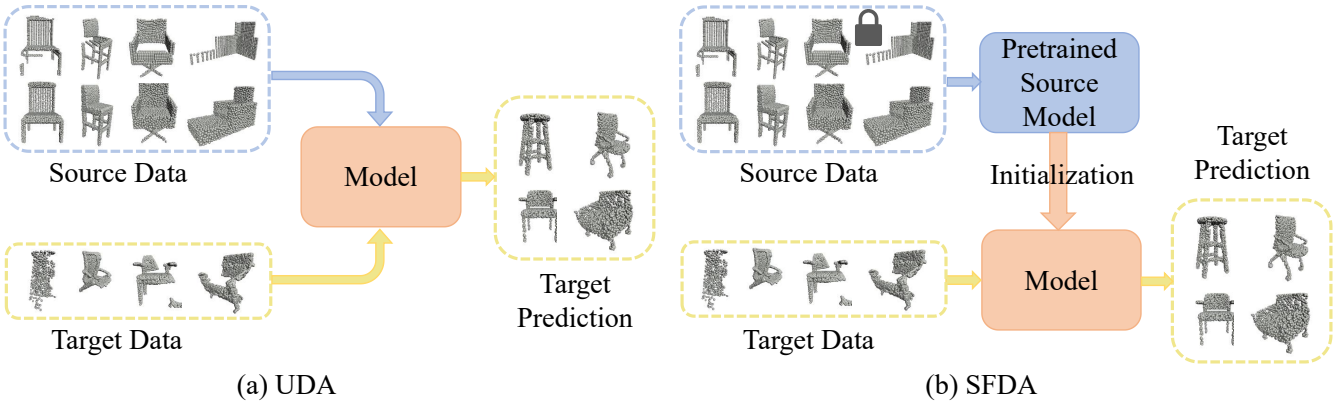

Fig. 1. Illustration of (a) unsupervised domain adaptation (UDA) and (b) source-free domain adaptation (SFDA) in point cloud completion. The UDA methods rely on both source and target datasets during adaptation, while SFDA methods only have access to a pretrained source model and unlabeled target data. Both UDA and SFDA aim to generate complete point clouds that resemble the distribution of real data.

KITTI [24] and ScanNet [25], and synthetic datasets [26], [27]. Experimental results demonstrate that despite the simplicity of our approach, it is highly effective. It significantly enhances the performance of state-of-the-art point cloud completion networks in cross-domain data, outperforming existing UDA and unsupervised methods by a considerable margin.

In summary, our main contributions are as follows:

- We design the first source-free domain adaptation framework specifically tailored for point cloud completion, making a significant advancement in cross-domain point cloud completion;
- We propose coarse-to-fine point cloud distillation to effectively transfer geometry knowledge from the pretrained source model to the target model;
- We introduce partial-mask consistency training, which enables the learning of local geometry information from the target domain in a self-supervised manner and enhances the source model to reduce noise.

II. Method

A. Problem Description and Overview

In the scenario of source-free domain adaptation, we are given a point cloud completion network Φ^s trained on a source dataset. Our goal is to train a target model Φ^t with only access to the source pretrained model Φ^s and the target dataset $\left\{P_{in}^{i,t}\right\}_{i=1}^{N}$, which compromises N partial point clouds without ground truth. We masked point clouds $\left\{P_{in,i}^{t}\right\}_{i=0}^{K}$, where K represents the number of masked point clouds in partial-mask consistency training, and $P_{in,0}^{t}$ equals the original partial point cloud P_{in}^{t}. In our setting, the target model Φ^t shares the same architecture with Φ^s. Following previous point cloud completion methods [1], [2], [28], [29], [7], [5], [3], [4], [30], [31], [6], we assume the network generates results in a coarse-to-fine manner, where a coarse shape P_{coarse}^{s} ($\left\{P_{coarse,i}^{t}\right\}_{i=0}^{K}$) is first produced and then refined to the final result P_{fine}^{s} ($\left\{P_{fine,i}^{t}\right\}_{i=0}^{K}$).

The overall framework of PointSFDA is illustrated in Fig. 2. Our approach comprises two core components: Coarse-to-fine Point Cloud Distillation and Partial-Mask Consistency Training. The former serves as a knowledge transfer module to exploit the learned domain-invariant information within Φ^s, while the latter focuses on learning the target-relevant information.

B. Coarse-to-Fine Point Cloud Distillation

In the absence of source data, most SFDA methods rely on knowledge distillation from pretrained source models [32], [33], [34], [35], [36]. However, direct transfer of class labels or features is insufficient for point cloud completion, as it fails to capture the underlying geometric structure. We observe that global shape structures within a category remain largely domain-invariant, motivating a coarse-to-fine distillation strategy that transfers geometry knowledge at both coarse and fine levels.

Our framework comprises two sequential stages: (1) Coarse-level distillation, which aligns low-resolution global shapes to establish structural correspondence, and (2) Fine-level distillation, which refines local details while preserving global fidelity through asymmetric matching.

a) Fine-level distillation: To guide fine-scale generation, source point clouds are downsampled using Farthest Point Sampling (FPS), retaining global structure while suppressing noise. The similarity to target predictions is measured using the Unidirectional Chamfer Distance (UCD):

$$\mathcal{L}_{fine} = \sum_{i=0}^{K} UCD\left(FPS(P_{fine}^{s}), P_{fine,i}^{t}\right), \tag{1}$$

where

$$UCD(X,Y) = \frac{1}{|X|}\sum_{x\in X}\min_{y\in Y}\|x-y\|. \tag{2}$$

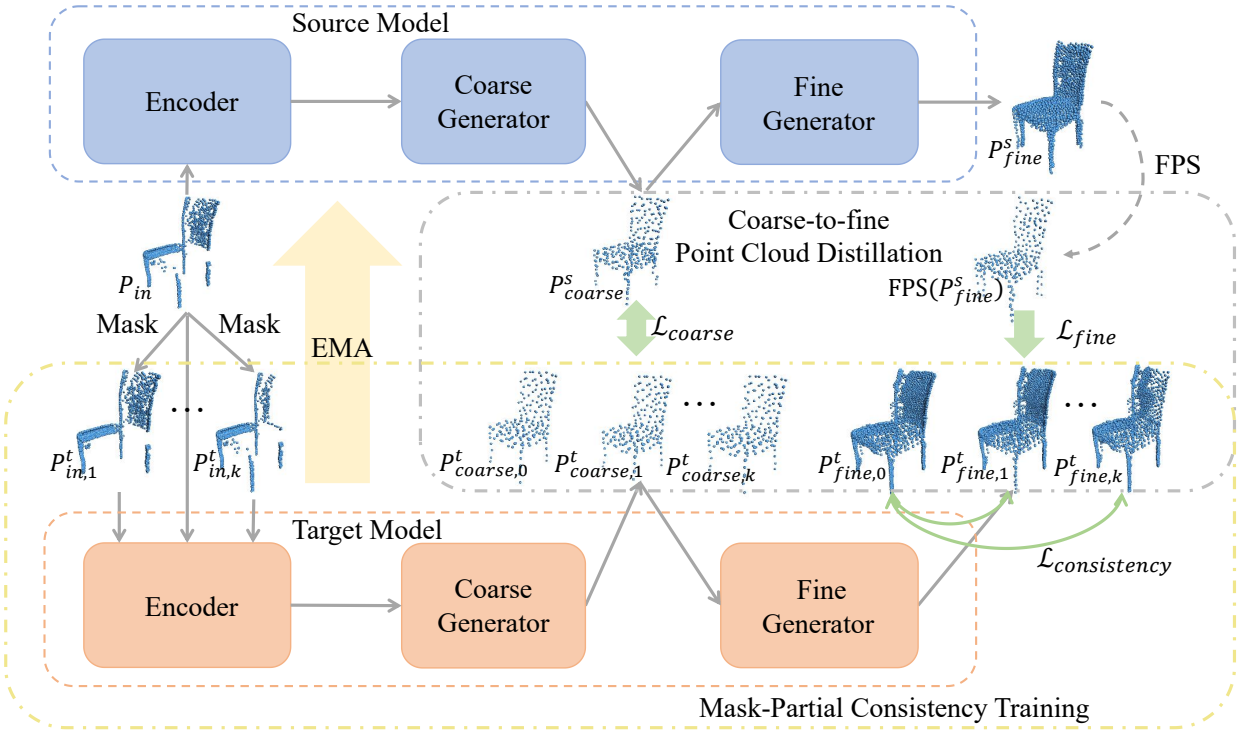

Fig. 2. Overview of the proposed PointSFDA. A partial point cloud from the target domain is directly used as input for the source model and masked k times ($k = 2$ in this figure) before being fed into the target model. Coarse-to-fine Point Cloud Distillation is employed in the coarse point cloud and fine point cloud to directly transfer the geometry information across domains. Mask-partial Consistency Training learns target data geometry information through point-wise predictive consistency across various mask augmentations.

Unlike symmetric Chamfer Distance, UCD only measures the deviation from source to target, preventing overfitting to noisy or misaligned local patterns.

b) Coarse-level distillation: Coarse point clouds, generated at intermediate stages of the completion network, serve as global shape representations. We enforce alignment across domains via the symmetric Chamfer Distance (CD):

$$\mathcal{L}_{coarse} = \sum_{i=0}^{K} CD\big(P^s_{coarse}, P^t_{coarse,i}\big)\,, \tag{3}$$

with

$$CD(X,Y) = \frac{1}{|X|}\sum_{x\in X}\min_{y\in Y}\|x-y\| + \frac{1}{|Y|}\sum_{y\in Y}\min_{x\in X}\|y-x\|. \tag{4}$$

By combining asymmetric fine-level guidance (UCD) with symmetric coarse-level alignment (CD), the proposed approach learns domain-invariant geometry while mitigating noise from local discrepancies.

C. Partial-Mask Consistency Training

While global knowledge transferred from the source model enables the target model to capture overall structure, the reconstruction of local details remains suboptimal without target-specific information. To address this, we introduce Partial-Mask Consistency Training (PMCT), which enhances local geometric learning by enforcing prediction consistency under masking augmentations.

Given a target input point cloud P_{in}, we uniformly divide it into K segments ($e.g., K = 8$) and randomly mask one to generate variants $\{P^t_{in,i}\}_{i=0}^{K}$. Passing these through the target model yields predictions $\{P^t_{fine,i}\}_{i=0}^{K}$. Consistency is enforced by minimizing the Chamfer Distance between the full prediction $P^t_{fine,0}$ and each masked prediction:

$$\mathcal{L}_{consistency} = \sum_{i=1}^{K} CD\big(P^t_{fine,0}, P^t_{fine,i}\big)\,. \tag{5}$$

Masking discards different local structures across variants, encouraging the model to infer missing regions and thereby strengthening its perception of local geometry.

To leverage knowledge learned by the target model for rectifying the source model and further reducing domain-induced noise, we continuously update the source model parameters Φ^s using the target model parameters Φ^t. Inspired by Mean Teacher [37], we adopt an Exponential Moving Average (EMA) update rule:

$$\Phi^s \leftarrow \eta\Phi^s + (1-\eta)\Phi^t, \tag{6}$$

where momentum $\eta \in [0,1]$ controls the update rate. This feedback mechanism allows the source and target models to iteratively refine each other throughout training.

D. Overall loss function

In addition to the loss functions mentioned above given by Equations (1), (3) and (5), the final output point cloud

needs to preserve the structure of the partial point cloud. To achieve this, we employ the partial match loss, i.e.,

$$\mathcal{L}_{partial} = \sum_{i=0}^{K} UCD\left(P_{in}, P_{fine,i}^{t}\right). \quad (7)$$

During the joint training, the overall loss is defined as:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{fine} + \lambda_2 \mathcal{L}_{coarse} + \lambda_3 \mathcal{L}_{consistency} + \lambda_4 \mathcal{L}_{partial} \quad (8)$$

where λ_1, λ_2, λ_3, and λ_4 are tradeoff hyperparameters.

III. Experiments

A. Implementation Details

Our method is implemented using PyTorch [38] and trained on a single NVIDIA GeForce RTX 4090 GPU. In the Coarse-to-Fine Point Cloud Distillation, the output point cloud P_{fine}^{s} from the source model is downsampled to 256 points. In the Partial-Mask Consistency Training, we set the number of masked variants $K = 1$ for computational efficiency while maintaining effective regularization. The loss weights are empirically configured as $\lambda_1 = 1$, $\lambda_2 = 1$, $\lambda_3 = 10^2$, and $\lambda_4 = 10^2$, balancing the contributions of global distillation, local consistency, and model refinement objectives during training. We apply the proposed PointSFDA to three representative point cloud completion networks: PCN [1], SnowflakeNet [4], and AdaPoinTr [5]. In all the experiments, the input and output point clouds are downsampled to 2,048 points.

We evaluate the proposed PointSFDA framework on three representative point cloud completion architectures: PCN [1], SnowflakeNet [4], and AdaPoinTr [5]. Across all experiments, both input partial point clouds and target completed point clouds are uniformly downsampled to 2,048 points for consistency and fair comparison.

B. Dataset and Evaluation Metrics

Following OptDE [14], we use the CRN [3] dataset as the source domain dataset in all experiments. CRN is derived from ShapeNet [39], which consists of eight categories: plane, car, cabinet, chair, lamp, sofa, table, and watercraft, totaling 30,174 partial-complete pairs. We utilize the same partitioning strategy as OptDE [14] for training and evaluation.

To evaluate PointSFDA in real-world scans, we first conduct experiments on KITTI [24] and ScanNet [25]. The cars in KITTI and the chairs, lamps, sofas, and tables in ScanNet are used. To facilitate the quantitative evaluation of performance on real datasets, we utilize the ground truth provided by the ScanSalon dataset [15], where the high-quality complete 3d models are created by artists.

In addition, we test PointSFDA on 3D-FUTURE [26] and ModelNet [27] for evaluation under different kinds of domain shifts. 3D-FUTURE incorporates photo-realistic synthetic images alongside high-resolution textured 3D CAD furniture shapes crafted by professional designers. As a result, samples in 3D-FUTURE closely mimic real-life objects. The partial and complete shapes within 3D-FUTURE are captured from five distinct viewpoints. ModelNet is derived from ModelNet40 [27], where The complete shapes are formed from randomly sampled points on the surface, while partial shapes are generated through virtual scanning techniques. The complete and incomplete point clouds in 3DFUTURE and ModelNet are both sampled to 2048 points, to align with CRN. We employ Chamfer Distance with L_2 normalization as the evaluation metric.

We conduct experiments on KITTI [24] and ScanNet [25] for real data evaluation. We compare our method with self-supervised methods P2C [13], ACL-SPC [12], and UDA method OptDE [14].

The quantitative results are summarized in Tab. I, where PCN, AdaPoinTr, and SnowflakeNet denote these models pretrained on the CRN dataset without adaptation. SFDA-PCN, SFDA-AdaPoinTr, and SFDA-SnowflakeNet represent the corresponding models adapted by the proposed PointSFDA.

TABLE I

Completion results on Real dataset in terms of per-point L_2 Chamfer Distance$\times 10^4$ (lower is better).

Method	Car	Chair	Lamp	Sofa	Table	Avg.
ACL-SPC [12]	31.94	20.03	35.05	31.09	25.08	28.64
P2C [13]	25.47	20.30	51.59	38.08	22.47	31.58
OptDE [14]	26.89	18.18	40.62	20.09	28.15	26.79
PCN [1]	37.09	25.33	66.46	18.14	60.43	41.49
SFDA-PCN	34.91	16.13	32.69	15.32	27.42	25.29
AdaPoinTr [5]	34.52	19.18	70.41	15.90	46.72	37.35
SFDA-AdaPoinTr	**20.51**	13.49	**30.32**	**11.17**	**21.85**	**19.47**
SnowflakeNet [4]	32.39	15.69	51.00	12.45	33.53	29.01
SFDA-SnowflakeNet	29.89	**12.81**	39.52	12.35	21.97	23.31

C. Real Data Evaluation

It can be observed that, although the source models are well-trained on the synthetic CRN dataset, the performance is far from satisfactory when tested on real-world data, typically performing worse than the self-supervised and UDA methods. After adaptation, all three models have shown significant improvements. In particular, our method can reduce the CD value of AdaPoinTr by 47.9%. Meanwhile, without access to the source data, all three models achieve superior performance compared to the state-of-the-art UDA method OptDE. This further outlines the ability of PointSFDA to effectively adapt to the target domain data.

In addition, we present a visual comparison in Fig. 3. Notably, when encountering a domain gap, AdaPoinTr exhibits errors in both global structure and geometric details. For example, the armrests of the chair model are merged. With the assistance of PointSFDA, the network avoids this issue and produces a result closer to the ground truth. Meanwhile, it is evident that our SFDA-AdaPoinTr significantly outperforms ACL-SPC, P2C, and OptDE by a large margin.

D. Synthetic Data Evaluation

In this study, we compare our method with existing unsupervised point cloud methods, including Pcl2Pcl [9],

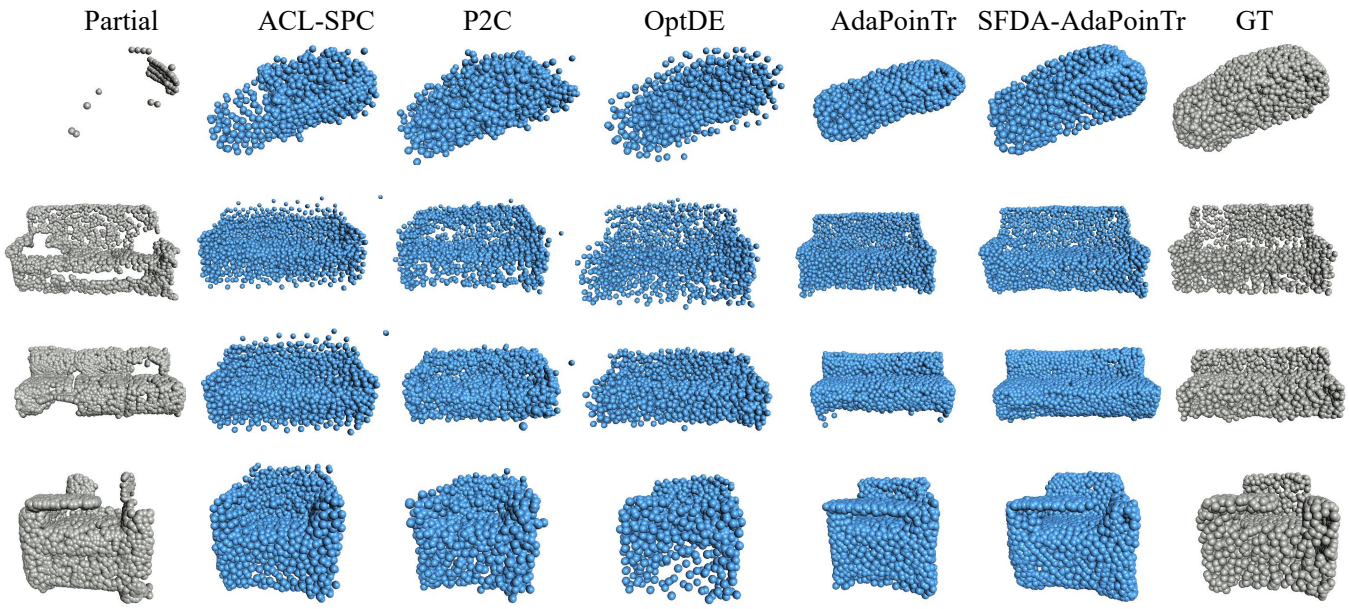

Fig. 3. Visual comparison with recent methods(ACL-SPC [12], P2C [13], OptDE [14], AdaPoinTr [5]) on the test set of KITTI and ScanNet.

ShapeInversion [10], and Cycle4completion [8], as well as self-supervised point cloud completion methods including P2C [13], ACL-SPC [12], and UDA methods like OptDE [14], DAPoinTr [40].

The quantitative results for the 3DFUTURE dataset and ModelNet dataset are presented in Tab. II and Tab. III, respectively. It is clear that our methods surpass other approaches by a significant margin. Our outstanding performance across different datasets demonstrates its capacity to adapt to diverse domain gaps.

TABLE II

COMPLETION RESULTS ON THE 3DFUTURE DATASET IN TERMS OF PER-POINT L_2 CHAMFER DISTANCE$\times 10^4$ (LOWER IS BETTER).

Method	Cabinet	Chair	Lamp	Sofa	Table	Avg.
Pcl2pcl [9]	57.23	43.91	157.86	63.23	141.92	92.83
ShapeInversion [10]	38.54	26.30	48.57	44.02	108.60	53.21
Cycle4Completion [8]	32.62	34.08	77.19	43.05	40.00	45.39
P2C [13]	43.01	24.57	26.52	41.60	22.69	31.68
ACL-SPC [12]	70.12	23.87	31.75	28.74	25.38	35.97
OptDE [14]	28.37	21.87	29.92	37.98	26.81	28.99
DAPointr [40]	18.46	17.60	27.91	23.08	24.71	22.35
PCN [1]	39.13	28.16	86.34	33.62	71.37	51.72
SFDA-PCN	24.73	17.94	27.54	19.60	22.43	22.45
AdaPoinTr [5]	26.96	22.71	50.45	25.24	45.26	34.12
SFDA-AdaPoinTr	16.72	12.31	17.44	14.15	14.47	15.08
SnowflakeNet [4]	19.43	17.08	32.84	19.20	29.26	23.80
SFDA-SnowflakeNet	**15.82**	**11.98**	**15.04**	**13.43**	**13.93**	**14.04**

We present some randomly selected qualitative results of 3DFUTURE and ModelNet in Fig. 4. Planes and cars are from ModelNet, while chairs, sofas, and tables are from the 3DFUTURE dataset. It is evident that ACL-SPC and P2C tend to generate average shapes. Additionally, P2C tends to produce a significant amount of noise points. While OptDE effectively restores the overall shape, it fails to complete the engine of the airplane and the armrests of the chair.

Due to the domain gap, SnowflakeNet incorrectly completes the overall shape, as seen in the erroneous chair legs and the incorrect size of the tabletop. In contrast, our SFDA-SnowflakeNet completes entire shapes and recovers finer details, outperforming the aforementioned methods.

E. Ablation Study

In this section, we conduct extensive ablation studies on the key components of our PointSFDA framework. All experiments in our ablation study utilize SnowflakeNet as the backbone and are conducted on the 3DFUTURE dataset.

1) Ablation on Coarse-to-fine Point Cloud Distillation: We start from the baseline model A, where only the partial match loss $L_{partial}$ and L_{fine} are used. Subsequently, we incorporate L_{coarse} in variant B, which uses coarse-to-fine point cloud distillation without partial-mask consistency training. The improvement indicates that L_{coarse} can effectively transfer the global structure. Variant C only employs partial-mask consistency training in a self-supervised manner. Compared to our method, variant C shows an average increase of 4.65 in CDL_2, indicating the effectiveness of the coarse-to-fine point cloud distillation. In Variant D, we replace L_{coarse} with $L_{feature}$, which minimizes the cosine similarity between the global features of the source model and the target model. The performance drop indicates that the distillation of global features can transfer the global structure to some extent, but its effectiveness is not as good as our point cloud distillation.

2) Ablation on Partial-Mask Consistency Training: Additionally, to validate the effectiveness of Partial-Mask Consistency Training (PMCT), we utilize Mean-Teacher (MT)

TABLE III

COMPLETION RESULTS ON MODELNET DATASET IN TERMS OF PER-POINT L_2 CHAMFER DISTANCE$\times 10^4$ (LOWER IS BETTER).

Method	Plane	Car	Chair	Lamp	Sofa	Table	Avg.
Pcl2pcl [9]	18.53	17.54	43.58	126.80	38.78	163.62	68.14
ShapeInversion [10]	3.78	15.66	22.25	60.42	22.25	125.31	41.61
Cycle4Completion [8]	5.77	11.85	26.67	83.34	22.82	21.47	28.65
P2C [13]	4.80	19.66	17.68	44.69	32.26	12.83	21.99
ACL-SPC [12]	5.75	11.73	43.08	106.29	25.62	16.89	34.89
OptDE [14]	2.18	9.80	14.71	39.74	19.43	**9.75**	15.94
DAPoinTr [40]	2.38	8.04	13.83	33.26	12.72	12.51	13.79
PCN [1]	5.12	9.38	25.45	75.03	19.24	81.12	35.89
SFDA-PCN	2.49	9.56	14.31	48.02	17.39	17.89	18.28
AdaPoinTr [5]	2.39	8.06	20.15	45.37	15.88	55.93	24.63
SFDA-AdaPoinTr	1.99	7.80	**9.89**	**28.53**	11.74	12.41	12.06
SnowflakeNet [4]	2.50	7.93	14.55	40.95	12.79	24.44	17.19
SFDA-SnowflakeNet	**1.72**	**7.59**	10.43	28.60	**10.71**	10.75	**11.63**

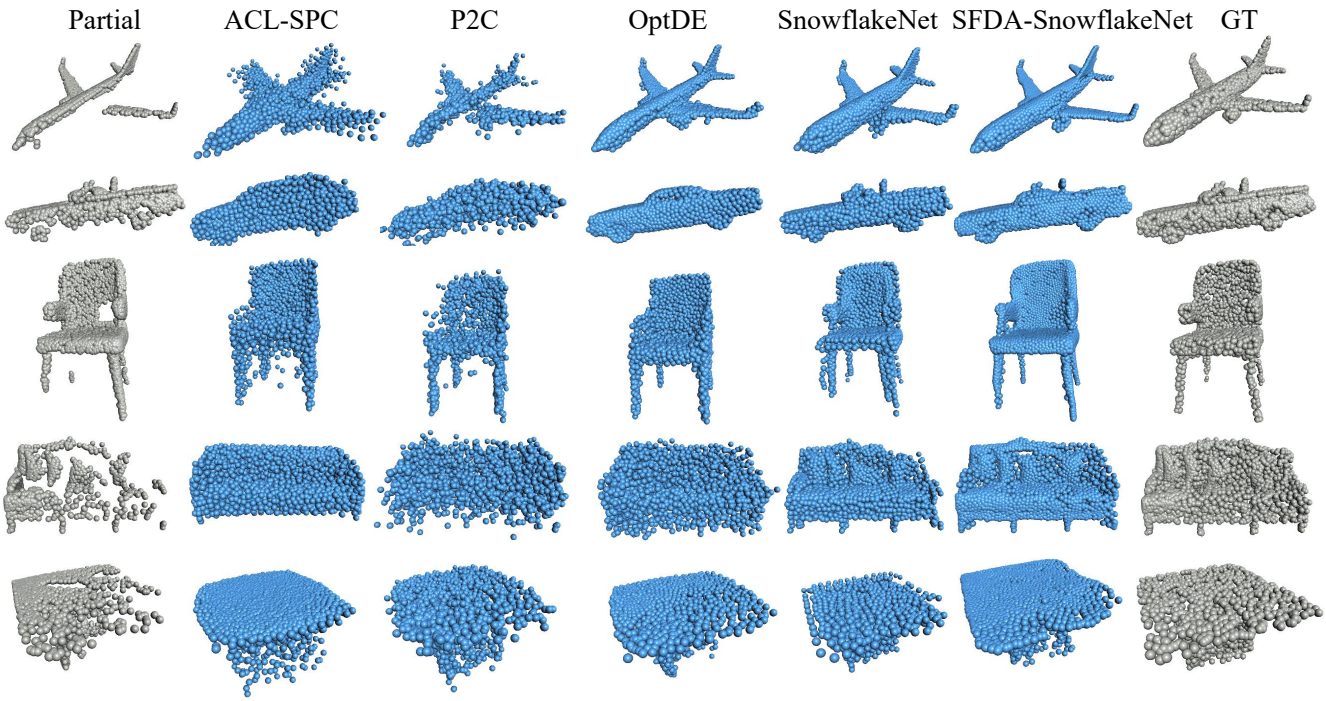

Fig. 4. Visual comparison with recent methods(ACL-SPC [12], P2C [13], OptDE [14], SnowflakeNet [4]) on the test set of 3D-FUTURE and ModelNet. The plane and car come from ModelNet, while the chair, sofa, and table are from 3D-FUTURE.

as a replacement in variant E. Our method demonstrates a significant improvement compared to Variant E, with an average increase in CDL_2 of 4.59. Meanwhile, comparing Variant E with Variant B, it can be observed that adding a Mean-Teacher to Coarse-to-fine point cloud distillation provides insignificant benefits.

Moreover, our visualization in Fig. 5 demonstrates that exclusively employing Coarse-to-fine point cloud distillation leads to inaccuracies in local structures, such as noise points near chair legs, attributed to domain gap. However, the addition of PMCT significantly alleviates this problem.

3) Ablation on Mask Number: We conducted an ablation study concerning the number of masked point clouds (K). We set K to 1, 2, 3, and 4 respectively. The results of the per-point L_2 Chamfer Distance $\times 10^4$ are presented in Tab. V. The results indicate that more masked input cannot bring any performance improvement. Thus to save on storage and computational expenses, we set $K = 1$ for all experiments.

4) Ablation on the Number of Sampled Point Clouds: We conducted an ablation study on the number of sampled point clouds $FPS\left(P^s_{fine}\right)$ in Coarse-to-fine Point Cloud Distillation. The results are presented in Tab. VI. We observe that with the increase in sampled numbers, the performance in most categories decreases to some extent. Thus, we sample 256 points to represent the global structure in our method.

TABLE IV

ABLATION STUDY ON 3DFUTURE DATASET IN TERMS OF PER-POINT L_2 CHAMFER DISTANCE$\times 10^4$ (LOWER IS BETTER). WE UTILIZE SNOWFLAKENET AS THE BACKBONE. MT AND PMCT REFER TO MEAN-TEACHER AND PARTIAL-MASK CONSISTENCY TRAINING, RESPECTIVELY.

Variants	L_{fine}	L_{coarse}	$L_{feature}$	MT	PMCT	Cabinet	Chair	Lamp	Sofa	Table	Avg.
A	✓					18.19	13.64	26.08	16.83	21.04	19.16
B	✓	✓				18.52	13.56	25.39	16.39	20.99	18.97
C					✓	18.48	19.74	17.79	20.46	16.99	18.69
D	✓		✓		✓	16.35	13.32	16.39	14.87	15.27	15.24
E	✓	✓		✓		16.92	13.81	25.15	17.40	20.85	18.63
Ours	✓	✓			✓	**15.82**	**11.96**	**15.04**	**13.43**	**13.93**	**14.04**

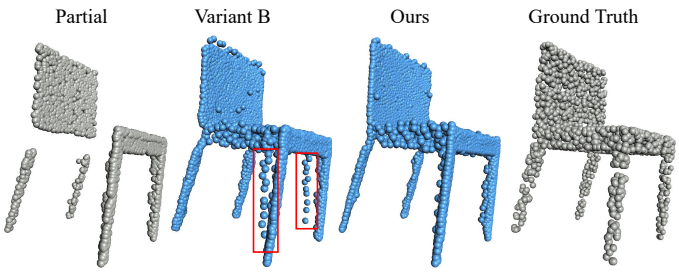

Fig. 5. The ablation study of PMCT. In Variant B, only coarse-to-fine point cloud distillation is used, without employing PMCT. Notably, Variant B may recover false chair legs due to the domain gap.

TABLE V

ABLATION STUDY ON THE 3DFUTURE DATASET REGARDING THE NUMBER OF MASKED POINT CLOUDS K, MEASURED IN TERMS OF PER-POINT L_2 CHAMFER DISTANCE$\times 10^4$ (LOWER IS BETTER).

K	Cabinet	Chair	Lamp	Sofa	Table	Avg.
$K=1$	15.82	11.98	15.04	13.43	13.93	14.04
$K=2$	15.98	12.07	15.22	13.50	14.43	14.24
$K=3$	15.96	12.73	15.00	13.49	13.65	14.17
$K=4$	15.90	12.05	15.47	13.50	14.68	14.32

5) Ablation on Mask Strategy: We also conducted an ablation study regarding mask strategy. The results are presented in Tab. VII. In the "Partition" strategy, we partition the input point cloud P_{in} into 8 equal parts in 3D space and randomly mask one of them each time to obtain masked point clouds. In the "View" strategy, we select one random point and then mask out a fixed number of points nearest to it. To ensure a fair comparison, the number of masked points is set to $1/8$ of the total number of points. As shown in Tab. VII, the "Partition" strategy exhibits a slight advantage over the "View" strategy. Consequently, we chose to employ the "Partition" strategy in our method.

TABLE VI

ABLATION STUDY ON THE 3DFUTURE DATASET REGARDING THE NUMBER OF SAMPLED POINT CLOUDS, MEASURED IN TERMS OF PER-POINT L_2 CHAMFER DISTANCE$\times 10^4$ (LOWER IS BETTER).

Sample Number	Cabinet	Chair	Lamp	Sofa	Table	Avg.
128	15.86	11.92	14.92	13.68	14.31	14.14
256	**15.82**	11.98	15.04	**13.43**	**13.93**	**14.04**
512	15.87	**11.78**	15.29	13.64	14.27	14.17
1024	16.18	12.26	15.06	13.54	13.94	14.20
2048	16.35	12.67	**14.70**	13.56	14.03	14.26

TABLE VII

ABLATION STUDY ON THE 3DFUTURE DATASET REGARDING THE METHOD OF MASKING, MEASURED IN TERMS OF PER-POINT L_2 CHAMFER DISTANCE$\times 10^4$ (LOWER IS BETTER).

Masking Method	Cabinet	Chair	Lamp	Sofa	Table	Avg.
Partition	**15.82**	**11.98**	**15.04**	13.43	13.93	**14.04**
View	16.96	12.07	16.62	**12.97**	**13.25**	14.37

IV. CONCLUSIONS

We propose PointSFDA for source-free domain adaptation of point cloud completion. PointSFDA is a general architecture applicable to all point cloud completion networks that generate point clouds in a coarse-to-fine manner. PointSFDA can effectively transfer geometric knowledge from the pre-trained source model to the target model by Coarse-to-fine Point Cloud Distillation. Meanwhile, we introduce partial-mask consistency training to learn the local structure information of the target domain. Extensive experiments demonstrate that PointSFDA can significantly improve the performance of cross-domain point cloud completion.

REFERENCES

[1] W. Yuan, T. Khot, D. Held, C. Mertz, and M. Hebert, "Pcn: Point completion network," in *2018 international conference on 3D vision (3DV)*. IEEE, 2018, pp. 728–737.

[2] J. Zhang, J. Shao, J. Chen, D. Yang, B. Liang, and R. Liang, "Pfnet: an unsupervised deep network for polarization image fusion," *Optics letters*, vol. 45, no. 6, pp. 1507–1510, 2020.

[3] 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*, 2020, pp. 790–799.

[4] 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*, 2021, pp. 5499–5509.

[5] X. Yu, Y. Rao, Z. Wang, J. Lu, and J. Zhou, "Adapointr: Diverse point cloud completion with adaptive geometry-aware transformers," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.

[6] Z. Zhu, H. Chen, X. He, W. Wang, J. Qin, and M. Wei, "Svdformer: Complementing point cloud via self-view augmentation and self-structure dual-generator," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 14 508–14 518.

[7] S. Li, P. Gao, X. Tan, and M. Wei, "Proxyformer: Proxy alignment assisted point cloud completion with missing part sensitive transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9466–9475.

[8] 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*, 2021, pp. 13 080–13 089.

[9] X. Chen, B. Chen, and N. J. Mitra, "Unpaired point cloud completion on real scans using adversarial training," *arXiv preprint arXiv:1904.00069*, 2019.

[10] J. Zhang, X. Chen, Z. Cai, L. Pan, H. Zhao, S. Yi, C. K. Yeo, B. Dai, and C. C. Loy, "Unsupervised 3d shape completion through gan inversion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1768–1777.

[11] H. Mittal, B. Okorn, A. Jangid, and D. Held, "Self-supervised point cloud completion via inpainting," *arXiv preprint arXiv:2111.10701*, 2021.

[12] S. Hong, M. Yavartanoo, R. Neshatavar, and K. M. Lee, "Acl-spc: Adaptive closed-loop system for self-supervised point cloud completion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9435–9444.

[13] R. Cui, S. Qiu, S. Anwar, J. Liu, C. Xing, J. Zhang, and N. Barnes, "P2c: Self-supervised point cloud completion from single partial clouds," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 14 351–14 360.

[14] J. Gong, F. Liu, J. Xu, M. Wang, X. Tan, Z. Zhang, R. Yi, H. Song, Y. Xie, and L. Ma, "Optimization over disentangled encoding: Unsupervised cross-domain point cloud completion via occlusion factor manipulation," in *European Conference on Computer Vision*. Springer, 2022, pp. 517–533.

[15] Y. Wu, Z. Yan, C. Chen, L. Wei, X. Li, G. Li, Y. Li, S. Cui, and X. Han, "Scoda: Domain adaptive shape completion for real scans," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17 630–17 641.

[16] Y. Liu, W. Zhang, and J. Wang, "Source-free domain adaptation for semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1215–1224.

[17] J. Liang, D. Hu, and J. Feng, "Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation," in *International conference on machine learning*. PMLR, 2020, pp. 6028–6039.

[18] J. Liang, D. Hu, Y. Wang, R. He, and J. Feng, "Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 11, pp. 8602–8617, 2021.

[19] M. Litrico, A. Del Bue, and P. Morerio, "Guiding pseudo-labels with uncertainty estimation for source-free unsupervised domain adaptation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7640–7650.

[20] F. Fleuret *et al.*, "Uncertainty reduction for model adaptation in semantic segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 9613–9623.

[21] S.-Y. Lo, P. Oza, S. Chennupati, A. Galindo, and V. M. Patel, "Spatio-temporal pixel-level contrastive learning-based source-free domain adaptation for video semantic segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 10 534–10 543.

[22] S. Li, M. Ye, X. Zhu, L. Zhou, and L. Xiong, "Source-free object detection by learning to overlook domain style," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 8014–8023.

[23] V. VS, P. Oza, and V. M. Patel, "Instance relation graph guided source-free domain adaptive object detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 3520–3530.

[24] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.

[25] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "Scannet: Richly-annotated 3d reconstructions of indoor scenes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5828–5839.

[26] H. Fu, R. Jia, L. Gao, M. Gong, B. Zhao, S. Maybank, and D. Tao, "3d-future: 3d furniture shape with texture," *International Journal of Computer Vision*, pp. 1–25, 2021.

[27] 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*, 2015, pp. 1912–1920.

[28] 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*, vol. 34, no. 07, 2020, pp. 11 596–11 603.

[29] 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*, 2021, pp. 12 498–12 507.

[30] H. Zhou, Y. Cao, W. Chu, J. Zhu, T. Lu, Y. Tai, and C. Wang, "Seedformer: Patch seeds based point cloud completion with upsample transformer," in *European conference on computer vision*. Springer, 2022, pp. 416–432.

[31] S. Zhang, X. Liu, H. Xie, L. Nie, H. Zhou, D. Tao, and X. Li, "Learning geometric transformation for point cloud completion," *International Journal of Computer Vision*, pp. 1–21, 2023.

[32] H. Yu, J. Huang, Y. Liu, Q. Zhu, M. Zhou, and F. Zhao, "Source-free domain adaptation for real-world image dehazing," in *Proceedings of the 30th ACM International Conference on Multimedia*, 2022, pp. 6645–6654.

[33] N. Karim, N. C. Mithun, A. Rajvanshi, H.-p. Chiu, S. Samarasekera, and N. Rahnavard, "C-sfda: A curriculum learning aided self-training framework for efficient source free domain adaptation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 24 120–24 131.

[34] D. Hegde, V. Kilic, V. Sindagi, A. B. Cooper, M. Foster, and V. M. Patel, "Source-free unsupervised domain adaptation for 3d object detection in adverse weather," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 6973–6980.

[35] S. Tang, Y. Shi, Z. Ma, J. Li, J. Lyu, Q. Li, and J. Zhang, "Model adaptation through hypothesis transfer with gradual knowledge distillation," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5679–5685.

[36] W. Chen, L. Lin, S. Yang, D. Xie, S. Pu, and Y. Zhuang, "Self-supervised noisy label learning for source-free unsupervised domain adaptation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 10 185–10 192.

[37] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," *Advances in neural information processing systems*, vol. 30, 2017.

[38] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, vol. 32, 2019.

[39] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, *et al.*, "Shapenet: An information-rich 3d model repository," *arXiv preprint arXiv:1512.03012*, 2015.

[40] Y. Li, Q. Zhou, J. Gong, Y. Zhu, R. Dazeley, X. Zhao, and X. Lu, "Dapointr: Domain adaptive point transformer for point cloud completion," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 5, 2025, pp. 5066–5074.