

# DAPoinTr: Domain Adaptive Point Transformer for Point Cloud Completion

Yinghui Li<sup>1\*</sup>, Qianyu Zhou<sup>2\*</sup>, Jingyu Gong<sup>3</sup>, Ye Zhu<sup>1</sup>, Richard Dazeley<sup>1</sup>,  
Xinkui Zhao<sup>4†</sup>, Xuequan Lu<sup>5†</sup>

<sup>1</sup>School of Information Technology, Deakin University, Australia;

<sup>2</sup>Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China;

<sup>3</sup>School of Computer Science and Technology, East China Normal University, Shanghai, China;

<sup>4</sup>School of Software Technology, Zhejiang University, Ningbo, China;

<sup>5</sup>Department of Computer Science and Software Engineering, The University of Western Australia, Australia.

s218161821@deakin.edu.au, zhouqianyu@sjtu.edu.cn, jygong@cs.ecnu.edu.cn, ye.zhu@deakin.edu.au,

richard.dazeley@deakin.edu.au, zhaoxinkui@zju.edu.cn, bruce.lu@uwa.edu.au

## Abstract

Point Transformers (PoinTr) have shown great potential in point cloud completion recently. Nevertheless, effective domain adaptation that improves transferability toward target domains remains unexplored. In this paper, we delve into this topic and empirically discover that direct feature alignment on point Transformer’s CNN backbone only brings limited improvements since it cannot guarantee sequence-wise domain-invariant features in the Transformer. To this end, we propose a pioneering Domain Adaptive Point Transformer (DAPoinTr) framework for point cloud completion. DAPoinTr consists of three key components: Domain Query-based Feature Alignment (DQFA), Point Token-wise Feature alignment (PTFA), and Voted Prediction Consistency (VPC). In particular, DQFA is presented to narrow the global domain gaps from the sequence via the presented domain proxy and domain query at the Transformer encoder and decoder, respectively. PTFA is proposed to close the local domain shifts by aligning the tokens, *i.e.*, point proxy and dynamic query, at the Transformer encoder and decoder, respectively. VPC is designed to consider different Transformer decoders as multiple of experts (MoE) for ensembled prediction voting and pseudo-label generation. Extensive experiments with visualization on several domain adaptation benchmarks demonstrate the effectiveness and superiority of our DAPoinTr compared with state-of-the-art methods.

**Code** — <https://github.com/Yinghui-Li-New/DAPoinTr>

## 1 Introduction

3D point cloud data plays a pivotal role across diverse domains, including autonomous vehicles, robotics, augmented and virtual reality, *etc.* Point cloud completion (PCC) aims to predict the complete geometric shape given a partial input. Previous works (Yang et al. 2018; Tchapmi et al. 2019; Huang et al. 2020; Xie et al. 2020; Yu et al. 2021, 2023) have significantly enhanced the performance of fully-supervised PCC models, using training datasets composed of pairs of

partial and complete point clouds. Recently, PoinTr (Yu et al. 2021) adopted a Transformer encoder-decoder architecture and has reformulated PCC as a set-to-set translation problem, showing excellent performance within the same domains. However, when transferring these models from source domains to target domains, they would suffer from substantially degraded performance on target datasets. This is mainly due to the existing domain gaps, *e.g.*, different occlusions, viewpoints, sensor resolutions, and light reflection, across different domains.

To address this issue, unsupervised domain adaptation (UDA) techniques are introduced into point cloud completion to enhance the transferability toward real domains. This involves several categories: adversarial learning (Chen, Chen, and Mitra 2019; Zhang et al. 2021), data reconstruction (Wen et al. 2021), disentangled learning (Gong et al. 2022), self-supervised learning (Hong et al. 2023). Nevertheless, a common limitation of these UDA PCC methods is their reliance on Convolution Neural Networks (CNNs) and suffer from limited receptive fields as layers are stacked, which can be problematic for capturing the broader context required for learning domain-agnostic features. With the recent surge of point Transformers, is it possible to empower them with such a capability to perform accurate completion in cross-domain scenarios?

A straightforward idea to address this issue is to directly apply adversarial alignment to features extracted by the DGCNN (Phan et al. 2018) backbone (referred to as DA-DGCNN) of PoinTr (Yu et al. 2021). As illustrated in Figure 1(a), such a manner does enhance the point Transformer’s cross-domain performance, but the improvement is very limited when transferring from CRN (Wang, Ang Jr, and Lee 2020) to ModelNet (Wu et al. 2015). We deduce that this limitation arises because aligning feature distributions on the DGCNN backbone does not ensure sequence-wise domain-invariant features in the subsequent Transformer, which are essential for accurate completion. Thus, the sequence features from the Transformer encoder remain domain-separated and not well-aligned, leading to less decent performance of the point Transformers due to these shifted sequence features. In Figure 1(a), we also provide

\*These authors contributed equally.

†Corresponding Authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

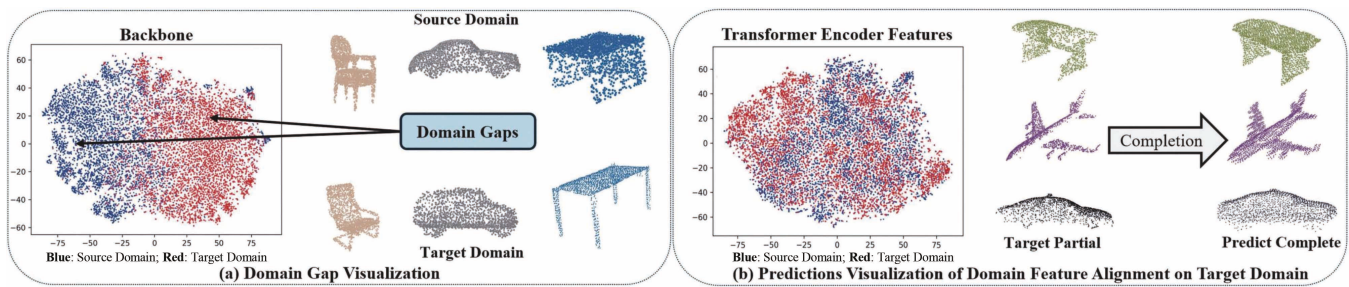


Figure 1: (a) Visualization of the domain discrepancy among objects of the same category from different domains where there are significant variations in topology, geometric patterns, and feature distribution. We observe that directly applying adversarial alignment to CNN backbone features of PoinTr (Yu et al. 2021) does not well align the source domains and target domains since it does not ensure learning domain-invariant sequence features of the point cloud. (b) In contrast, we propose a Domain Adaptive Point Transformer (DAPoinTr) framework for point cloud completion that well aligns the cross-domain distributions (visualization of Transformer encoder features) and manages to generate complete shapes of partial point cloud input.

the visualization of the domain discrepancy among objects of the same category from different domains where there are significant variations in topology and geometric patterns.

Motivated by the aforementioned analysis, we introduce a pioneering framework, namely Domain Adaptive Point Transformer (DAPoinTr), for the point cloud completion task. DAPoinTr aims to align the sequence-wise point features with Transformer encoder-decoder architecture to enhance the adaptability toward the real target domains. It consists of three key components: Domain Query-based Feature Alignment (DQFA), Point Token-wise Feature alignment (PTFA), and Voted Prediction Consistency (VPC). Specifically, DQFA is presented to narrow the global domain gaps in layout and inter-sketch relationships from the sequence via the presented domain proxy and domain query at the Transformer encoder and decoder, respectively. A domain proxy is designed to encode global features that reflect object layout and aggregate domain-specific features from the whole sequence. Besides, the domain query encodes some sketch relationships for adaptation and fuses context features from each dynamic query in the sequence. Secondly, PTFA is proposed to close the local domain shifts by aligning the token at the Transformer encoder and decoder, respectively. Concretely, Point Proxy Feature Alignment is presented to align each token in the encoder sequence, *i.e.*, point proxy, to alleviate domain gaps caused by local shape, appearance, *etc.* Besides, Dynamic Query Feature Alignment is proposed to align each token in the decoder sequence *i.e.*, dynamic query, and to close domain gaps at the object level. Finally, VPC is designed to consider different Transformer decoders as multiple of experts (MoE) to vote for the ensemble prediction, alleviating the domain shifts between the source and target domain. Since predictions from different experts performed differently, we use the voted prediction consistency result as the threshold value for selecting qualified predictions to generate pseudo-labels to further improve the robustness and transferability.

In summary, our contributions are three-fold:

- We present a pioneering framework, Domain Adaptive Point Transformer (DAPoinTr) for the point cloud com-

pletion task. To the best of our knowledge, this is the first work that studies the transferability of the point Transformer (PoinTr) in DA point cloud completion.

- We present Domain Query-based Feature Alignment and Point Token-wise Feature Alignment for aligning sequence features on a global and local level, respectively. Besides, we design a Voted Prediction Consistency to further regularize the sequence features and improve the discriminability of the point transformers.
- Extensive experiments with visualization on several challenging DA benchmarks, *e.g.*, KITTI, ScanNet, and MatterPort for real-world scans, and 3D-FUTURE and ModelNet for synthetic data, demonstrate our method’s superiority compared with other state-of-the-art methods.

## 2 Related Work

**Point Cloud Completion (PCC).** PCC aims to generate complete shapes from partial point clouds. PCN (Yuan et al. 2018) was proposed to use the FoldingNet (Yang et al. 2018) to transform 2D points into 3D surfaces by emulating the deformation process of a 2D plane. Based on this, many techniques (Tchapmi et al. 2019; Xie et al. 2020; Xiang et al. 2021; Wang, Ang, and Lee 2021) have been developed, focusing on achieving higher resolution and enhanced robustness in PCC. For instance, AtlasNet (Groueix et al. 2018) and 3D-EPN (Dai, Ruizhongtai Qi, and Niessner 2017) incorporated convolutional neural networks (CNNs) to reconstruct complete shapes. However, they suffer from limited receptive fields as the layers are deeply stacked. To tackle this, SnowflakeNet (Xiang et al. 2021) was presented by conceptualizing PCC as a progressive growth from specific parent points in a snowflake-like manner, allowing for the more organic development of global structures. Despite their gratifying progress, most of them struggle to alleviate domain gaps when applying them to real-world target domains.

**Domain Adaptation for Point Cloud Completion.** Unsupervised domain adaptation (UDA) (Chen, Chen, and Mitra 2019; Jiang et al. 2024b) has recently been introduced into PCC to narrow domain shifts and improve the adaptability

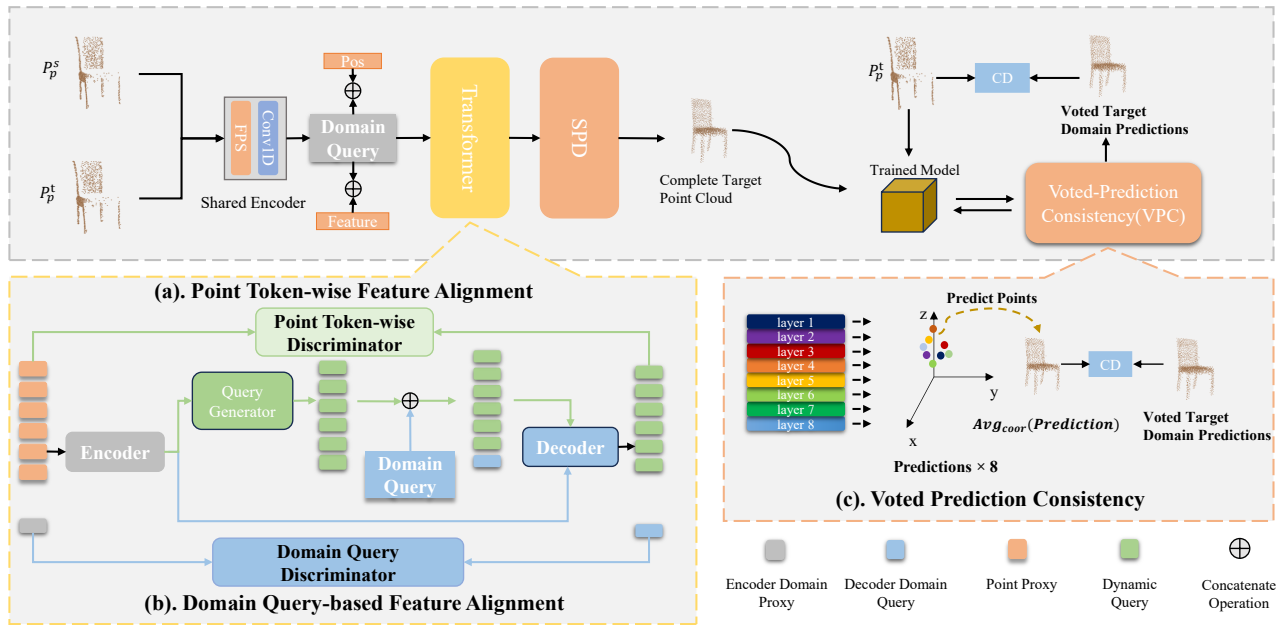


Figure 2: The framework of DAPoinTr for domain adaptive point cloud completion, including three key components: (a) Point Token-wise Feature Alignment (PTFA) is proposed to close the local domain shifts by aligning the tokens, *i.e.*, point proxy and dynamic query, at the Transformer encoder and decoder, respectively. (b) Domain Query-based Feature Alignment (DQFA) is presented to narrow the global domain gaps from the sequence via the presented domain proxy and domain query at the Transformer encoder and decoder, respectively. Finally, Voted Prediction Consistency (VPC) is designed to consider different Transformer decoders as multiple of experts (MoE) to vote for the ensembled prediction and pseudo-label generation.

toward target domains. The mainstream UDA PCC methods aim to learn domain-invariant features and are mainly categorized into adversarial learning (Chen, Chen, and Mitra 2019; Zhang et al. 2021), data reconstruction (Wen et al. 2021), disentangled learning (Gong et al. 2022), self-supervised learning (Hong et al. 2023), and dual mixup consistency regularized learning (Liu et al. 2024). Nonetheless, these approaches heavily rely on Convolution Neural Networks (CNNs) and tend to learn local features, which can be problematic for capturing the broader global context required for learning domain-agnostic features. Thus, it is in urgent need to investigate UDA PCC with global modeling and high transferability based on vision Transformers.

**Vision Transformers.** Transformers (Vaswani et al. 2017) have proven to be very effective in modeling long-range dependencies and have been widely explored in various vision tasks (Dosovitskiy et al. 2020; Liu et al. 2021; Carion et al. 2020; Jiang et al. 2024a; Zhou et al. 2023; He et al. 2021). Thanks to both self-attention and cross-attention mechanisms, Transformers have a strong capability in sequence modeling and improving the information interaction between the Transformer encoder and the decoders. Recently, PoinTr (Yu et al. 2021) adopted a Transformer encoder-decoder architecture and has been presented to reformulate PCC as a set-to-set translation problem, showing excellent performance within specific domains. Unfortunately, no relevant works have investigated the transferability of point Transformer toward various target domains. To the best of our knowledge, this is the first work that studies

the transferability of point Transformer in UDA PCC.

### 3 Methodology

In this section, we concentrate on unsupervised domain adaptation (UDA) for point cloud completion (PCC), where the training dataset comprises paired partial-complete point clouds from the labeled source domain and partial input from the unlabeled target domain. The primary objective is to narrow the domain shift between the source and the target domain. To this end, we present Domain Adaptive Point Transformer (DAPoinTr) framework for UDA point cloud completion that aligns the sequence-wise point features with Transformer encoder-decoder architecture to enhance the adaptability toward the real target domains. As shown in Figure 2, DAPoinTr consists of three key components: Domain Query-based Feature Alignment (DQFA), Point Token-wise Feature alignment (PTFA), and Voted Prediction Consistency (VPC). Specifically, DQFA is presented to narrow the global domain gaps from the sequence via the presented domain proxy and domain query at the Transformer encoder and decoder, respectively. Secondly, PTFA is proposed to close the local domain shifts by aligning the tokens, *i.e.*, point proxy and dynamic query, at the Transformer encoder and decoder, respectively. Finally, VPC is designed to consider different Transformer decoders as multiple of experts (MoE) to vote for the ensembled prediction and pseudo-label generation.

### 3.1 Domain Query-based Feature Alignment

To ensure domain invariance in the sequence features of point cloud, we propose Dynamic Query-based Feature Alignment (DQFA) to align source and target features from a global perspective. Specifically, on the Transformer encoder side, a *domain proxy* embedding denoted as  $q_d^{enc}$  is defined to encode the global context of the sequence, aggregates useful domain-specific features of the entire token sequences, and narrows the global domain gaps in layout and inter-sketch relationships of the point cloud. During the encoding phase, *domain proxy*  $q_d^{enc}$  is integrated with point proxy sequences and then fed forwarded to the Transformer encoder for adversarial alignment. Specifically, the initial input  $\mathcal{X}_E$  is constructed by concatenating  $q_d^{enc}$  with the sequence of point proxy,  $f^i$  and then adding the position encoding  $Pos_d^{enc}$  generated by the DGCNN. This can be formulated as follows:

$$\mathcal{X}_E = [q_d^{enc}; f^1, f^2, \dots, f^N] + [Pos_d^{enc}; pos^1, pos^2, \dots, pos^N] \quad (1)$$

To align domain distributions, we feed the domain proxy into the domain-query discriminator and expect that the learned features of the domain proxy cannot be distinguished by the discriminator.  $\lambda$  denotes the domain label (0 for the source domain and 1 for the target domain):

$$\mathcal{L}_{enc_q} = \lambda \log D_{enc_q}(\mathcal{X}_E) + (1-\lambda) \log(1 - D_{enc_q}(\mathcal{X}_E)) \quad (2)$$

Similarly, a *domain query*  $q_d^{dec}$  is presented to encode some sketch relationships for adaptation and fuse the context features from each dynamic query generated by the Query Generator in the sequence.  $q_d^{dec}$  is concatenated with dynamic queries to construct  $X_D$  before being fed into the Transformer decoder, which can be formulated as follows:

$$\mathcal{X}_D = [q_d^{dec}; f'^1, f'^2, \dots, f'^N] + [pos_d'^{dec}; pos'^1, pos'^2, \dots, pos'^N] \quad (3)$$

where  $f'$  means the global feature and  $pos'$  represents the position embedding generated from Transformer encoder. During the decoding phase, we input the domain query to the aforementioned discriminator to align distributions:

$$\mathcal{L}_{dec_q} = \lambda \log D_{dec_q}(\mathcal{X}_D) + (1-\lambda) \log(1 - D_{dec_q}(\mathcal{X}_D)) \quad (4)$$

### 3.2 Point Token-wise Feature Alignment

Although the Domain Query-based Feature Alignment (DQFA) can effectively bridge domain gaps from a global perspective, it is less adept at tackling the domain shifts associated with local shape. To address this limitation, we introduce a Point Token-wise Feature Alignment (PTFA) to close the local domain shifts by aligning the tokens at the Transformer encoder and decoder, respectively.

Concretely, Point Proxy Feature Alignment (PPFA) is presented to align each token in the Transformer encoder sequence, *i.e.*, *point proxy*, to alleviate domain gaps caused by local shape, appearance, *etc.* Specifically, each point proxy in the Transformer encoder sequence is fed into the Point Token-wise Discriminator for adversarial feature alignment:

$$\mathcal{L}_{enc_k} = -\frac{1}{N} \sum_{i=1}^N (\lambda \log D_{enc_k}(X_p) + (1-\lambda) \log(1 - D_{enc_k}(X_p))) \quad (5)$$

where  $X_p$  indicates the point proxies and  $\lambda$  indicates domain label (0 for the source domain and 1 for the target domain).

In addition, Dynamic Query Feature Alignment (DQFA) is proposed to align the *dynamic query* at the object levels, thus narrowing the domain gaps on the Transformer decoder side via the Point Token-wise Discriminator  $D_{dec_k}$ :

$$\mathcal{L}_{dec_k} = -\frac{1}{N} \sum_{i=1}^N (\lambda \log D_{dec_k}(X_q) + (1-\lambda) \log(1 - D_{dec_k}(X_q))) \quad (6)$$

where  $X_p$  represents the dynamic query and  $\lambda$  indicates domain label (0 for source domain and 1 for target domain).

**Remarks.** Although the point token-wise sequence feature alignment is adopted on both the encoder and decoder side of point Transformers, it has different implications. Concretely, each token in the Transformer encoder sequence, *i.e.*, point proxy, encodes the local features, and thus Point Proxy Feature Alignment (PPFA) focuses on narrowing the domain gaps caused by that local shape, appearance *etc.* In contrast, each token in the Transformer decoder sequence, *i.e.*, dynamic queries, captures the local structural information at the object level, and therefore, Dynamic Query Feature Alignment (DQFA) concentrates more on alleviating the domain shifts caused by sketch information.

### 3.3 Voted Prediction Consistency

Due to the various performances of the point cloud predictions at different layers within the Transformer decoder, we introduce Voted Prediction Consistency (VPC) that considers different Transformer decoders as multiple of experts (MoE) to vote for the ensembled prediction, alleviating the domain shifts between the source and target domain.

To be specific, the coordinates  $Pred_l$  of the point clouds predicted by each decoder layer are meticulously preserved, and the mean of these coordinates  $M_{mean}$  is calculated within the 3D coordinate system. This process enforces a consistency constraint across the outputs of each decoder layer to optimize the quality of the recovered shape.

$$\mathcal{L}_{cons} = CD(M_{mean}, Pred_l) \quad (7)$$

In addition, we use the results from the Voted Prediction Consistency as a threshold criterion for selecting high-quality predictions to generate pseudo-labels to further improve the robustness and transferability.

This process serves to further regularize the sequence features, enhancing the discriminability and generalization capabilities of our proposed approach. By systematically refining the quality of pseudo-labels incorporated into the training process, we ensure that only the most representative and accurate predictions are selected to alleviate domain gaps and improve the performance of target samples. Therefore, the total loss function can be formulated as follows:

$$\mathcal{L}_{total} = \alpha(\mathcal{L}_{enc_q} + \mathcal{L}_{dec_q}) + \beta(\mathcal{L}_{enc_k} + \mathcal{L}_{dec_k}) + \gamma \mathcal{L}_{cons} \quad (8)$$

where  $\alpha$ ,  $\beta$  and  $\gamma$  are the weights balancing these losses.

Methods	Avg	Cabinet	Chair	Lamp	Sofa	Table
Pcl2Pcl	92.83	57.23	43.91	157.86	63.23	141.92
ShapeInv.	53.21	38.54	26.30	48.57	44.02	108.60
Cycle4Co.	45.39	32.62	34.08	77.19	43.05	40.00
ACL-SPC	35.97	70.12	23.87	31.75	28.74	25.38
OptDE	28.99	28.37	21.87	29.92	37.98	26.81
DAPoinTr (Ours)	<b>22.35</b>	<b>18.46</b>	<b>17.60</b>	<b>27.91</b>	<b>23.08</b>	<b>24.71</b>

Table 1: Cross-domain completion results on 3D-FUTURE. We take Chamfer Distance (CD) $\downarrow$  as the metric to evaluate the performance, and the scale factor is  $10^4$ . Lower is better.

## 4 Experiments

### 4.1 Experimental Setting

**Datasets.** Following protocols of previous UDA PCC methods (Chen, Chen, and Mitra 2019; Gong et al. 2022), we use 3D data from CRN (Wang, Ang Jr, and Lee 2020) as the source domain, and the datasets, including Real-World Scans, 3D-FUTURE (Fu et al. 2021) and ModelNet (Wu et al. 2015), as target domains. As for *CRN* dataset, we take 26,863 samples of CRN from shared categories between CRN and other datasets for DA. For the benchmark of *Real-World Scans*, we evaluate the performance on ScanNet (Dai et al. 2017), MatterPort3D (Chang et al. 2017), and KITTI (Geiger, Lenz, and Urtasun 2012), where the tables and chairs in ScanNet and MatterPort3D, and cars in KITTI are used for evaluation. We re-sample the input scans to 2,048 points for unpaired training and inference to match the virtual dataset. As for *3D-FUTURE*, it only contains indoor furniture models that appear like real objects. We obtain partial shapes and complete ones with 2,048 points from 5 different viewpoints as the target domain. As for *ModelNet*, a subset extracted from the original ModelNet40 dataset (Wu et al. 2015), 2,048 points are taken for both partial and complete shapes to match CRN, and 6 shared categories between ModelNet40 and CRN are used for evaluation.

**Implementation Details.** All experiments were conducted on an RTX 4090 with 64GB RAM. We use the refinement module of PoinTr (Yu et al. 2021) as the backbone of our PCC network. For the training, we employ the initial learning rate of  $2 \times 10^{-4}$  and a weight decay of  $5 \times 10^{-5}$ . The batch size is set to 2. To balance losses, weights of  $\alpha$ ,  $\beta$ , and  $\gamma$  are set as 0.025, 0.25, and 0.01 respectively. Following previous works, we adopt Unidirectional Chamfer Distance (UCD), Unidirectional Hausdorff Distance (UHD), and Chamfer Distance (CD) as the metrics for evaluation.

### 4.2 Comparison Results

In this section, we conduct extensive experimental comparisons on the widely-used real-world and synthetic benchmarks, including KITTI, ScanNet, and MatterPort for real-world scans, and 3D-FUTURE and ModelNet for synthetic data, to evaluate the performance of our method.

**Comparison Methods.** We compare our method with several state-of-the-art UDA point cloud completion techniques, including Pcl2Pcl (Chen, Chen, and Mi-

Methods	Avg	Plane	Car	Chair	Lamp	Sofa	Table
Pcl2Pcl	68.14	18.53	17.54	43.58	126.80	38.78	163.62
ShapeInv.	41.61	3.78	15.66	22.25	60.42	22.25	125.31
Cycle4Co.	28.65	5.77	11.85	26.67	83.34	22.82	21.47
ACL-SPC	34.89	5.75	11.73	43.08	106.29	25.62	16.89
OptDE	15.94	<b>2.18</b>	9.80	14.71	39.74	19.43	<b>9.75</b>
DAPoinTr (Ours)	<b>13.79</b>	2.38	<b>8.04</b>	<b>13.83</b>	<b>33.26</b>	<b>12.72</b>	12.51

Table 2: Cross-domain completion results on ModelNet. We take Chamfer Distance (CD) $\downarrow$  as the metric to evaluate the performance, and the scale factor is  $10^4$ . Lower is better.

Methods	ScanNet		MatterPort3D		KITTI
	Chair	Table	Chair	Table	Car
Pcl2Pcl	17.3/10.1	9.1/11.8	15.9/10.5	6.0/11.8	9.2/14.1
ShapeInv.	3.2/10.1	3.3/11.9	3.6/10.0	3.1/11.8	2.9/13.8
Cycle4Co.	5.1/6.4	3.6/5.9	8.0/8.4	4.2/6.8	3.3/5.8
ACL-SPC	1.4/4.7	1.8/5.1	1.8/4.8	2.1/4.9	2.0/4.9
OptDE	2.6/5.5	1.9/4.6	3.0/5.5	1.9/5.3	1.6/3.5
DAPoinTr (Ours)	<b>1.1/2.7</b>	<b>0.96/2.7</b>	<b>1.3/2.9</b>	<b>1.2/2.8</b>	<b>0.45/1.8</b>

Table 3: Cross-domain completion results on Real-World Scans. We take [UCD $\downarrow$ /UHD $\downarrow$ ] as the metric to evaluate the performance, and the scale factor is  $10^4$  and  $10^2$  for UCD and UHD, respectively. UCD: Unidirectional Chamfer Distance. UHD: Unidirectional Hausdorff Distance. A lower value of UCD or UHD is better.

tra 2019), ShapeInversion (Zhang et al. 2021), Cycle4Completion (Wen et al. 2021), Optde (Gong et al. 2022) and ACL-SPC (Hong et al. 2023).

**Results on 3D FUTURE dataset.** As illustrated in Table 1, our DAPoinTr shows substantial superiority to state-of-the-art UDA PCC approaches on the 3D-FUTURE dataset, achieving significantly better performance across all categories. Notably, we achieved an average improvement of 6.64 in Chamfer Distance (CD) across all categories compared to the 2nd best method OptDE. Additionally, unlike previous UDA PCC methods, which often struggle to bridge domain gaps, especially in cabinet and sofa categories, our results indicate substantial enhancements in these categories, as evidenced by significant margins of improvement (see the 3rd and 6th columns in Table 1). This is because our approach can better tackle domain discrepancies by aligning feature distribution between source and target domain, as well as leveraging the capacity of Transformer to capture both global and local features.

**Results on ModelNet dataset.** Table 2 demonstrates that our method outperforms state-of-the-art methods in most categories on ModelNet. Specifically, DAPoinTr achieves a remarkable improvement in the category of Lamp and Sofa in terms of the mean Chamfer Distance. This superiority indicates that the innovative designs of our model can effectively alleviate domain shifts and close domain gaps at the object level. This remarkable improvement underscores the robustness and adaptability of the innovative design elements incorporated within our model.

Methods	PTFA	DQFA	VPC	Cabinet	Chair	Lamp	Sofa	Table	Avg
Baseline* (PoinTr)				32.88	24.21	41.76	33.76	35.27	33.58
Our DAPoinTr*	✓	✓	✓	31.62	20.01	29.01	32.97	28.72	28.47
Baseline (PoinTr w SPD)				20.82	21.68	29.43	24.68	32.09	25.74
Our DAPoinTr	✓			20.65	18.77	32.26	24.39	26.67	24.55
		✓		20.37	18.24	29.80	25.02	26.13	23.91
	✓	✓		20.21	18.11	29.13	23.84	25.01	23.26
	✓	✓	✓	<b>18.46</b>	<b>17.60</b>	<b>27.91</b>	<b>23.08</b>	<b>24.71</b>	<b>22.35</b>

Table 4: Ablation studies on each proposed module of our DAPoinTr on 3D-FUTURE dataset. We take Chamfer Distance (CD) $\downarrow$  as the metric to evaluate the performance, and the scale factor is  $10^4$ . Lower is better.

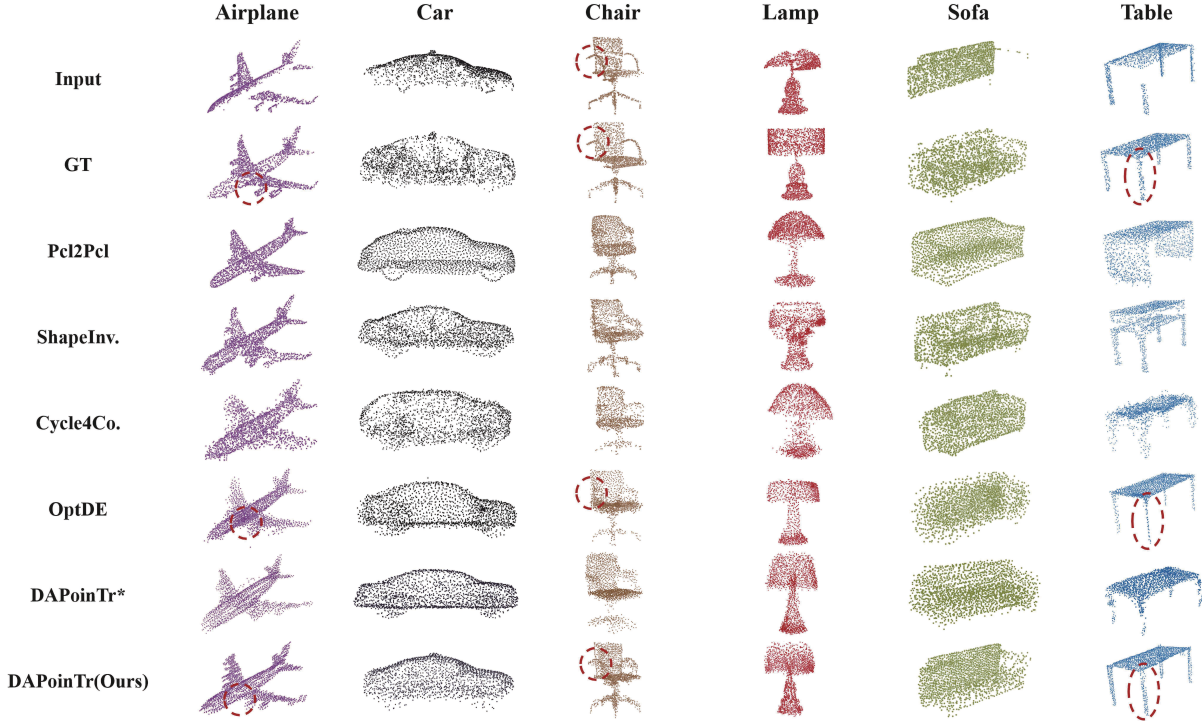


Figure 3: Visualization comparisons with state-of-the-art PCC methods on ModelNet dataset.

**Results on Real-World Scans.** In this benchmark, we utilize the UCD and UHD as evaluation metrics. As shown in Table 3, it is obvious that our approach outperforms previous methods across all real-world scan datasets in terms of UCD and UHD metrics by a large margin, which indicates our model can better maintain the overall geometry of partial inputs and effectively align the feature distribution and bridge the domain gaps in output predictions.

### 4.3 Ablation Studies

In this section, we conduct ablation studies to verify the contribution of each proposed module, as shown in Table 4. In all ablation experiments, we use the CRN dataset as the source domain and 3D-FUTURE as the target domain.

Initially, we employed the original PoinTr (Yu et al. 2021) (denoted as Baseline\*) to train on the source domain and

directly tested it on the target domain. However, we observe that the original PoinTr produces undesirable performance (CD: 33.58) on the target domain due to the lack of learning domain-invariant sequence features. In contrast, by adding our proposed components into Baseline\*, our DAPoinTr\* significantly improves transferability and manages to achieve a better completion performance (CD: 28.47).

Considering that PoinTr employs FoldingNet (Yang et al. 2018) to decode the recovered complete point cloud, we observed that FoldingNet struggles with effective recovery when handling complex shapes. Inspired by SnowFlakeNet (Xiang et al. 2021), we replace the FoldingNet in PoinTr with the SPD decoder from SnowFlakeNet (denoted as Baseline) to reconstruct the complete point cloud. The results in the 3rd row of Table 4 confirm that the SPD decoder demonstrates better performance, making it a competent so-

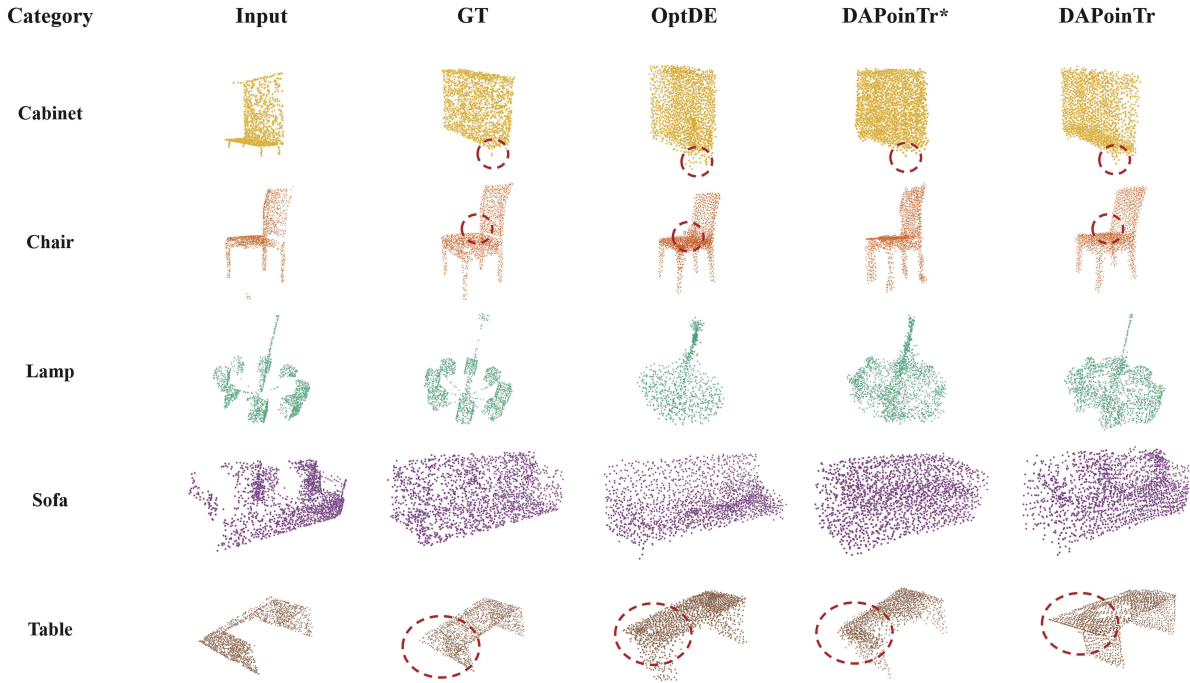


Figure 4: Visualization comparisons with state-of-the-art PCC methods on 3D-FUTURE dataset.

lution for complex geometry.

Based on this superior baseline, we conduct detailed ablation studies to reveal the impact of each proposed component. By adding PTFA into the baseline, we effectively decrease the average performance from 25.74 to 24.55, demonstrating the effectiveness of narrowing local domain shifts. By further adding DQFA, we achieve 23.91 of the average CD performance, demonstrating the effectiveness in narrowing global domain shifts. Besides, when using both PTFA and DQFA, the results are much better than using merely one. In challenging categories such as sofas and chairs, as shown in Table 4, the integration of PTFA and DQFA effectively diminishes the impact of domain gaps and promotes the learning of domain-invariant sequence features. This is because each module is indispensable in enhancing the model’s proficiency in capturing both local details and global structures, and in boosting the model’s adaptability in the realm of UDA PCC. Finally, when further adding VPC, the average CD result is further decreased to 22.35. The improvement is particularly notable in the categories of cabinets and tables, which show significant enhancements over previous UDA PCC methods. These results confirm that these individual components are complementary and together they significantly promote the performance.

#### 4.4 Visualization and Analysis

##### Visualization of Completed Point Clouds on ModelNet.

As depicted in Figure 3, we visualize the qualitative results on ModelNet. When comparing the visualization results of our DAPoinTr with the previous state-of-the-art method, OptDE, our approach consistently generates more accurate

completions with superior fine-grained local details. This improvement is particularly evident in the cases of chairs and tables, as highlighted in the second and final rows of the figure. These results demonstrate the enhanced capability of our method in rendering more detailed and realistic structures in complex object categories.

**Visualization of Completed Point Clouds on 3D-FUTURE.** Figure 4 visualizes completed point clouds on the 3D-FUTURE dataset. Our approach excels in capturing both local details and global structures. Our DAPoinTr consistently reconstructs more accurate complete shapes, maintaining superior structural integrity and a more even distribution of points compared to other state-of-the-art methods. This demonstrates our robustness to delivering better reconstructions across various complex scenarios.

## 5 Conclusion

In this paper, we introduce Domain Adaptive Point Transformer (DAPoinTr), a novel and pioneering framework aimed at enhancing the transferability of point Transformers (PoinTr) for point cloud completion across varying domains. Through the integration of several new modules including Domain Query-based Feature Alignment (DQFA), Point Token-wise Feature Alignment (PTFA), and Voted Prediction Consistency (VPC), our DAPoinTr effectively addresses both global and local domain discrepancies, resulting in superior structural integrity and point distribution. Extensive experiments on the synthetic and real datasets demonstrate remarkable improvements over existing methods, particularly in complex object categories.

## References

- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, 213–229.
- Chang, A.; Dai, A.; Funkhouser, T.; Halber, M.; Niessner, M.; Savva, M.; Song, S.; Zeng, A.; and Zhang, Y. 2017. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*.
- Chen, X.; Chen, B.; and Mitra, N. J. 2019. Unpaired point cloud completion on real scans using adversarial training. *arXiv preprint arXiv:1904.00069*.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5828–5839.
- Dai, A.; Ruizhongtai Qi, C.; and Niessner, M. 2017. Shape Completion Using 3D-Encoder-Predictor CNNs and Shape Synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5868–5877.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fu, H.; Jia, R.; Gao, L.; Gong, M.; Zhao, B.; Maybank, S.; and Tao, D. 2021. 3d-future: 3d furniture shape with texture. *International Journal of Computer Vision*, 129: 3313–3337.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3354–3361.
- Gong, J.; Liu, F.; Xu, J.; Wang, M.; Tan, X.; Zhang, Z.; Yi, R.; Song, H.; Xie, Y.; and Ma, L. 2022. Optimization over disentangled encoding: Unsupervised cross-domain point cloud completion via occlusion factor manipulation. In *European Conference on Computer Vision*, 517–533.
- Groueix, T.; Fisher, M.; Kim, V. G.; Russell, B. C.; and Aubry, M. 2018. A papier-mâché approach to learning 3d surface generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 216–224.
- He, L.; Zhou, Q.; Li, X.; Niu, L.; Cheng, G.; Li, X.; Liu, W.; Tong, Y.; Ma, L.; and Zhang, L. 2021. End-to-end video object detection with spatial-temporal transformers. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1507–1516.
- Hong, S.; Yavartanoo, M.; Neshatavar, R.; and Lee, K. M. 2023. 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*, 9435–9444.
- Huang, Z.; Yu, Y.; Xu, J.; Ni, F.; and Le, X. 2020. PF-Net: Point Fractal Network for 3D Point Cloud Completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7662–7670.
- Jiang, J.; Zhou, Q.; Li, Y.; Lu, X.; Wang, M.; Ma, L.; Chang, J.; and Zhang, J. J. 2024a. DG-PIC: Domain Generalized Point-In-Context Learning for Point Cloud Understanding. In *European Conference on Computer Vision*, 455–474.
- Jiang, J.; Zhou, Q.; Li, Y.; Zhao, X.; Wang, M.; Ma, L.; Chang, J.; Zhang, J. J.; and Lu, X. 2024b. PCoTTA: Continual Test-Time Adaptation for Multi-Task Point Cloud Understanding. *Advances in Neural Information Processing Systems*.
- Liu, F.; Gong, J.; Zhou, Q.; Lu, X.; Yi, R.; Xie, Y.; and Ma, L. 2024. CloudMix: Dual Mixup Consistency for Unpaired Point Cloud Completion. *IEEE Transactions on Visualization and Computer Graphics*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Phan, A. V.; Le Nguyen, M.; Nguyen, Y. L. H.; and Bui, L. T. 2018. Dgcnn: A convolutional neural network over large-scale labeled graphs. *Neural Networks*, 108: 533–543.
- 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.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Wang, X.; Ang, M. H.; and Lee, G. H. 2021. Voxel-based network for shape completion by leveraging edge generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13189–13198.
- Wang, X.; Ang Jr, M. H.; and Lee, G. H. 2020. Cascaded refinement network for point cloud completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 790–799.
- Wen, X.; Han, Z.; Cao, Y.-P.; Wan, P.; Zheng, W.; and Liu, Y.-S. 2021. 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*, 13080–13089.
- 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.
- 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.
- Xie, H.; Yao, H.; Zhou, S.; Mao, J.; Zhang, S.; and Sun, W. 2020. Grnet: Gridding residual network for dense point cloud completion. In *European Conference on Computer Vision*, 365–381.

- Yang, Y.; Feng, C.; Shen, Y.; and Tian, D. 2018. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 206–215.
- 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.
- Yu, X.; Rao, Y.; Wang, Z.; Lu, J.; and Zhou, J. 2023. AdapointR: Diverse point cloud completion with adaptive geometry-aware transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45: 14114–14130.
- Yuan, W.; Khot, T.; Held, D.; Mertz, C.; and Hebert, M. 2018. Pcn: Point completion network. In *International Conference on 3D Vision (3DV)*, 728–737.
- Zhang, J.; Chen, X.; Cai, Z.; Pan, L.; Zhao, H.; Yi, S.; Yeo, C. K.; Dai, B.; and Loy, C. C. 2021. Unsupervised 3d shape completion through gan inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1768–1777.
- Zhou, Q.; Li, X.; He, L.; Yang, Y.; Cheng, G.; Tong, Y.; Ma, L.; and Tao, D. 2023. TransVOD: end-to-end video object detection with spatial-temporal transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6): 7853–7869.