

PointCG

Self-supervised Point Cloud Learning via Joint Completion and Generation

Liu, Yun; Li, Peng; Yan, Xuefeng; Nan, Liangliang; Wang, Bing; Chen, Honghua; Gong, Lina; Zhao, Wei; Wei, Mingqiang

DOI

[10.1109/TVCG.2025.3526257](https://doi.org/10.1109/TVCG.2025.3526257)

Publication date

2025

Document Version

Final published version

Published in

IEEE Transactions on Visualization and Computer Graphics

Citation (APA)

Liu, Y., Li, P., Yan, X., Nan, L., Wang, B., Chen, H., Gong, L., Zhao, W., & Wei, M. (2025). PointCG: Self-supervised Point Cloud Learning via Joint Completion and Generation. *IEEE Transactions on Visualization and Computer Graphics*, 31(10), 6648-6660. <https://doi.org/10.1109/TVCG.2025.3526257>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.

PointCG: Self-Supervised Point Cloud Learning via Joint Completion and Generation

Yun Liu , Peng Li , Xuefeng Yan , Liangliang Nan , Bing Wang , Honghua Chen , Lina Gong , Wei Zhao , and Mingqiang Wei , *Senior Member, IEEE*

Abstract—The core of self-supervised point cloud learning lies in setting up appropriate pretext tasks, to construct a pre-training framework that enables the encoder to perceive 3D objects effectively. In this article, we integrate two prevalent methods, masked point modeling (MPM) and 3D-to-2D generation, as pretext tasks within a pre-training framework. We leverage the spatial awareness and precise supervision offered by these two methods to address their respective limitations: ambiguous supervision signals and insensitivity to geometric information. Specifically, the proposed framework, abbreviated as PointCG, consists of a Hidden Point Completion (HPC) module and an Arbitrary-view Image Generation (AIG) module. We first capture visible points from arbitrary views as inputs by removing hidden points. Then, HPC extracts representations of the inputs with an encoder and completes the entire shape with a decoder, while AIG is used to generate rendered images based on the visible points' representations. Extensive experiments demonstrate the superiority of the proposed method over the baselines in various downstream tasks. Our code will be made available upon acceptance.

Index Terms—PointCG, self-supervised learning, hidden point completion, arbitrary-view image generation, point clouds.

I. INTRODUCTION

SELF-SUPERVISED representation learning (SSRL) aims to fully exploit the statistical and structural knowledge

Received 8 May 2024; revised 14 December 2024; accepted 23 December 2024. Date of publication 6 January 2025; date of current version 5 September 2025. This work was supported in part by the National Natural Science Foundation of China under Grant T2322012, Grant 62172218, and Grant 62032011, in part by the National Defense Basic Scientific Research Program of China under Grant JCKY2020605C003, in part by the Shenzhen Science and Technology Program under Grant JCYJ20220818103401003 and Grant JCYJ20220530172403007, and in part by the Guangdong Basic and Applied Basic Research Foundation under Grant 2022A1515010170. Recommended for acceptance by D. Ritchie. (*Corresponding author: Xuefeng Yan.*)

Yun Liu, Peng Li, Honghua Chen, Lina Gong, Wei Zhao, and Mingqiang Wei are with the School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, Nanjing 210095, China, and also with the Shenzhen Institute of Research, Nanjing University of Aeronautics and Astronautics, Shenzhen 518055, China (e-mail: yun.liu.lydia@gmail.com; pengl@nuaa.edu.cn; chenhonghua.cn@gmail.com; gonglina@nuaa.edu.cn; weizhao0120@nuaa.edu.cn; mingqiang.wei@gmail.com).

Xuefeng Yan is with the School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, Nanjing 210095, China, and also with the Collaborative Innovation Center of Novel Software Technology and Industrialization, Nanjing 210000, China (e-mail: fzjm_313@nuaa.edu.cn).

Liangliang Nan is with the Urban Data Science Section, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: lianliang.nan@gmail.com).

Bing Wang is with the Department of Aeronautical and Aviation Engineering, Hong Kong Polytechnic University, Hongkong, China (e-mail: bingwang@polyu.edu.hk).

Digital Object Identifier 10.1109/TVCG.2025.3526257

inherent in unlabeled datasets, enabling the encoder of the pre-training model to extract informative and discriminative representations. The pre-trained encoder can be subsequently applied to various downstream tasks such as classification, denoising [1], [2], segmentation [3], and object detection [4]. The core of SSRL lies in the design of appropriate pretext tasks aimed at aiding the encoder in achieving a full perception and understanding of the inputs.

Based on the tasks employed, existing self-supervised pre-training methods can be broadly classified into two paradigms: contrastive learning and generative learning, both of which have attained great success in processing 2D images [6], [7], [8] and 3D point clouds [5], [9], [10], [11], [12], [13]. Compared to contrastive learning, generative learning is considered as a more data-efficient pre-training method, capable of capturing the patterns of the inputs with relatively limited data volume [14]. Therefore, it is highly favored in the context of data scarcity within the field of 3D vision, where masked point modeling [5], [9], [10], [12], [15] and 3D-to-2D generation [13], [16] stand out as two representative generative learning methods. Among them, masked point modeling drives the model to predict arbitrary missing parts based on the remaining points. Accomplishing this task requires a thorough understanding of the spatial properties and global-local context of point clouds. 3D-to-2D generation employs a cross-modal pretext task which translates a 3D object point cloud to its diverse forms of 2D rendered images (e.g., silhouette, depth, contour). Pre-training with pixel-wise precise supervision drives the backbone to perceive the fine-grained edge details of 3D objects.

However, both of the above methods have their own limitations. As revealed in [17], [18], [19], due to the irregularity of point clouds, commonly used point set similarity metrics (e.g., Chamfer Distance and Earth Mover's Distance) in masked point modeling cannot provide explicit point-to-point supervision between ground truth and generated point clouds. The lack of precise correspondence results in limited feature representation capability of the pre-trained backbone network. Conversely, 3D-to-2D generation [13], [16] alleviates the issue of insufficient supervision signals by utilizing regular 2D images as the generation objective, offering pixel-wise precise supervision. However, relying solely on images from limited views as ground truth may overlook the structural information from occluded point sets, diminishing the backbone's perception of the spatial properties of point clouds. As shown in Fig. 1, masked point modeling exhibits subpar performance in reconstructing

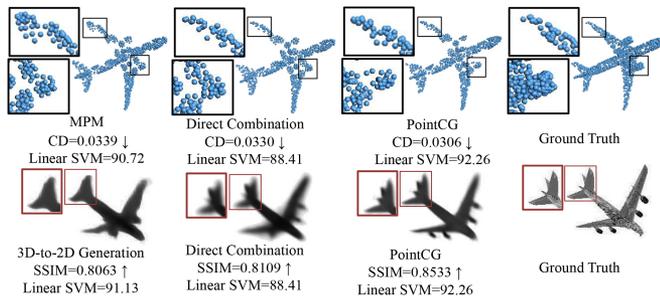


Fig. 1. Qualitative and quantitative comparison of models using different pretext tasks. Chamfer Distance (CD) and Structural Similarity Index (SSIM) are employed as the quantitative metrics. For the masked point modeling (MPM) task, we utilize the method proposed in Point-MAE [5] with the inputs of visible points from arbitrary views (see Section III-A). For the 3D-to-2D generation task, we define the pretext task as generating images from arbitrary views. The result of the model using only MPM exhibits group clustering at the edges, while our method yields sharpened and clear edges that closely align with the ground truth. The model relying solely on 3D-to-2D generation fails to capture three-dimensional structural information, while our method can effectively preserve the geometric structure. Directly combining both tasks generates point clouds and images superior to using only MPM or 3D-to-2D generation (Direct Combination) but with lower Linear-SVM accuracy.

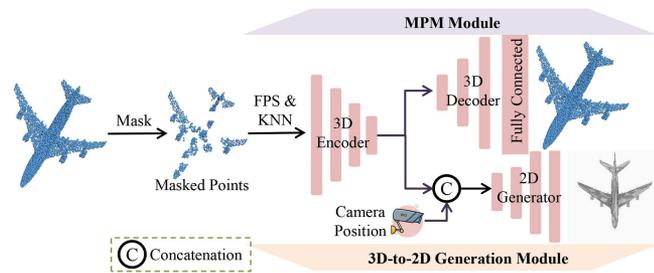


Fig. 2. The framework of direct combination integrates MPM and 3D-to-2D generation by sharing a common encoder for feature extraction and using separate decoders for 3D reconstruction and 2D generation tasks.

some challenging areas (e.g., edges) due to the lack of point-to-point supervision. Besides, 3D-to-2D generation yields images lacking three-dimensional structural information, attributed to the lack of explicit geometric guidance. These observations collectively indicate the models' inadequate perception of the inputs, consequently reducing their performance on downstream tasks.

Based on the aforementioned analysis, an intuitive method is to combine these two pretext tasks (Fig. 2) to retain their individual merits while compensating for their respective limitations. However, as shown in Fig. 1, while the model directly combining both tasks outperforms those relying solely on MPM or 3D-to-2D generation in generating high-quality point clouds or images, its Linear-SVM accuracy is lower (88.41% versus 90.72% and 91.13%). We argue that the encoder's involvement in concurrently generating content for two modalities leads to confusion. Furthermore, to accomplish both tasks, the model shifts its training focus toward the decoder, which is typically discarded after pre-training. This phenomenon diminishes the feature extraction capability of the encoder, ultimately reducing the Linear SVM accuracy.

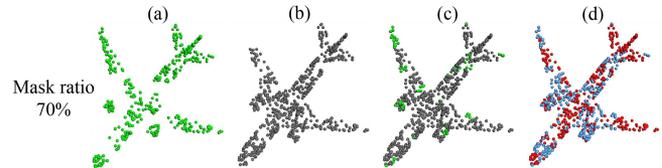


Fig. 3. Visualization of the unmasked points (a), the masked points (b), the completed point cloud composed of green unmasked points and gray masked points (c), and the completed point cloud in blue with overlapping points highlighted in red (d).

To address these issues, we propose PointCG, a framework that effectively integrates masked point modeling and 3D-to-2D generation tasks. This framework incorporates a **Hidden Point Completion (HPC)** module and an **Arbitrary-view Image Generation (AIG)** module. Existing MAE-based MPM methods often employ a random masking strategy based on Farthest Point Sampling (FPS) and K-Nearest Neighbor (KNN) techniques. However, the inputs of unmasked patches (Fig. 3(a)) preserve the overall shape of an object and exhibit substantial overlap with the target points (highlighted in red in Fig. 3(d)). The leakage of overall structure and point location information enables the model to reconstruct the object without a holistic comprehension of the entire structure, which limits the learning capacity of the encoder during pre-training. To overcome this limitation, we select the visible points from arbitrary views by removing hidden points as input and introduce the HPC module to complete the point clouds. For the 3D-to-2D generation task, we employ the arbitrary-view image generation as the pretext task, which generates the image from an arbitrary view based on the representations of visible points extracted by the encoder. Furthermore, the cross-modal feature alignment is introduced to align the feature spaces of point clouds and images, which enables simultaneous content generation across both modalities and refocuses the training on the encoder. Specifically, we extract features from both the input point clouds and their corresponding rendered 2D images, encouraging feature proximity for the same instance while maintaining feature separation for different instances.

Through the effective integration of HPC and AIG, the pre-trained encoder achieves a comprehensive understanding of 3D objects and can extract high-quality 3D representations. We evaluate our model and the proposed modules with a variety of downstream tasks and ablation studies. We further demonstrate that informative representations can be effectively learned from the restricted points, and such representations facilitate effortless masked point modeling and arbitrary-view image generation.

II. RELATED WORK

A. Self-Supervised Representation Learning

Self-supervised representation learning aims to derive robust and general representations from unlabeled datasets, which can be broadly classified into two categories based on the types of pretext tasks: contrastive learning and generative learning.

Contrastive learning-based methods (e.g., BYOL [20], SimSiam [8], DINO [21], STRL [11], CrossPoint [22]) define the

augmented views of a sample as positive samples, while considering other instances as negative samples, thereby constructing discriminative tasks. Generative learning-based methods (e.g., GPT [23], Point-BERT [9], Point-MAE [5], OcCo [12], MaskFeat3D [24]) are based on the intuition that effective feature abstractions contain sufficient information to reconstruct the original geometric structures [13]. In the point cloud processing community, where 3D assets are relatively scarce, generative learning has garnered widespread attention due to its data efficiency [5], [6], [15]. Among them, MAE stands out as one of the representative paradigms. It involves masking a substantial portion of input data, followed by the use of an encoder to extract informative representations and a decoder to reconstruct explicit features (e.g., pixels or points) or implicit features (e.g., discrete tokens). Taking Point-BERT [9], MaskFeat3D [24], and IAE [25] as examples, each of these methods utilizes the visible groups as input after masking and reconstructs the positions of masked points, surface normals, and surface variations, as well as the implicit features of the masked points. However, after random masking [5] or partial occlusion [15], the visible groups often retain the overall structure of the object (Fig. 8), and there are substantial overlap regions between input and target patches (Fig. 3). The leakage of overall structure and point location information will reduce the difficulty in reconstructing masked patches, thus limiting the learning and inference capabilities of the encoder.

To avoid the leakage of the object’s overall shape and minimize overlap, we simulate scanners to capture visible points from arbitrary views as input. Our approach is conceptually aligned with OcCo [12], which employs the Z-Buffer algorithm [26] to select visible points from multiple views, subsequently completing the original point clouds with an encoder-decoder architecture. The Z-Buffer algorithm addressed within rendering relies on two assumptions: the points satisfy sampling criteria (e.g., Nyquist condition) and the points are associated with normals (or the normals can be estimated) [27]. However, our method seeks rigorous theoretical support for visibility computation without requiring normal estimation, point rendering, or surface reconstruction. Therefore, we employ the Hidden Point Removal (HPR) operator to compute visibility in a more robust manner.

B. Cross-Modal Learning

Recently, cross-modal learning [28] has been a popular research topic, aiming at extracting informative representations from multiple modalities such as images, audio, and point clouds. It has the potential to enhance the performance of various tasks, including image recognition, speech recognition, and point cloud analysis.

In point cloud analysis, a variety of methods have been proposed for cross-modal learning, such as CrossPoint [22], PointMCD [29], TAP [16], and PointVST [13]. CrossPoint [22] establishes cross-modal contrastive learning between images and point clouds, demonstrating that the correspondence between images and points can enhance 3D object understanding. PointMCD [29] obtains a powerful point encoder by aligning

the multi-view visual and geometric descriptors generated by a pretrained image encoder and a learnable point encoder. Both CrossPoint [22] and PointMCD [29] are based on the contrastive paradigm and rely heavily on extensive 3D-2D paired data. Generative methods, such as TAP [16] and PointVST [13], generate images from specific views based on the input point clouds. These methods use regular 2D images as generation objectives to provide precise supervision.

In this paper, we follow the generative learning paradigm and propose a unified pre-training framework with two complementary pretext tasks: hidden point completion and arbitrary-view image generation. The spatial awareness provided by 3D completion addresses the geometric insensitivity inherent in image supervision, as shown in the second line of column one in Fig. 1. Additionally, we demonstrate the mutual enhancement between the two pretext tasks through various experiments.

III. METHODOLOGY

As illustrated in Fig. 4, PointCG mainly consists of a hidden point completion (HPC) module and an arbitrary-view image generation (AIG) module. Specifically, we begin by selecting the visible points from arbitrary views as the inputs (Section III-A), and then introduce an asymmetric Transformer-based encoder-decoder architecture for extracting representations and completing hidden points (Section III-B). Finally, we generate arbitrary-view images (Section III-C) based on the aligned representations extracted by the encoder. In the following, we will delve into the details of these modules.

A. Data Organization

Given a complete point cloud $P = \{p_i | 1 \leq i \leq N\} \in \mathbb{R}^3$, we randomly select the camera position $C = [azimuth, elevation, distance]$, where *distance* is fixed at 1.0. *azimuth* and *elevation* are randomly chosen within the range of $[0, 2\pi]$. The HPR operator [27] is employed to determine whether p_i is visible from C . It mainly consists of two steps: inversion transformation and convex hull construction.

Inversion transformation: We employ spherical flip [30] to reflect each point $p_i \in P$ to the spherical surface (denoted as $\hat{p}_i \in \hat{P}$) along the ray from C through p_i to the spherical surface, as illustrated in Fig. 5(a).

Convex hull construction: The visible points from C inverted on the spherical surface are situated on the convex hull of $\hat{P} \cup C$. Therefore, we need to compute the collection of triangular planes, which make up the convex hull. Then we extract all vertices (magenta points in \hat{P} in Fig. 5(b)) of the convex hull and project them back onto the original point cloud to obtain the visible points P_v (magenta points in P in Fig. 5(b)). The remaining points of the original point cloud are hidden from C , denoted as P_h .

B. Occluded Point Completion

For each input, we employ the FPS and KNN to divide the visible points P_v into patches with v centers. Simultaneously, we extract h central points from the hidden points P_h and retrieve

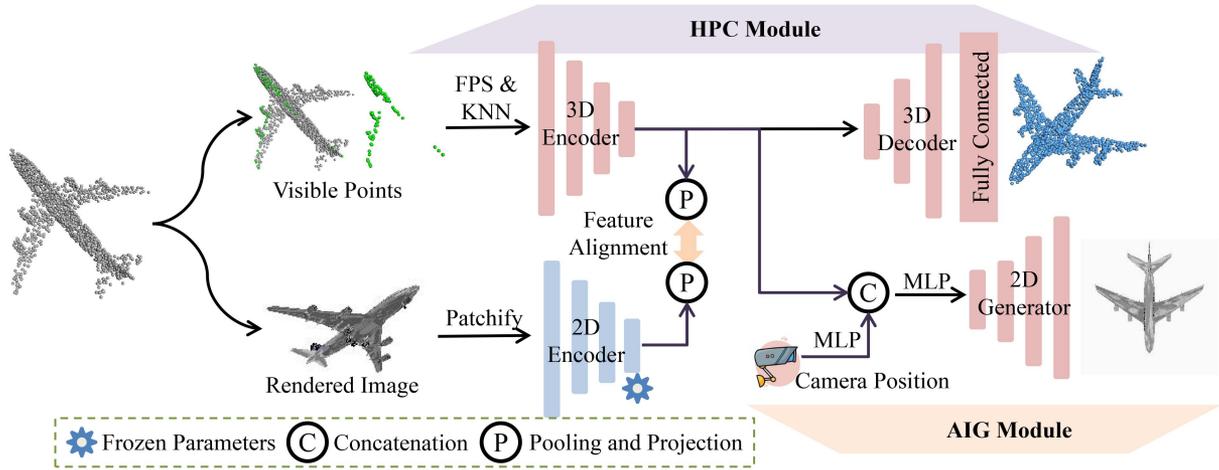


Fig. 4. Overview of PointCG. PointCG effectively integrates two prevalent methods, masked point modeling (MPM) and 3D-to-2D generation, as pretext tasks within a pre-training framework. In detail, we first capture visible points with the HPR [27] operator. Then we utilize an encoder-decoder architecture to extract features from these visible points and complete the original point clouds through the hidden point completion (HPC) module. The arbitrary-view image generation (AIG) module generates images based on the aligned representations of visible points. Note that the input images for feature alignment are randomly selected and do not need to match the target images of image generation.

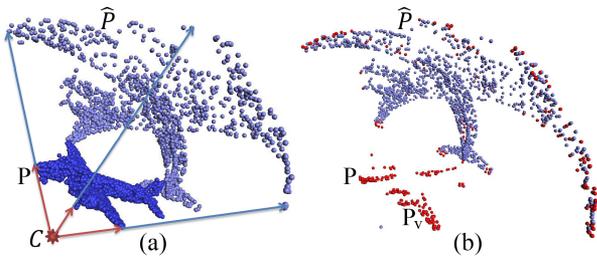


Fig. 5. Visualization of the original points P in blue, points after spherical flipping \hat{P} in light blue and visible points from C in magenta.

k nearest neighbor points from the complete point cloud as the target patches P_{GT} . Then, the visible patches are projected into tokens $T^v \in \mathbb{R}^{v \times d}$ with a lightweight PointNet [31], where d is the dimension of features. Subsequently, a learnable Multi-Layer Perceptron (MLP) is adopted to embed the visible and hidden centers into positional tokens denoted as T_L^v and T_L^h , respectively. Finally, we extract representations T_E by an encoder and capture the tokens T_D with a decoder for completing the original point clouds:

$$T_E = \text{Encoder}(T^v, T_L^v), T_E \in \mathbb{R}^{v \times g \times d} \quad (1)$$

$$T_D = \text{Decoder}(Cat(T_E, T_H), Cat(T_L^v, T_L^h))$$

$$T_H \in \mathbb{R}^{h \times d}, T_D \in \mathbb{R}^{(v+h) \times d} \quad (2)$$

where T_H represents the hidden tokens, which is initialized by duplicating a learnable token of dimension d . We concatenate the visible points' features T_E and T_H , as well as the positional tokens T_L^v and T_L^h as the inputs of the decoder. Based on the outputs T_D of the decoder, we will reconstruct the k nearest neighbors of h center points by a reconstruction head of a fully connected (FC) layer:

$$P_{pre} = \text{Reshape}(\text{FC}(T_D)),$$

$$P_{pre} \in \mathbb{R}^{h \times k \times 3} \quad (3)$$

where P_{pre} denotes as the predicted hidden point patches.

Loss function: The Chamfer distance [32] is employed as the reconstruction loss:

$$\begin{aligned} \mathcal{L}_{CD} = & \frac{1}{|P_{pre}|} \sum_{\hat{x} \in P_{pre}} \min_{x \in P_{GT}} \|\hat{x} - x\|_2^2 \\ & + \frac{1}{|P_{GT}|} \sum_{x \in P_{GT}} \min_{\hat{x} \in P_{pre}} \|\hat{x} - x\|_2^2 \end{aligned} \quad (4)$$

where $P_{GT} \in \mathbb{R}^{h \times k \times 3}$ denotes the reconstruction targets.

C. Arbitrary-View Image Generation

1) *Feature Alignment:* To shift the pre-training focus towards enhancing the encoder for better 3D understanding, we employ the feature alignment module to build correspondence between images and point clouds in the feature space.

During pre-training, a pre-trained CLIP-visual [33] module f is used to extract features from the rendered image I_i . Then, the image features $f(I_i)$ and the 3D features $T_i \in \text{Max}(T_E)$ are projected into the invariant space with functions g and h , respectively, resulting in $\mathcal{H}_i = g(f(I_i))$ and $\mathcal{Z}_i = h(T_i)$.

Loss function: In the invariant space, we aim to maximize the similarity between \mathcal{Z}_i and \mathcal{H}_i when they correspond to the same objects. The cross-modal instance discrimination loss \mathcal{L}_{CM} can be formulated as:

$$\begin{aligned} \mathcal{L}_{CM} = & \frac{1}{2M} \sum_{i=1}^M [l(i, \mathcal{Z}, \mathcal{H}) + l(i, \mathcal{H}, \mathcal{Z})] \\ l(i, \mathcal{Z}, \mathcal{H}) = & -\log \frac{\exp(s(\mathcal{Z}_i, \mathcal{H}_i) / \tau)}{\sum_{k=1, k \neq i}^M \exp(s(\mathcal{Z}_i, \mathcal{Z}_k) / \tau) + \sum_{k=1}^M \exp(s(\mathcal{Z}_i, \mathcal{H}_k) / \tau)} \end{aligned} \quad (5)$$

where M is the mini-batch size. τ is the temperature co-efficient, and $s(\cdot)$ denotes the cosine similarity function.

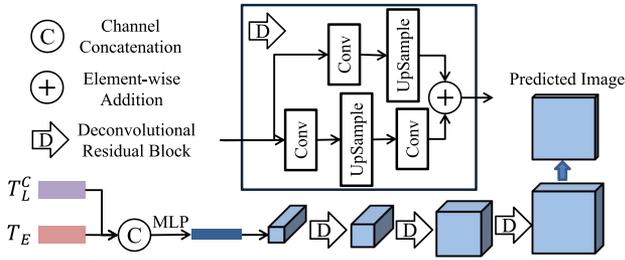


Fig. 6. The image generator consists of several deconvolutional residual blocks, generating the image from the view of L_C .

2) *Image Generation*: AIG generates rendered images from arbitrary views based on the visible points' representations extracted by the encoder.

In pre-training, we randomly select a rendered image as the target and capture the corresponding view parameters L_C as the input. L_C comprises azimuth ϕ , elevation λ , and distance ρ , described as $L_C = \text{Cat}(\phi, \lambda, \rho)$. To enhance flexibility, we apply several learnable transformation layers for positional embedding tokens T_L^C . Then, we concatenate T_E with T_L^C and encode the combined representation using MLP (\mathcal{G}_θ): $T_G^I = \mathcal{G}_\theta \text{Cat}(T_E, T_L^C)$. Finally, we design an image generator (Fig. 6) to generate rendered images based on T_G^I .

Specifically, we start by reshaping T_G^I from its original vectorized representation to a 2D feature map. This feature map is then passed through a series of deconvolutional residual blocks and three parallel convolutional blocks to generate the rendered image G_I^p .

Loss function: We utilize the \mathcal{L}_1 loss as the content loss for image generation:

$$\mathcal{L}_1 = \frac{1}{n} * \sum_{i=1}^n |GT_I - G_I^p| \quad (6)$$

where G_I^p and GT_I represent the predicted and GT images, respectively, while n denotes the number of sample points. Besides, we incorporate multi-scale frequency reconstruction (MSFR) loss [34] as the auxiliary loss alongside the content loss \mathcal{L}_1 to reduce the differences in the frequency space. MSFR loss measures the \mathcal{L}_1 distance between multi-scale GT and predicted images in the frequency domain:

$$\mathcal{L}_{MSFR} = \frac{1}{n} * \sum_{i=1}^n |\mathcal{F}(GT_I) - \mathcal{F}(G_I^p)| \quad (7)$$

where \mathcal{F} denotes the fast Fourier transform (FFT) that transfers the image signal to the frequency domain. MSFR can effectively maintain contrast in high-frequency regions, complementing the ability of \mathcal{L}_1 to preserve colors and luminance [35]. The image generation loss is given by:

$$\mathcal{L}_G = \alpha * \mathcal{L}_1 + \beta * \mathcal{L}_{MSFR} \quad (8)$$

where the contribution of \mathcal{L}_1 and \mathcal{L}_{MSFR} can be adjusted by modifying the values of α and β .

Our loss function during pre-training is formulated as:

$$\mathcal{L} = \omega * \mathcal{L}_{CD} + \phi * \mathcal{L}_{CM} + \psi * \mathcal{L}_G \quad (9)$$

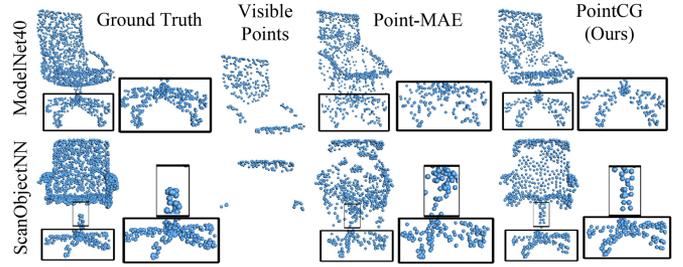


Fig. 7. 3D completion with pre-trained models of Point-MAE [5] and our method, based on the unmasked points from ModelNet40 [39] (top row) and ScanObjectNN [40] (bottom row).

where \mathcal{L}_{CD} enforces 3D completion, \mathcal{L}_{CM} introduces 3D-2D correspondence, and \mathcal{L}_G ensures image generation.

IV. EXPERIMENT

In this section, we present extensive experiments to demonstrate the effectiveness of our method. We begin by introducing the pre-training process on ShapeNet55 [36]. Then, we showcase its performance on 3D completion tasks in Section IV-A. Then, in Sections IV-B, IV-C, and IV-D, we follow the previous works to conduct experiments of object classification, part segmentation, and semantic segmentation. Finally, we validate the effectiveness of our modules through various ablation studies in Section IV-F.

In the following tables, "Pre-T" indicates whether the model is initialized with a pre-trained model, while "Rep." signifies that the result is reproduced with the official code. Please note that we reproduce experiments of Point-MAE [5] and Point-M2AE [10] with their official codes, and all settings are consistent with our experimental configuration.

Pre-training: We pre-train the encoder on ShapeNet55 [36], which contains 52,470 clean 3D models, covering 55 common object categories. The input point number N is 2,048, and the rendered images have a size of $224 \times 224 \times 3$. The encoder and decoder include 12 and 4 standard Transformer blocks, respectively. Each Transformer block has 384 hidden dimensions with 6 heads. We employ the AdamW optimizer [37] and cosine learning rate decay [38]. The initial learning rate is set to 0.001, and the weight decay is 0.05.

A. 3D Completion

Completion based on visible points from random views: To assess the effectiveness of our self-supervised model initialized with pre-trained weights, we randomly select an instance from synthetic dataset ModelNet40 [39] and real-world dataset ScanObjectNN [40] separately and reconstruct the original point clouds. The visualization results of Point-MAE [5] and our model are shown in Fig. 7.

Compared to Point-MAE, our method not only completes the chair's pivot axis and the five-pronged base with greater fidelity but also obtains smoother surface structures. Our method achieves remarkable performance in reconstructing both synthetic and real-world data with visible points from arbitrary views.

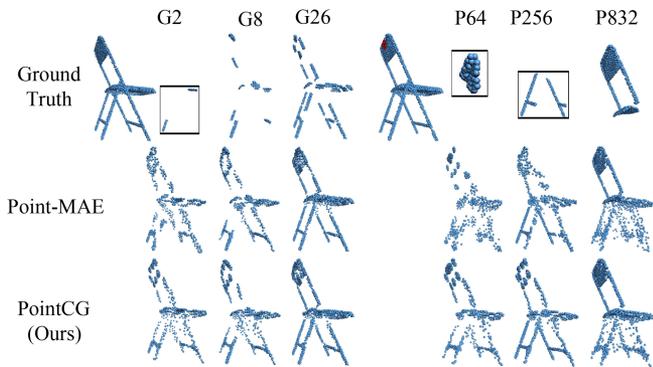


Fig. 8. 3D completion with pre-trained models of Point-MAE [5] and our model based on partial points from ModelNet40 [39].

Completion based on grouped patches and partial points: To demonstrate the robustness and generalizability of our method in the 3D completion task, we devise two methods to obtain grouped patches and partial points as inputs. The reconstruction results are visualized in Fig. 8. In the first method, we obtain 2, 8, and 26 central points via FPS, and then acquire 32 neighboring points with KNN to form groups [5], [10] denoted by ‘G2’, ‘G8’, and ‘G26’. In the second method, we randomly select one group, consisting of 64, 256, and 832 points, denoted as ‘P64’, ‘P256’, and ‘P832’, respectively.

For ‘G2’ and ‘G8’, both Point-MAE and our method can complete the overall structure, but our method excels in recovering finer geometric details and sharper edges. The ‘G26’ column retains sufficient information, and the results are satisfactory for both methods. For ‘P64’ and ‘P256’, our method successfully completes entire structures and captures many local details. In contrast, the results of Point-MAE appear quite blurry. Although the ‘G26’ and ‘P832’ have the same number of input points, the reconstruction results of Point-MAE based on ‘P832’ are significantly lower than ‘G26’.

As shown, the inputs obtained by random masking retain objects’ structural information while exposing the coordinates of target points. These MAE methods employing random masking strategy, such as Point-MAE, exhibit poorer reconstruction performance when the inputs are partial and lack of complete structural integrity. Our method, however, excels in extracting informative representations and demonstrates strong inference capability, leading to superior reconstruction performance, even from highly partial point data.

B. Shape Classification

To assess the discrimination of the representations extracted by the pre-trained encoder, we validate the encoder on the shape classification task using the ModelNet40 [39] and ScanObjectNN [40] datasets.

Shape classification on synthetic data: ModelNet40 [39] contains 12,311 clean 3D CAD models, covering 40 object categories. We fine-tune the pre-trained encoder, and the results are presented in Table I. Our method achieves 94.03% with global fine-tuning, surpassing the reproduced version of

TABLE I
ACCURACY OF SHAPE CLASSIFICATION ON MODELNET40

Methods	Pre-T	booktitle/year	Acc (Vote).
DGCNN [41]	-	ACM/2019	92.9
RSCNN [42]	-	CVPR/2019	93.6
PointTransformer [43]	-	ICCV/2021	93.7
DGCNN+OcCo [12]	Y	ICCV/2021	93.0
DGCNN+STRL [11]	Y	ICCV/2021	93.1
DGCNN+MAE3D [15]	Y	TMM/2023	93.4
Point-BERT [9]	Y	CVPR/2022	93.2
MaskPoint [44]	Y	ECCV/2022	93.8
Point-MAE [5]	Y	ECCV/2022	93.8
Joint-MAE [45]	Y	CoRR/2023	94.0
Point-MAE (Rep.)	Y	ECCV/2022	93.21
PointCG (Point-MAE)	Y	-	94.03
Point-M2AE [10]	Y	NeurIPS/2022	94.0
Point-M2AE (Rep.)	Y	NeurIPS/2022	93.59
PointCG (Point-M2AE)	Y	-	94.11

TABLE II
ACCURACY OF LINEAR-SVM ON MODELNET40

Methods	Pre-T	Acc.
FoldingNet [46]	-	88.4
DGCNN+Jiasaw [47]	Y	90.6
DGCNN+OcCo [12]	Y	89.2
DGCNN+CrossPoint [22]	Y	91.2
FoldingNet+PointMCD [29]	Y	89.8
Point-BERT [9]	Y	87.4
Joint-MAE [45]	Y	92.4
Point-MAE [5]	Y	91.0
PointCG (Point-MAE)	Y	92.26
Point-M2AE [10]	Y	92.9
Point-M2AE (Rep.)	Y	92.63
PointCG (Point-M2AE)	Y	92.92

Point-MAE (Rep.) (93.21%) by 0.82% and the publicly released accuracy by 0.23%. To validate the effectiveness of our architecture, we incorporate Point-M2AE as the backbone and evaluate its classification performance on ModelNet40. The classification results outperform the outcomes of the reproduced Point-M2AE model.

Besides, we also attempt to freeze the parameters of our pre-trained model, and validate it with a Linear-SVM classifier in Table II.

Our method outperforms Point-MAE [5] and Point-M2AE [10] by margins of +1.46% and +0.29%, respectively. The results highlight the superior quality of the 3D representation learned by our method.

Shape classification on the real-world data: Evaluating a pre-trained model’s performance on real-world datasets is crucial, as real-world scenes tend to be more complex than synthetic ones. We follow the common practice to evaluate our model on three variants: ‘OBJ-BG’, ‘OBJ-ONLY’, and ‘PB-T50-RS’ of ScanObjectNN [40].

To further validate the effectiveness of our design, we follow PointVST [13] with DGCNN as the encoder to construct a pre-training network (DGCNN+PointCG) and evaluate it on the ‘PB-T50-RS’ variant. Additionally, we reproduce TAP [16] and PointVST [13] using their official pretrained models.

As presented in Table III, our method outperforms CrossPoint [22], as well as the reproduced TAP [16] and PointVST [13], all using DGCNN [41] as the encoder. When

TABLE III
ACCURACY OF SHAPE CLASSIFICATION ON SCANOBJECTNN

Methods	OBJ-BG	OBJ-ONLY	PB-T50-RS
DGCNN [41]	82.8	86.2	78.1
DGCNN+MAE3D [15]	87.7	88.4	86.2
DGCNN+CrossPoint [22]	-	-	86.2
DGCNN+TAP [16]	-	-	86.6
DGCNN+TAP [16](Rep.)	-	-	86.54
DGCNN+PointVST [13]	-	-	89.3
DGCNN+PointVST [13] (Rep.)	-	-	87.6
DGCNN+PointCG	-	-	87.90
Transformer+OcCo [12]	84.85	85.54	78.79
Transformer+TAP [16]	90.36	89.50	85.67
Point-BERT [9]	87.43	88.12	83.07
Joint-MAE [45]	90.94	88.86	86.07
Point-MAE [5]	90.02	88.29	85.18
PointCG (Point-MAE)	91.16	88.99	86.47
Point-M2AE [10]	91.22	88.81	86.43
Point-M2AE (Rep.)	90.87	88.12	85.39
PointCG (Point-M2AE)	91.19	88.72	86.41

utilizing Point-MAE or Point-M2AE as the backbone, the classification results exhibit a significant improvement over the reproduced Point-MAE and Point-M2AE. These results underscore the discriminative power of the representations extracted by the encoder, even in complex real-world scenes.

Few-shot Learning: Following previous works [5], [9], [12], [47], we conduct few-shot learning experiments using the pre-trained model on ModelNet40 [39]. We adopt n -way, m -shot setting, where n denotes the number of classes randomly selected from the dataset, and m represents the number of objects randomly sampled for each class. This yields $n \times m$ objects for training. For evaluation, we randomly select 20 unseen objects from each of n classes.

The results with settings of $n \in \{5, 10\}$ and $m \in \{10, 20\}$ are presented in Table IV. As shown, our method consistently outperforms the baselines in nearly all few-shot settings, with minimal deviation. This highlights the robustness and generalization of the representations extracted by the PointCG encoder, even in data-limited scenarios.

C. Part Segmentation

The task of part segmentation aims to predict more fine-grained class labels for every instance. We conduct part segmentation on ShapeNetPart [49], which comprises 16,881 samples shared by 16 categories, annotated with 50 parts in total. As illustrated in Table V, our method achieves competitive results and outperforms others in eleven categories.

Visualization of part segmentation: Fine-grained part segmentation holds immense practical value. To highlight the clear advantage of our method in this task, we visualize the results and compare them with Point-MAE [5] in Fig. 9. As depicted in the third line, our method accurately segments the fuselage and tail fin of the airplane, along with the earphone and headband. This reveals the capability of our method to capture discriminative features of points belonging to distinct sections within the same instance.

D. Semantic Segmentation

Large-scale indoor datasets introduce more complexities as they cover larger scenes in real-world environments with noise

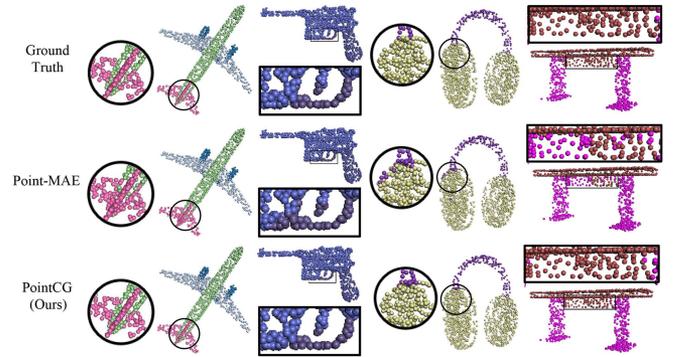


Fig. 9. Visualization of the part segmentation results on the test set of ShapeNetPart.

and outliers. We evaluate the performance of our pre-trained model on the 3D semantic segmentation task using the Stanford large-scale 3D Indoor Spaces (S3DIS) [50] dataset. S3DIS includes data from 6 indoor areas, comprising a total of 272 rooms. We fine-tune the pre-trained model with Area 1-5 and evaluate it with Area 6. The results for each category are shown in Table VI. Our method outperforms Point-MAE [5] across all categories except ‘beam’ and ‘board’. The results underscore our model’s capability to extract contextual and semantic information, which is crucial for producing fine-grained segmentation outcomes.

In reference to the semantic segmentation experiments of STRL [11], we fine-tune our pre-trained model on one area in Area 1-5, followed by evaluation on Area 6. We extend the experiments of Point-MAE [5] based on the pre-trained model released in the official code and present the mean IoU across all class categories $mIoU(\%)$ and the classification accuracy $Acc(\%)$ in Table VII. Our model exhibits a significant improvement in accuracy and $mIoU$ compared to STRL and Point-MAE. These results demonstrate the capability of our model to extract contextual and semantic information, leading to fine-grained segmentation results.

E. Indoor 3D Object Detection

To validate the effectiveness of our method in scene-level prediction tasks, we conduct object detection experiments on ScanNetV2 [51], a 3D indoor scene dataset with rich annotations, including 1,513 scenes across 18 object classes. The dataset includes semantic labels, per-point instances, and both 2D and 3D bounding boxes. Following TAP [16], we adopt 3DETR [52] as a baseline, construct the PointCG network, and pre-train on the object-level dataset ShapeNet55 [36].

As shown in Table VIII, our method demonstrates superior performance compared to both the baseline 3DETR [52] and TAP [16] in both $AP_{0.25}$ and $AP_{0.5}$ metrics. This improvement indicates that the encoder trained by PointCG effectively captures discriminative information and generalizes well to complex scenes even when pre-trained with object-level datasets.

F. Ablation Study

To investigate the architectural designs of our method, we conduct comprehensive ablation studies with Point-MAE as the

TABLE IV
FEW-SHOT OBJECT CLASSIFICATION ON MODELNET40

Methods	5-way, 10-shot	5-way, 20-shot	10-way, 10-shot	10-way, 20-shot
Transformer [49]	87.8 ± 5.2	93.3 ± 4.3	84.6 ± 5.5	89.4 ± 6.3
Transformer+OcCo [12]	94.0 ± 3.6	95.9 ± 2.3	89.4 ± 5.1	92.4 ± 4.6
Point-BERT [9]	94.6 ± 3.1	96.3 ± 2.7	91.0 ± 5.4	92.7 ± 5.1
Point-MAE [5]	96.3 ± 2.5	97.8 ± 1.8	92.6 ± 4.1	95.0 ± 3.0
PointCG (Point-MAE)	96.7 ± 2.1	98.0 ± 1.3	93.1 ± 3.6	95.8 ± 2.6
Point-M2AE [10]	96.8 ± 1.8	98.3 ± 1.4	92.3 ± 4.5	95.0 ± 3.0
PointCG (Point-M2AE)	97.0 ± 1.9	98.4 ± 1.6	92.8 ± 3.8	95.5 ± 2.9

We conduct 10 independent experiments for each setting and report mean accuracy (%) with standard deviation.

TABLE V
PART SEGMENTATION RESULTS ON SHAPENETPART

Methods	$mIoU_C$	$mIoU_I$	aero plane	bag	cap	car	chair	ear phone	guitar	knife	lamp	laptop	motor bike	mug	pistol	rocket	skate board	table
Transformer [9]	83.42	85.1	82.9	85.4	87.7	78.8	90.5	80.8	91.1	87.7	85.3	95.6	73.9	94.9	83.5	61.2	74.9	80.6
Point-BERT [9]	84.11	85.6	84.3	84.8	88.0	79.8	91.0	81.7	91.6	87.9	85.2	95.6	75.6	94.7	84.3	63.4	76.3	81.5
Point-MAE [5]	84.19	86.1	84.3	85.0	88.3	80.5	91.3	78.5	92.1	87.4	86.1	96.1	75.2	94.6	84.7	63.5	77.1	82.4
PointCG (Point-MAE)	84.48	86.2	84.4	86.1	88.4	80.7	91.3	81.2	91.8	88.3	85.9	95.9	75.7	94.9	85.1	63.7	76.5	81.8

We report the mean IoU across all part categories $mIoU_C(\%)$ and the mean IoU across all instances $mIoU_I(\%)$, as well as the IoU(%) for categories.

TABLE VI
SEMANTIC SEGMENTATION RESULTS ON S3DIS

Methods	$oAcc$	$mAcc$	$mIoU_C$	ceiling	floor	wall	beam	column	window	door	table	chair	sofa	bookcase	board	clutter
Point-MAE [5]	90.29	82.18	73.14	94.4	96.8	77.6	82.0	68.5	78.4	79.0	73.3	82.9	45.4	52.6	55.5	64.6
PointCG (Point-MAE)	92.22	85.17	76.58	96.5	98.4	81.6	81.9	82.0	82.9	83.5	75.1	84.7	49.6	57.5	52.6	69.1

The mean class-wise intersection over union ($mIoU_C(\%)$), mean class-wise accuracy ($mAcc(\%)$), overall accuracy ($oAcc(\%)$), and the $IoU(\%)$ of each class are reported.

TABLE VII
SEMANTIC SEGMENTATION ON S3DIS

Fine-tuning Area	Method	Acc.	mIoU
Area 1 (3687 samples)	STRL [11]	85.28	59.15
	Point-MAE [5]	89.03	71.92
	PointCG (Point-MAE)	90.29	74.28
Area 2 (4440 samples)	STRL [11]	72.37	39.21
	Point-MAE [5]	76.72	47.13
	PointCG (Point-MAE)	78.31	48.95
Area 3 (1650 samples)	STRL [11]	79.12	51.88
	Point-MAE [5]	84.09	64.29
	PointCG (Point-MAE)	85.21	65.63
Area 4 (3662 samples)	STRL [11]	73.81	39.28
	Point-MAE [5]	77.34	45.15
	PointCG (Point-MAE)	78.07	47.30
Area 5 (6852 samples)	STRL [11]	77.28	49.53
	Point-MAE [5]	80.56	51.46
	PointCG (Point-MAE)	81.79	54.04

TABLE VIII
SCENE-LEVEL OBJECT DETECTION ON SCANNETV2 [51]

Methods	Pre-T	$AP_{0.25}$	$AP_{0.5}$
VoteNet	-	58.6	33.5
3DETR	-	62.1	37.9
3DETR+TAP	ShapeNet	63.0(+0.9)	41.4(+3.5)
3DETR+PointCG	ShapeNet	63.21(+1.11)	42.17(+4.27)

Average precision at 0.25 IoU thresholds ($AP_{0.25}$) and 0.5 IoU thresholds ($AP_{0.5}$) of detection are reported.

backbone model and elucidate the individual contribution of each module.

Effectiveness of the components: As shown in Table IX, we validate the effectiveness of each module by enhancing and replacing modules on the baseline. We adopt Point-MAE [5] for comparison and utilize hidden point completion (HPC) as

TABLE IX
ABLATION STUDIES ON THE INTRODUCED MODULES

Methods	Linear-SVM.	Acc.	Acc+Vote
Point-MAE (Rep.) [5]	-	92.22	93.21
(a) HPC	91.15	92.76	93.47
(b) + Feature Alignment	91.23	92.99	93.51
(c) + AIG	92.26	93.52	94.03
(d) PointCG (Vit-B/32)	92.21	93.05	93.97
(e) PointCG (ResNet50 [54])	91.01	92.75	93.39
(f) PointCG (Grayscale image)	91.75	93.19	93.52
(g) PointCG (depth map)	91.09	92.78	93.43

the baseline (a). In (b), we add the feature alignment module with pre-trained Vit-B/16 [33]. Based on (b), we incorporate the arbitrary-view image generation (AIG) module, constituting our PointCG, denoted as (c). As shown in Table IX, while the inclusion of the feature alignment module based on HPC does not substantially improve classification accuracy, the exclusion of this module from PointCG yields a diminished Linear-SVM accuracy of 88.41%. We further replace Vit-B/16 with Vit-B/32 (d) and ResNet50 [53] (e). While Vit-B/32 outperforms Vit-B/16 in the image domain, it does not improve the accuracy of classification. The model with ResNet50 [53] exhibits comparatively poorer performance.

To assess the impact of color in rendered images, we extract grayscale images from the rendered ones and pre-train with them in (f). This operation leads to a decrease in shape classification. Grayscale images may potentially lose structural or finer details inherent in the original images. In experiment (g), we extract depth maps from point clouds following PointCLIP [54] and pre-train with them instead of rendered images. The classification result shows poor performance.

TABLE X
ABLATION STUDIES ON THE INTRODUCED MODULES

Methods	Epoch time (s)	Inference time (s)
Point-MAE (Rep.) [5]	54.73-57.59	36.85
HPC	395.18-399.26	36.65
PointCG w/o AIG	413.96-419.47	37.65
PointCG w/o Feature Alignment	405.52-411.22	35.67
PointCG w/o 3D Completion	520.97-524.73	36.79
PointCG (Point-MAE)	534.56-538.72	37.37

Training and inference times of each model.

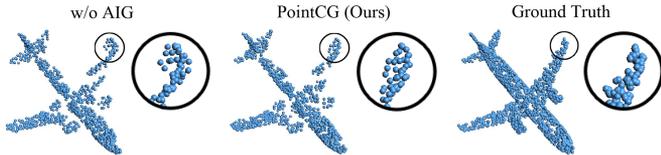


Fig. 10. Visualization of the 3D completion results based on the pre-trained models of PointCG without AIG (w/o AIG) and PointCG on ModelNet40 [39].

Training and inference time: Table X presents pre-training and inference times for the classification task of each module. The results demonstrate that the pre-training time is notably longer than that of the baseline, whereas the inference time for classification remains comparable.

The pre-training of HPC requires approximately 397s per epoch. Compared to the baseline (Point-MAE), the primary time consumption arises from the data organization module, which is responsible for obtaining visible points from arbitrary views as inputs. Among the components, the 3D completion module has the shortest pre-training duration, while the feature alignment and AIG modules demand considerably more time.

Visualization of 3D Completion: To validate the positive impact of AIG on the 3D completion task, we exclude AIG from PointCG (w/o AIG) and pre-train the network. The 3D completion results of this variant and PointCG are shown in Fig. 10. As depicted, the edges of the aircraft wings are sharpened with our method. Without AIG, the completion edges exhibit point groups, attributed to solely relying on 3D completion, while the completion targets are point clusters.

Visualization of the generated images: To validate the provision of geometric structural information by 3D completion, we showcase the results of image generation after excluding the completion module and compare them with PointCG’s results in Fig. 11. As depicted in line (c), outcomes exhibit significant artifacts and lack local structural information. Specifically, in the case of the aircraft tail wing, only the wing’s shape is generated, omitting volumetric structural details. However, our method in line (d) successfully predicts the accurate shapes of the objects.

Camera perspectives: To identify more suitable inputs and predicted images, we design experiments with different camera positions as inputs, and the results are reported in Table XI. (a) takes the left-view image as input and the front-view as target, and (b) takes the left-view image as input and an arbitrary-view image as target. In (c), we use images from three different views (front, left, and top views) as inputs and an arbitrary-view image

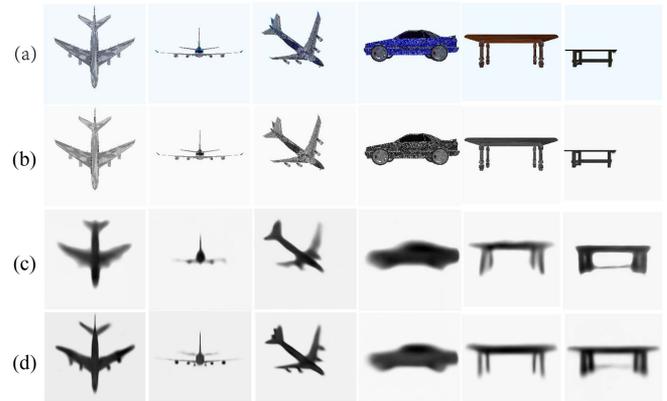


Fig. 11. Visualization of the generated images based on pre-trained models. We present the original rendered images in row (a) and corresponding grayscale one in row (b). The row (c) shows the results of PointCG without 3D completion. The results based on PointCG are shown in row (d).

TABLE XI
ABLATION STUDIES ON CAMERA VIEWS

Input/Prediction	Linear-SVM	Acc.	Acc+Vote
(a) Left view / front view	91.13	92.66	93.31
(b) Left view / arbitrary view	91.65	92.83	93.68
(c) Front, left, and top views / arbitrary view	92.22	93.03	93.40
(d) Arbitrary view / arbitrary view	92.26	93.52	94.03

Pre-training and shape classification with images from different views.

TABLE XII
ABLATION STUDY ON THE COMPLETENESS OF INPUTS

Input	Linear-SVM	Acc	Acc+Vote
(a) One arbitrary view	92.26	93.52	94.03
(b) One arbitrary view + 2 patches	91.73	93.14	93.55
(c) One arbitrary view + 8 patches	91.41	92.85	93.25
(d) Two arbitrary views	91.02	92.75	93.27

We pre-train the model based on inputs with varying levels of completeness.

as target. Both the input and target of (d) are from arbitrary views, yielding the optimal results. Therefore, we utilize two arbitrary-view images as the input and target, respectively.

Pre-training with more complete inputs: We posit that if the inputs of 3D completion contain more structural information and more overlap areas with the targets, the completion task will be accomplished more easily. This leads to a reduced training intensity for the backbone, thereby diminishing the backbone’s perception of 3D objects. We design this study based on different inputs. The inputs in (a) are derived from a single arbitrary view. The inputs for (b) and (c) involve the addition of two and eight extra patches, respectively, to the input of (a). The points from two arbitrary views serve as the inputs for (d). Results are reported in Table XII.

In cases (b) and (c), the accuracy of the Linear-SVM during pre-training decreases as more patches are included in the inputs. The classification results via fine-tuning also show a decline. In case (d), inputs from two views offer more structural information about the input objects, which leads to a significant decrease in shape classification. This experiment reveals that as additional structural information is progressively included in the input, shape classification accuracy steadily decreases. This indicates

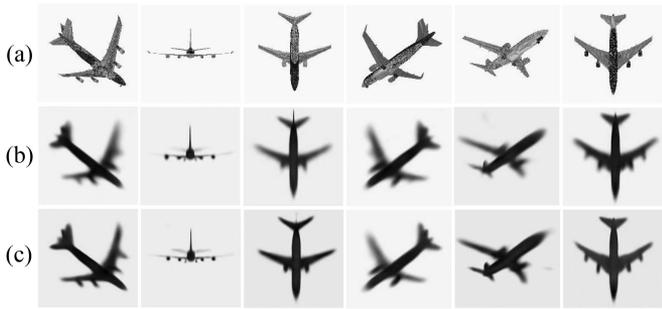


Fig. 14. Visualization of the generated images. We present the grayscale versions of the original rendered images in row (a), the generated images with multi-scale representations in row (b), and the generated images by PointCG in row (c).

V. CONCLUSION

In this paper, we propose PointCG, a unified framework with hidden points completion and arbitrary-view image generation for self-supervised point cloud learning. Completion and generation based on partial points prompt the encoder to extract high-quality representations with 3D structural intricacies and alleviate ambiguous supervision. Thereby, our method achieves notable enhancements over baseline methods and outperforms similar methods in classification and reconstruction tasks on real datasets. We expect our pre-trained models will benefit a wide range of 3D tasks, including 3D object detection, semantic segmentation, and visual grounding.

While performing well across multiple tasks, there is much room for improving PointCG. This includes expanding the modality of images to incorporate other formats such as language or audio. Furthermore, while we focus on instance-level tasks, scene-level understanding is crucial in real-world applications. Therefore, our future studies may involve delving into applications in scene understanding and investigating interactions among multiple modalities for 3D reasoning.

REFERENCES

- [1] H. Chen, M. Wei, Y. Sun, X. Xie, and J. Wang, "Multi-patch collaborative point cloud denoising via low-rank recovery with graph constraint," *IEEE Trans. Vis. Comput. Graph.*, vol. 26, no. 11, pp. 3255–3270, Nov. 2020.
- [2] H. Chen, Z. Wei, X. Li, Y. Xu, M. Wei, and J. Wang, "RePCD-Net: Feature-aware recurrent point cloud denoising network," *Int. J. Comput. Vis.*, vol. 130, no. 3, pp. 615–629, 2022.
- [3] Y. Liu, X. Yan, Z. Li, Z. Chen, Z. Wei, and M. Wei, "PointGame: Geometrically and adaptively masked autoencoder on point clouds," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5705312.
- [4] L. Gu et al., "PointSee: Image enhances point cloud," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 9, pp. 6291–6308, Sep. 2024.
- [5] Y. Pang, W. Wang, F. E. H. Tay, W. Liu, Y. Tian, and L. Yuan, "Masked autoencoders for point cloud self-supervised learning," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 604–621.
- [6] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. B. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 15979–15988.
- [7] K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 9726–9735.
- [8] X. Chen and K. He, "Exploring simple siamese representation learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 15750–15758.
- [9] X. Yu, L. Tang, Y. Rao, T. Huang, J. Zhou, and J. Lu, "Point-BERT: Pre-training 3D point cloud transformers with masked point modeling," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 19291–19300.
- [10] R. Zhang et al., "Point-M2AE: Multi-scale masked autoencoders for hierarchical point cloud pre-training," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 27061–27074.
- [11] S. Huang, Y. Xie, S. Zhu, and Y. Zhu, "Spatio-temporal self-supervised representation learning for 3D point clouds," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 6515–6525.
- [12] H. Wang, Q. Liu, X. Yue, J. Lasenby, and M. J. Kusner, "Unsupervised point cloud pre-training via occlusion completion," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9762–9772.
- [13] Q. Zhang and J. Hou, "PointVST: Self-supervised pre-training for 3D point clouds via view-specific point-to-image translation," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 10, pp. 6900–6912, Oct. 2024.
- [14] Z. Qi et al., "Contrast with reconstruct: Contrastive 3D representation learning guided by generative pretraining," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 28223–28243.
- [15] J. Jiang, X. Lu, L. Zhao, R. Dazeley, and M. Wang, "Masked autoencoders in 3D point cloud representation learning," 2022, *arXiv:2207.01545*.
- [16] Z. Wang, X. Yu, Y. Rao, J. Zhou, and J. Lu, "Take-A-Photo: 3D-to-2D generative pre-training of point cloud models," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 5617–5627.
- [17] W. Feng, J. Zhang, H. Cai, H. Xu, J. Hou, and H. Bao, "Recurrent multi-view alignment network for unsupervised surface registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 10297–10307.
- [18] J. Sauder and B. Sievers, "Self-supervised deep learning on point clouds by reconstructing space," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 12942–12952.
- [19] Y. Zeng, Y. Qian, Z. Zhu, J. Hou, H. Yuan, and Y. He, "CorrNet3D: Unsupervised end-to-end learning of dense correspondence for 3D point clouds," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 6052–6061.
- [20] J. Grill et al., "Bootstrap your own latent - A new approach to self-supervised learning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 21271–21284.
- [21] M. Caron et al., "Emerging properties in self-supervised vision transformers," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9630–9640.
- [22] M. Afham, I. Dissanayake, D. Dissanayake, A. Dharmasiri, K. Thilakarathna, and R. Rodrigo, "CrossPoint: Self-supervised cross-modal contrastive learning for 3D point cloud understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 9892–9902.
- [23] A. Radford et al., "Language models are unsupervised multitask learners," *OpenAI Blog*, vol. 1, no. 8, 2019, Art. no. 9.
- [24] S. Yan et al., "3D feature prediction for masked-autoencoder-based point cloud pretraining," in *Proc. Int. Conf. Learn. Representations*, 2024, pp. 1–12.
- [25] S. Yan et al., "Implicit autoencoder for point-cloud self-supervised representation learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 14484–14496.
- [26] W. StraBer, "Schnelle Kurven-und Flachendarstellung auf graphischen Sichtgeraten," PhD dissertation, Fachbereich Technische Univ. Berlin, Berlin, Germany, 1974.
- [27] S. Katz, A. Tal, and R. Basri, "Direct visibility of point sets," *ACM Trans. Graph.*, vol. 26, no. 3, 2007, Art. no. 24.
- [28] H. Chen, Z. Wei, Y. Xu, M. Wei, and J. Wang, "ImLoveNet: Misaligned image-supported registration network for low-overlap point cloud pairs," in *Proc. ACM SIGGRAPH Conf.*, 2022, pp. 29:1–29:9.
- [29] Q. Zhang, J. Hou, and Y. Qian, "PointMCD: Boosting deep point cloud encoders via multi-view cross-modal distillation for 3D shape recognition," *IEEE Trans. Multimedia*, to be published, doi: [10.1109/TMM.2023.3286981](https://doi.org/10.1109/TMM.2023.3286981).
- [30] S. Katz, G. Leifman, and A. Tal, "Mesh segmentation using feature point and core extraction," *Vis. Comput.*, vol. 21, no. 8–10, pp. 649–658, 2005.
- [31] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 77–85.
- [32] H. Fan, H. Su, and L. J. Guibas, "A point set generation network for 3D object reconstruction from a single image," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2463–2471.
- [33] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.

- [34] S. Cho, S. Ji, J. Hong, S. Jung, and S. Ko, "Rethinking coarse-to-fine approach in single image deblurring," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 4621–4630.
- [35] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for image restoration with neural networks," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 47–57, Mar. 2017.
- [36] A. X. Chang et al., "ShapeNet: An information-rich 3D model repository," 2015, *arXiv:1512.03012*.
- [37] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. Int. Conf. Learn. Representations*, 2019, pp. 1–8.
- [38] I. Loshchilov and F. Hutter, "SGDR: Stochastic gradient descent with warm restarts," in *Proc. Int. Conf. Learn. Representations*, 2017, pp. 1–16.
- [39] Z. Wu et al., "3D ShapeNets: A deep representation for volumetric shapes," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1912–1920.
- [40] M. A. Uy, Q. Pham, B. Hua, D. T. Nguyen, and S. Yeung, "Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 1588–1597.
- [41] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph CNN for learning on point clouds," *ACM Trans. Graph.*, vol. 38, no. 5, pp. 146:1–146:12, 2019.
- [42] Y. Liu, B. Fan, S. Xiang, and C. Pan, "Relation-shape convolutional neural network for point cloud analysis," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 8895–8904.
- [43] H. Zhao, L. Jiang, J. Jia, P. H. S. Torr, and V. Koltun, "Point transformer," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 16239–16248.
- [44] H. Liu, M. Cai, and Y. J. Lee, "Masked discrimination for self-supervised learning on point clouds," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 657–675.
- [45] Z. Guo, R. Zhang, L. Qiu, X. Li, and P. Heng, "Joint-MAE: 2D-3D joint masked autoencoders for 3D point cloud pre-training," in *Proc. Int. Joint Conf. Artif. Intell.*, 2023, pp. 791–799.
- [46] Y. Yang, C. Feng, Y. Shen, and D. Tian, "FoldingNet: Point cloud auto-encoder via deep grid deformation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 206–215.
- [47] C. Sharma and M. Kaul, "Self-supervised few-shot learning on point clouds," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 7212–7221.
- [48] A. Vaswani et al., "Attention is all you need," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.
- [49] L. Yi et al., "A scalable active framework for region annotation in 3D shape collections," *ACM Trans. Graph.*, vol. 35, no. 6, pp. 210:1–210:12, 2016.
- [50] I. Armeni et al., "3D semantic parsing of large-scale indoor spaces," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 1534–1543.
- [51] A. Dai, A. X. Chang, M. Savva, M. Halber, T. A. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3D reconstructions of indoor scenes," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2432–2443.
- [52] I. Misra, R. Girdhar, and A. Joulin, "An end-to-end transformer model for 3D object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 2886–2897.
- [53] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [54] R. Zhang et al., "PointCLIP: Point cloud understanding by CLIP," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 8542–8552.



Peng Li is currently working toward the PhD degree with the School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, China. He has published several papers on the *IEEE Transactions on Image Processing*, *AAAI*, *PG*, etc. His research interests include deep learning, image processing, and computer vision.



Xuefeng Yan received the PhD degree from the Beijing Institute of Technology, in 2005. He is a professor with the School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics (NUAA), China. He was the visiting scholar with Georgia State University in 2008 and 2012. His research interests include intelligent computing, MBSE/complex system modeling, simulation, and evaluation.



Liangliang Nan received the PhD degree in mechatronics engineering from the Graduate University of the Chinese Academy of Sciences, China, in 2009. He is currently an associate professor with the Delft University of Technology (TU Delft), The Netherlands, where he is leading the AI Laboratory on 3D Urban Understanding (3DUU). His research interests include computer graphics, computer vision, 3D geoinformation, and machine learning.



Bing Wang received the DPhil/PhD degree from the Department of Computer Science, University of Oxford, in 2022. He is an assistant professor with the Department of Aeronautical and Aviation Engineering, Hong Kong Polytechnic University. His research interests broadly lie in the design of intelligent perception solutions for autonomous systems, and the development of reliable 3D scene understanding algorithms on mobile robotics operating in the real world.



Yun Liu received the BSc degree from Huaiyin Normal University, in 2013, and the MSc degree from Jiangsu University, in 2016. She is currently working toward the PhD degree with the Nanjing University of Aeronautics and Astronautics. Her research interests include 3D computer vision and self-supervised learning.



Honghua Chen received the PhD degree from the Nanjing University of Aeronautics and Astronautics (NUAA), Nanjing, China, in 2022. His research interests include 3D measurement/vision and point cloud processing.



ing, software analysis, software testing, and mining software repositories.

Lina Gong received the BE degree in software engineering from the China University of Petroleum, China, and the PhD degree in computer software and theory from the China University of Mining and Technology, China. She is currently an associate professor with the College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, China. She also studied as a visitor one year with the Software Analysis and Intelligence Lab (SAIL), School of Computing, Queen's University, Canada. Her research interests include machine learning,



on Multimedia Computing, Communications, and Applications, The Visual Computer Journal, Journal of Electronic Imaging, and a guest editor of the IEEE Transactions on Multimedia. His research interests focus on 3D vision, computer graphics, and deep learning.

Mingqiang Wei (Senior Member, IEEE) received the PhD degree in computer science and engineering from the Chinese University of Hong Kong (CUHK), in 2014. He is a professor with the School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics (NUAA). Before joining NUAA, he served as an assistant professor with the Hefei University of Technology, and a post-doctoral fellow with CUHK. He was a recipient of the CUHK Young Scholar Thesis Awards in 2014. He is now an associate editor of the *ACM Transactions*



Wei Zhao received the PhD degree in biological science & medical engineering from the Southeast University, Nanjing, China, in 2024. He is a post-doctoral fellow with the School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics. His research interests focus on machine learning, genetic algorithm, and 3D structure optimization.