

# Take-A-Photo: 3D-to-2D Generative Pre-training of Point Cloud Models

Ziyi Wang<sup>1,2\*</sup> Xumin Yu<sup>1,2\*</sup> Yongming Rao<sup>1,2</sup> Jie Zhou<sup>1,2</sup> Jiwen Lu<sup>1,2†</sup>

<sup>1</sup>Department of Automation, Tsinghua University <sup>2</sup>BNRist

{wziyi22, yuxm20}@mails.tsinghua.edu.cn;

raoyongming95@gmail.com; {jzhou, lujiwen}@tsinghua.edu.cn

## Abstract

With the overwhelming trend of mask image modeling led by MAE, generative pre-training has shown a remarkable potential to boost the performance of fundamental models in 2D vision. However, in 3D vision, the over-reliance on Transformer-based backbones and the unordered nature of point clouds have restricted the further development of generative pre-training. In this paper, we propose a novel 3D-to-2D generative pre-training method that is adaptable to any point cloud model. We propose to generate view images from different instructed poses via the cross-attention mechanism as the pre-training scheme. Generating view images has more precise supervision than its point cloud counterpart, thus assisting 3D backbones to have a finer comprehension of the geometrical structure and stereoscopic relations of the point cloud. Experimental results have proved the superiority of our proposed 3D-to-2D generative pre-training over previous pre-training methods. Our method is also effective in boosting the performance of architecture-oriented approaches, achieving state-of-the-art performance when fine-tuning on ScanObjectNN classification and ShapeNet-Part segmentation tasks. Code is available at <https://github.com/wangzy22/TakeAPhoto>.

## 1. Introduction

Nowadays, pre-training fundamental models with self-supervised mechanisms has witnessed a thriving development in the computer vision community, given its low requirement in data annotation and its high transferability to downstream applications. Self-supervised pre-training aims to fully exploit the statistical and structural knowledge from abundant annotation-free data and empowers the fundamental models with robust representation ability. In 2D vision, self-supervised pre-training has shown strong potential and achieved remarkable progress on diverse backbones in various downstream tasks. Successful pre-training strategies




Figure 1: **Principle illustration of 3D-to-2D generative pre-training.** The photograph module explicitly encodes pose condition into 3D features from the backbone, and the 2D generator decodes pose-conditioned features into different view images.

in 2D domain such as contrastive learning [19, 7] and mask modeling [18, 3] have also been adopted to 3D point cloud analysis [58, 30, 23] in recent research.

However, the pre-training paradigm hasn't become the prevailing trend in 3D learning and architectural design is still the mainstream to reach a new state-of-the-art performance, which is considerably different from the dominant status of pre-training in 2D domain. In object-level analysis, generative pre-training inspired by MAE [18] has been studied thoroughly, but their performances still lag behind architecture-based methods like PointNeXt [35]. Two factors mainly contribute to the inferior status of generative pre-training in 3D learning. Since point clouds are unordered sets of point coordinates, there is no direct and precise supervision like one-to-one MSE loss between generated point clouds and their corresponding ground truth. Chamfer Distance supervision for point clouds only calculates a rough set-to-set matching and its imprecision has been widely discussed in [24, 50, 20]. Additionally, existing advanced generative pre-training methods in object analysis are limited to the Transformer-based backbone, and fail to be extended to other sophisticated point cloud models.

To alleviate the aforementioned problems, we propose a 3D-to-2D generative pre-training method for point cloud analysis that has higher preciseness in supervision and broader adaptation to different backbones. Instead of re-

\*Equal contribution. †Corresponding author.

constructing point clouds as previous literature [30, 60], we propose to generate view images of the input point cloud given the instructed camera poses. This is similar to taking photos of a realistic object from different perspectives to fully present its structure or internal relations. Therefore, we name our model *Take-A-Photo*, in short *TAP*. More specifically, we propose a pose-dependent photograph module that utilizes the cross-attention mechanism to explicitly encode pose conditions with 3D features from the backbone. Then a 2D generator decodes pose-conditioned features into view images that are supervised by rendered ground truth images. The principle illustration of TAP is shown in Figure 1. In the pose-dependent photograph module, we do not provide detailed projection relations from 3D points to 2D pixels, thus the cross-attention layers are encouraged to learn by themselves what those view images look like conditioned on given poses. Since the projection layout, part occlusion relation, faces colors that represent light reflections are largely distinct among view images, the proposed 3D-to-2D generative pre-training is a challenging pretext task that obliges 3D backbone to gain higher representation ability of geometrical knowledge and stereoscopic relations.

We conduct extensive experiments on various backbones and downstream tasks to verify the effectiveness and superiority of our proposed 3D-to-2D generative pre-training method. When pre-trained on synthetic ShapeNet [6] and transferred to real-world ScanObjectNN [45] classification, TAP brings consistent improvement to different backbone models and successfully outperforms previous point cloud generative pre-training methods based on Transformers backbone. With PointMLP [27] as the backbone, TAP achieves state-of-the-art performance on ScanObjectNN classification and ShapeNetPart [56] part segmentation among methods that do not include any pre-trained image or text model. We also conduct thorough ablation studies to discuss the architectural design of the TAP model and verify the individual contribution of each component.

In conclusion, the contributions of our paper can be summarized as follows: (1) We propose TAP, the first 3D-to-2D generative pre-training method that is adaptable to any point cloud model. TAP pre-training helps to exploit the potential of point cloud models on geometric structure comprehension and stereoscopic relation understanding. (2) We propose a Photograph Module where we derive mathematical formulations to encode pose conditions as query tokens in cross-attention layers. (3) TAP surpasses previous generative pre-training methods on the Transformers backbone and achieves state-of-the-art performance on ScanObjectNN classification and ShapeNetPart segmentation.

## 2. Related Work

**Point Cloud Analysis.** Point cloud analysis is a fundamental and important task in the realm of 3D vision. Current

literature has developed two principal methodologies to extract structural representations from 3D point clouds, namely point-based and voxel-based methods. Point-based methods [32, 33, 48, 44, 35, 27] process unordered points directly, introducing various techniques for local information aggregation. Existing point-based methods can be categorized into three types: SetAbstraction-based [32, 33, 35], DynamicGraph-based [48], and Attention-based [57, 58, 30, 23, 60], all of which focus on modeling the relationships between points. Owing to their exceptional ability to effectively preserve fine-grained geometric information, point-based methods are frequently employed for object-level tasks. On the other hand, voxel-based methods [28, 21, 39] partition the 3D space into ordered voxels and employ 3D convolutions for feature extraction. Voxel-based methods are primarily based on SparseConvolution [10, 16], which enables efficient convolution operations in 3D space through sparse convolutions. In exchange for faster processing speeds, voxel-based methods sacrifice their capacity to capture detailed local structures, making them more suitable for large-scale scene-level tasks rather than object-level tasks.

**Point Cloud Pre-training.** Pre-training has always been a focal point of research in the field of deep learning. Generally speaking, we usually distinguish pre-training methods based on the amount of annotation required, namely full-supervised pre-training [14, 59, 5], weakly-supervised pre-training [43, 4, 53], and unsupervised pre-training [19, 7, 8]. Among these methods, unsupervised pre-training has become the most popular approach, mainly due to its excellent transferability across tasks and its notable advantage of not relying on labeled data. Numerous researchers have proposed a variety of pretext tasks for unsupervised pre-training. Based on the type of pretext task employed, there are two prevailing pretext tasks for unsupervised pre-training. The first is contrastive learning, as exemplified by MoCo [19] and SimCLR [7]. The other method entails utilizing generative tasks to restore the data from partially or disrupted inputs, such as MAE [18] and BEiT [3].

Inspired by pre-training strategies in the image domain, there are more and more unsupervised pre-training methods being proposed for point cloud pre-training. PointContrast [54] embraces the principle of contrastive learning, whereas PointBERT [58] and PointMAE [30] integrate reconstruction pretext tasks. However, existing generative-based pre-training methods for point clouds only consider a single modality. In this paper, we propose a cross-modal generative-based pre-training strategy to achieve more effective pre-training.

**Cross-Modal Learning.** Recently, cross-modal learning has been a popular research topic, aiming at learning from multiple modalities such as images, audio and point clouds. It has the potential to enhance the performance of various tasks,

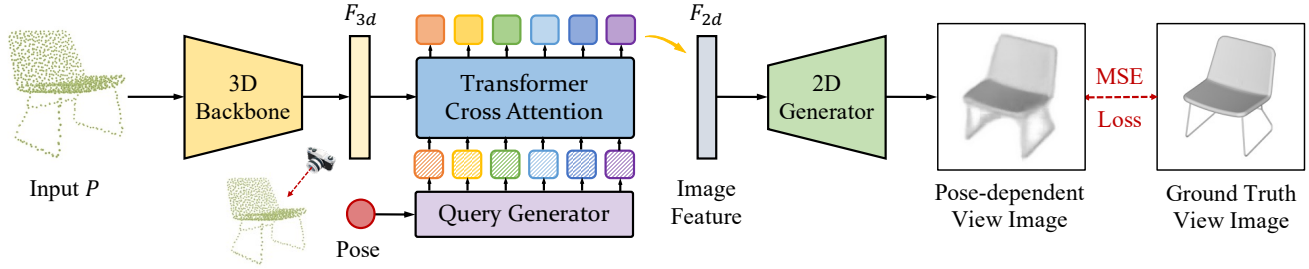


Figure 2: **The pipeline of TAP pre-training method.** We design a query generator to encode pose conditions and implement cross-attention layers to transform 3D point cloud features  $F_{3d}$  to 2D view image features  $F_{2d}$  according to pose instruction. The predicted pose-dependent view image from 2D generator is supervised by ground truth view image via MSE loss.

including visual recognition, speech recognition, and point cloud analysis. A variety of methods have been proposed for cross-modal learning, including multi-task learning [11], conditional generation [40], pre-training [34, 13] and tuning [49, 55].

In the realm of point cloud analysis, much previous work has explored this learning paradigm. Some literature leverages 2D data for 3D point cloud analysis, such as MVTN [17], MVCNN [42] and CrossPoint [2], proving that the multi-view images and the correspondence between images and points can be helpful for the 3D object understanding. Another line of the research, such as Image2Point [55] and P2P [49] take effort to adapt the models from 2D vision into 3D point cloud analysis, fully exploiting the relationship of 2D and 3D understanding. In this paper, we continue this learning paradigm in 3D vision domain, and for the first time propose the cross-modal generative pre-training scheme for point cloud models.

### 3. 3D-to-2D Generative Pre-training

#### 3.1. Preliminary: Generative Pre-training

Generative pre-training is a fundamental branch of pre-training methods that aims at reconstructing integral and complete data given partial or disrupted input. Mathematically, suppose  $x$  is a sample from raw data with no annotation. The pre-processing step  $T(\cdot)$  either erases part of  $x$  randomly or splits  $x$  into pieces and intermingles them to get  $\hat{x} = T(x)$ . The generative pre-training model  $M$  is designed to restore from those broken input  $\hat{x} = M(T(x))$  and the training loss function is designed to measure the reconstruction distance  $\mathcal{L} = D(\hat{x}, x)$ . In point cloud object analysis, earlier generative pre-training methods propose various pretext tasks as  $T$ , including deformation [1], jigsaw puzzles [41] and depth projection [47] to produce disarrayed or partial point clouds. Recently, inspired by MAE [18] in the image domain, generative pre-training in 3D domain mainly focuses on implementing random masking as  $T$  and utilizing

Transformers model as  $M$  for reconstruction [58, 30, 23, 60]. The reconstruction distance  $D$  is usually measured by the classical  $l_2$  Chamfer Distance:

$$D(\hat{x}, x) = \frac{1}{|\hat{x}|} \sum_{a \in \hat{x}} \min_{b \in x} \|a - b\|_2^2 + \frac{1}{|x|} \sum_{b \in x} \min_{a \in \hat{x}} \|a - b\|_2^2 \quad (1)$$

Besides Chamfer Distance between point clouds, some methods also exploit feature distance between latents [58] or occupancy value distance [23] as the loss function.

The exact reason why generative pre-training would help enhance the representation ability of backbone models still remains an open question. However, abundant experimental results have conveyed that predicting missing parts according to known parts demands high reasoning ability and global comprehension capacity of the model. What’s more, generative pre-training is more efficient and suitable for point cloud object analysis than contrastive pre-training, given that contrastive pre-training typically requires a large amount of training data to avoid trivial overfitting solutions but there has always been a data-starvation problem in point cloud object research field.

#### 3.2. Overall Pipeline

Different from the aforementioned generative pre-training methods that focus on uni-modal point cloud reconstruction, we propose a novel cross-modal pre-training approach of generating view images from instructed camera poses.

The overall architecture of our proposed TAP pre-training model is depicted in Figure 2. Our model takes as an input point cloud  $P \in \mathbb{R}^{N \times 3}$ , where  $N$  is the number of points in the input point cloud. The basic building block of TAP mainly consists of: 1) a *3D Backbone* that extracts 3D geometric features  $F_{3d} \in \mathbb{R}^{n \times C_{3d}}$ , where  $n$  is the number of downsampled center points and  $C_{3d}$  is the geometric feature dimension; 2) a *pose-dependent Photograph Module* that takes as inputs  $F_{3d}$  and pose matrix  $R \in \mathbb{R}^{3 \times 3}$ , and predicts view image features  $F_{2d}^R \in \mathbb{R}^{h \times w \times C_{2d}}$  conditioned on  $R$ , where  $h, w$  are height and width of predicted view image

feature map; 3) an *2D Generator* that decodes  $F_{2d}^R$  into an RGB image  $I_{\text{gen}}^R \in \mathbb{R}^{H \times W \times 3}$ , where  $H, W$  are height and width of the output view image.

As we place no restriction on  $F_{3d}$ , the *3D Backbone* can be arbitrarily chosen and adopted. Therefore, our TAP is more flexible and compatible than existing generative pre-training methods that are limited to Transformer-based architecture. Experimental results in Section 4 will later verify that TAP brings consistent improvement to all kinds of point cloud models. The technical designs of the *pose-dependent Photograph Module* will be thoroughly discussed in Section 3.3. The *2D Generator* consists of four Transpose Convolution layers to progressively upsample image resolution and decode RGB colors of each pixel.

### 3.3. Photograph Module

**Architectural Design.** As illustrated in Figure 2, we leverage cross-attention mechanism from Transformers [46] to build our *pose-dependent Photograph Module*.

$$\text{Attention}(Q, K, V) = \text{softmax} \left( \frac{QK^T}{\sqrt{d_k}} \right) V \quad (2)$$

where  $d_k$  is the scaling factor, and  $Q, K, V$  are queries, keys and values matrix. More specifically, we design a Query Generator  $\Phi$  to encode camera pose conditions into query tokens:  $Q = \Phi(R) \in \mathbb{R}^{hw \times C_{2d}}$ . We also design a Memory Builder  $\Theta$  to construct  $K$  and  $V$  from 3D geometric features:  $K = V = \Theta(F_{3d}) \in \mathbb{R}^{m \times C_{2d}}$ , where  $m$  is the number of memory tokens. The output sequence of the cross attention layers will be rearranged from  $hw \times C_{2d}$  to  $h \times w \times C_{2d}$ , forming the predicted view image features  $F_{2d}^R$ .

During the cross-attention calculation process, we do not explicitly provide any projection clues of which 3D points would project to which 2D pixel. Instead, the Photograph Module learns by itself how to arrange unordered 3D feature points to ordered 2D pixel grids, purely based on semantic similarities between 3D geometric features and our delicately-designed queries that reveal pose information. Since one sample will only have one set of memory tokens in 3D space but its view images from different poses are quite distinct from each other, learning to predict precise view images from instructed poses in a data-driven manner is not a trivial task. Therefore, during the end-to-end optimization process, the 3D backbone is trained to have a stronger perception of the object’s overall geometric structure and gain a higher representative ability of the stereoscopic relations. In this way, our proposed 3D-to-2D generative pre-training would help exploit the potential and enhance the strength of 3D backbone models.

**Query Generator.** The query generator  $\Phi$  is designed to encode pose condition  $R$  into 2D grid of shape  $h \times w$ . In object analysis, common practice is leveraging parallel light

shading to project 3D objects onto 2D grids, and pose matrix  $R$  here is used to rotate objects into desired angles before projection. Therefore, each 2D grid actually represents an optical line that starts from infinity, passes through 3D objects and ends at the 2D plane. As a consequence, we choose the direction and the origin points that the optical line goes through as the delegate of the query grid.

Before deriving formulations of optical lines for each grid, let us first revisit the parallel light shading process for better comprehension. Given 3D coordinates  $\mathbf{x} = (x, y, z)$  of a point cloud  $P$  and pose matrix  $R$ , rotation is first performed to align the object to the ideal pose position:

$$\mathbf{x}' = (x', y', z') = R\mathbf{x} \quad (3)$$

Then we just omit the final dimension  $z'$  and evenly split the first two dimensions  $(x', y')$  into 2D grids  $(u, v)$ :

$$u = \frac{x' - x_0}{g_h} + o_h, \quad v = \frac{y' - y_0}{g_w} + o_w \quad (4)$$

where  $(x_0, y_0)$  is the minimum value of  $(x', y')$ ,  $(g_h, g_w)$  is the grid size,  $(o_h, o_w)$  is the offset value to place the projected object at the center of the image.  $0 \leq u \leq h - 1, 0 \leq v \leq w - 1$  and  $(u, v)$  is a sampled pixel coordinate from the 2D grid.

Now let us begin to derive formulations of the optical line that passes through the query grid. We only know  $(u, v)$  for each grid and we want to reversely trace which 3D points  $(x, y, z)$  are on the same optical line during parallel light projection. According to Eq. 4:

$$\begin{aligned} x' &= g_h u + x_0 - o_h = \Psi_h(u) \\ y' &= g_w v + y_0 - o_w = \Psi_w(v) \end{aligned} \quad (5)$$

If we denote  $A = R^{-1}$  and  $A_{ij}$  as the element at  $i^{\text{th}}$  row and  $j^{\text{th}}$  column, then according to Eq. 3:

$$\begin{aligned} x &= A_{11}\Psi_h(u) + A_{12}\Psi_w(v) + A_{13}z' = \Omega_x(u, v) + A_{13}z' \\ y &= A_{21}\Psi_h(u) + A_{22}\Psi_w(v) + A_{23}z' = \Omega_y(u, v) + A_{23}z' \\ z &= A_{31}\Psi_h(u) + A_{32}\Psi_w(v) + A_{33}z' = \Omega_z(u, v) + A_{33}z' \end{aligned} \quad (6)$$

According to the definition of line’s parametric equation, Eq. 6 represents a line passing through the origin point  $O : (\Omega_x(u, v), \Omega_y(u, v), \Omega_z(u, v))$  with optical line direction  $\mathbf{d} = (A_{13}, A_{23}, A_{33})$ , where  $\Omega_x, \Omega_y, \Omega_z$  are  $xyz$  coordinates of  $O$  and their formulations are conditioned on  $u, v$ . Therefore, we concatenate the coordinate of origin point  $O$ , normalized direction  $\mathbf{d}^\dagger = \mathbf{d}/\|\mathbf{d}\|_2$  and normalized position  $(u/h, v/w)$  as positional embedding together to be the initial state of our query. A multi-layer-perceptron (MLP) module is later leveraged to map the 8-dim initial query to higher dimensional space.

**Memory Builder.** The memory builder takes  $F_{3d}$  as input to prepare for initial state of  $K, V$  in cross-attention layers. We

first concatenate aligned 3D coordinate  $P_{3d}$  with 3D features to enhance the geometric knowledge of  $F_{3d}$ :

$$\hat{F}_{3d} = \text{MLP}(\text{cat}(F_{3d}, P_{3d})) \quad (7)$$

Additionally, we initialize a learnable memory token  $T_{\text{pad}}$  as the pad token and concatenate it with  $\hat{F}_{3d}$  to obtain the initial state of  $K, V$ . The reason for concatenating a learnable pad token  $T_{\text{pad}}$  is that there are white background areas on the projected image (as shown in Figure 2). As  $F_{3d}$  only encodes foreground objects, we further need a learnable pad token to represent background regions. Otherwise, the cross-attention layers will be confused to learn how to combine foreground tokens into background features and this will inevitably diminish the pre-training effectiveness.

### 3.4. Objective Function

We perform per-pixel supervision with Mean Squared Error (MSE) loss between generated view image  $I_{\text{gen}}^R$  and ground truth image  $I_{\text{gt}}^R$ , aligned by camera pose  $R$ . For simplicity, we will omit  $R$  in later formulations. As the background of the rendered ground truth images is all white and reveals little information, we further design a compound loss to balance the weight between foreground regions and background regions:

$$\mathcal{L}(I_{\text{gen}}, I_{\text{gt}}) = w^{\text{fg}}\mathcal{D}^{\text{fg}} + w^{\text{bg}}\mathcal{D}^{\text{bg}} \quad (8)$$

$$\mathcal{D}^k(I_{\text{gen}}^k, I_{\text{gt}}^k) = \frac{1}{HW} \sum_{h,w} (I_{\text{gen}}^k(h,w) - I_{\text{gt}}^k(h,w))^2 \quad (9)$$

where  $k = \text{fg}$  (foreground),  $\text{bg}$  (background) and  $w^{\text{fg}}, w^{\text{bg}}$  are loss weights for foreground and background, respectively. Such per-pixel supervision is more precise than the ambiguous set-to-set Chamfer Distance introduced in Eq. 1.

## 4. Experiments

In this section, we first introduce the setups of our pre-training scheme. Then we evaluate our pre-training method on various point cloud backbones by fine-tuning them on different downstream tasks, such as point cloud classification on ModelNet and ScanObjectNN datasets, and part segmentation on the ShapeNetPart dataset. Finally, we provide in-depth ablation studies for the architectural design of our proposed TAP pre-training pipeline.

### 4.1. Pre-training Setups

**Data Setups** To align with previous research practices [58, 30, 54], we choose **ShapeNet** [6] that contains more than 50 thousand CAD models as our pre-training datasets. We sampled 1024 points from each 3D CAD model to form the point clouds, consistent with previous work. Since ShapeNet does not provide images for each point cloud, we use the

Table 1: **Classification results on the ScanObjectNN dataset.** We report the overall accuracy (%). The results with † are reproduced by PointNeXt [35] repository.

Method	OBJ_BG	OBJ_ONLY	PB_T50_RS
<i>Hierarchical Models with TAP Pre-training</i>			
†DGCNN [48]	-	-	86.1
+ TAP	-	-	86.6 (+0.5)
†PointNet++ [33]	-	-	86.2
+ TAP	-	-	86.8 (+0.6)
†PointMLP [27]	-	-	87.4
+ TAP	-	-	88.5 (+1.1)
<i>Standard Transformers with Generative Pre-training</i>			
w/o pre-training [46]	79.86	80.55	77.24
OcCo [47]	84.85	85.54	78.79
Point-BERT [58]	87.43	88.12	83.07
MaskPoint [23]	89.30	88.10	84.30
Point-MAE [30]	90.02	88.29	85.18
TAP (Ours)	<b>90.36</b>	<b>89.50</b>	<b>85.67</b>

rendered image from 12 surrounding viewpoints generated by MVCNN [42]. During our pre-training, the models are exclusively pre-trained with the training split following the practice of previous work [58].

**Architecture Setups** We conduct experiments on various point cloud encoders, including PointNet++ [33], DGCNN [48], PointMLP [27] and Transformers [58] for point cloud object classification. During the pre-training stage, the photograph module takes encoded point cloud features and pose conditions as inputs to generate a  $32 \times$  downsampled view image feature map of size  $7 \times 7$  from a specific viewpoint. Then the 2D generator progressively upsamples the image feature map to decode RGB view images of size  $224 \times 224$ . We do not alter the architecture of the point cloud backbone since the photograph module and the 2D generator are exclusively used during the pre-training phase and are dropped during the fine-tuning stage. In our experiment, the photograph module is a six-layer cross-attention block, with attention layer channels limited to 256 to enhance efficiency. During the pre-training task on ShapeNet, we utilized four simple transpose convolutions to upsample the reconstructed 2D feature map and predict the RGB value for each pixel.

**Implementation Details** The experiments of TAP pre-training and finetuning on various downstream tasks are implemented with PyTorch [31]. We utilize AdamW [26] optimizer and the CosineAnnealing learning rate scheduler [25] to pre-train the point cloud backbone for 100 epochs. We set the initial learning rate as  $5e^{-4}$  and weight decay as  $5e^{-2}$ . In our experiment, we train various point cloud backbones with 32 batch sizes on a single Nvidia 3090Ti GPU. The

Table 2: **Results comparisons with previous methods on the ScanObjectNN and ModelNet40 datasets.** The model parameters number (#Params) and overall accuracy (%) are reported. † denotes our reproduced results of PointMLP on the ModelNet40 dataset. Methods with \* introduce extra knowledge from pre-trained image models or pre-trained vision-language models. We do not compete with them for fair comparison and only list them for reference.

Method	#Params (M)	ScanObjectNN	ModelNet40
<i>Supervised Learning Only</i>			
PointNet [32]	3.5	68.0	89.2
PointNet++ [33]	1.5	77.9	90.7
Transformer [46]	22.1	77.24	91.4
DGCNN [48]	1.8	78.1	92.9
PointCNN [22]	0.6	78.5	92.2
DRNet [36]	-	80.3	93.1
SimpleView [15]	-	80.5±0.3	93.9
GBNet [37]	8.8	81.0	93.8
PRA-Net [9]	2.3	81.0	93.7
MVTN [17]	11.2	82.8	93.8
RepSurf-U [38]	1.5	84.3	<b>94.4</b>
PointMLP [27]	12.6	85.4±0.3	93.5†
PointNeXt [35]	1.4	87.7±0.4	93.7
<i>Transformers with Pre-training</i>			
OcCo [47]	22.1	78.8	92.1
Point-BERT [58]	22.1	83.1	93.2
MaskPoint [23]	22.1	84.3	93.8
Point-MAE [30]	22.1	85.2	93.8
Point-M2AE [60]	15.3	86.4	94.0
<i>With Pre-trained Image Model</i>			
ACT [13]*	22.1	88.2	93.7
I2P-MAE [61]*	-	90.1	93.7
ReCon [34]*	44.3	90.6	94.1
<i>Our Proposed TAP Pre-training</i>			
PointMLP + TAP	12.6	<b>88.5</b>	94.0

drop path rate of the cross-attention layer is set to 0.1. The foreground and background loss weights  $w^{fg}$ ,  $w^{bg}$  are set to 20 and 1. The detailed architecture of our simple 2D generator is:  $\text{TConv}(256, 128, 5, 4) \rightarrow \text{TConv}(128, 64, 3, 2) \rightarrow \text{TConv}(64, 32, 3, 2) \rightarrow \text{TConv}(32, 3, 3, 2)$ , where TConv stands for Transpose Convolution and the four numbers in the tuple denotes  $(C_{in}, C_{out}, \text{Kernel}, \text{Stride})$  respectively. During the fine-tuning stage, we perform a learning rate warming up for point cloud backbones with 10 epochs, and keep other settings unchanged for a fair comparison.

## 4.2. Downstream Tasks

In this section, we report the experimental results of various downstream tasks. We follow the previous work to conduct experiments of object classification on real-world ScanObjectNN and synthetic ModelNet40 datasets. We also verify the effectiveness of our pre-training method on the




Figure 3: **The visualization results of our proposed 3D-to-2D generative pre-training.** The first row displays view images generated by our TAP pre-training pipeline and the second row shows ground truth images. Our TAP can produce view images with appropriate shapes and reflection colors, demonstrating its ability in capturing geometric structure and stereoscopic knowledge.

part segmentation task with the ShapeNetPart dataset.

### 4.2.1 Object Classification

**Main Results.** To evaluate the effectiveness of our proposed TAP, we implement it with various point cloud architectures, including classical baselines such as PointNet++, DGCNN, and PointMLP, as well as widely used Standard Transformers backbone for existing generative pre-training methods. We follow the common practice to experiment our model on three variants of the ScanObjectNN dataset: 1) OBJ-ONLY: cropping object without any background; 2) OBJ-BG: containing the background and object; 3) PB-T50-RS: adopting various augmentations to the objects. We reported the results comparing with existing pre-training methods in Table 1.

As shown in the upper part of the table, we pre-train the graph-based architecture DGCNN, set-abstraction-based architecture PointNet++, and MLP-based architecture PointMLP with TAP, and observe consistent improvements across models. The results strongly convey that our proposed TAP can be successfully applied to various types of point cloud models and the proposed novel 3D-to-2D generative pre-training is effective regardless of the backbone architecture. Considering that nearly all existing generative pre-training methods are specially designed for Transformer-based architecture, our TAP is much superior in its wider adaptation and higher flexibility. Additionally, we also provide a detailed and fair comparison with previous work by implementing TAP with the Standard Transformers architecture in the lower part of the table, where no hierarchical designs or inductive bias is included. Our TAP outperforms previous pre-training methods in all three split settings, providing strong evidence that our 3D-to-2D generative pre-training strategy can also benefit attention-based architectures and surpass uni-modal generative pre-training competitors. It is worth noting that although TAP and many previous

Table 3: **Part segmentation results on the ShapeNetPart dataset.** We report the mean IoU across all part categories  $mIoU_C$ , the mean IoU across all instances  $mIoU_I$ , and the IoU for each category.

Methods	$mIoU_C$	$mIoU_I$	aero	bag	cap	car	chair	earphone	guitar	knife	lamp	laptop	motor	mug	pistol	rocket	skateboard	table
<i>Supervised Representation Learning Only</i>																		
PointNet [32]	80.4	83.7	83.4	78.7	82.5	74.9	89.6	73.0	91.5	85.9	80.8	95.3	65.2	93.0	81.2	57.9	72.8	80.6
PointNet++ [33]	81.9	85.1	82.4	79.0	87.7	77.3	90.8	71.8	91.0	85.9	83.7	95.3	71.6	94.1	81.3	58.7	76.4	82.6
DGCNN [48]	82.3	85.2	84.0	83.4	86.7	77.8	90.6	74.7	91.2	87.5	82.8	95.7	66.3	94.9	81.1	63.5	74.5	82.6
PointMLP [27]	84.6	86.1	83.5	83.4	87.5	80.5	90.3	78.2	92.2	88.1	82.6	96.2	77.5	95.8	85.4	64.6	83.3	84.3
KPCConv [44]	85.1	86.4	84.6	86.3	87.2	81.1	91.1	77.8	92.6	88.4	82.7	96.2	78.1	95.8	85.4	69.0	82.0	83.6
<i>Transformers with Uni-modal Generative Pre-training</i>																		
Point-BERT [58]	84.1	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 [30]	84.2	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
MaskPoint [23]	84.4	86.0	84.2	85.6	88.1	80.3	91.2	79.5	91.9	87.8	86.2	95.3	76.9	95.0	85.3	64.4	76.9	81.8
Point-M2AE [60]	84.9	86.5	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
<i>Our Proposed 3D-to-2D Generative Pre-training</i>																		
PointMLP+TAP	<b>85.2</b>	<b>86.9</b>	84.8	86.1	89.5	82.5	92.1	75.9	92.3	88.7	85.6	96.5	79.8	96.0	85.9	66.2	78.1	83.2

pre-training approaches have significantly improved the performance of Transformers in point cloud tasks, increasing the accuracy from 77.24 to 85.67, they still lag far behind advanced point cloud networks such as PointMLP. Therefore, TAP’s applicability to various models is an important characteristic that would benefit future research.

**Comparisons with Previous Methods.** To clearly demonstrate the high performance of our proposed TAP pre-training in object classification tasks, we compare TAP with existing methods on both synthetic ModelNet and real-world ScanObjectNN (the hardest PB-T50-RS variant) datasets in Table 2. We categorize existing methods into two types: (1) architecture-oriented methods, which focus on developing novel model architectures for 3D point clouds and do not involve any pre-training techniques, and (2) pre-training methods, which pay more attention to the pre-training strategy and whose backbone model are mostly limited to Transformer-based architecture. It’s worth noticing that methods marked with an asterisk (\*) incorporate additional knowledge from pre-trained image models or pre-trained vision-language models. To ensure a fair and unbiased comparison, we refrain from directly comparing our method with these approaches. However, we include them in the listing for reference purposes, acknowledging their existence and potential relevance in related research.

From the experimental results, we can see that accompanied by PointMLP backbone model, our proposed TAP pre-training achieves the best classification accuracy on ScanObjectNN and ModelNet40 among existing models (with no pre-trained knowledge from image or language like P2P [49]), demonstrating the effectiveness of our approach and validating the superiority of our 3D-to-2D cross-modal generative pre-training method over previous generative pre-training methods. Moreover, we also note that our proposed method has brought higher performance improvements on the ScanObjectNN dataset than on ModelNet40. This may

be attributed to the reason that the cross-modal generative pre-training has enhanced the network’s ability to understand point clouds from different views, which is beneficial for a more robust understanding of the real-scan data with more noise and disturbance in the ScanObjectNN dataset.

**Visualization Results.** Figure 3 shows the visualization results of TAP. The first row shows the generated view images while the second row displays the ground truth images for reference. The TAP method can successfully predict the accurate shape of the object and the RGB colors that represent light reflections in rendered images. Therefore, TAP is capable of capturing the geometric structure of 3D objects and reasoning occlusion relations from specific camera poses.

#### 4.2.2 Part Segmentation

Performing dense prediction is always a more challenging task compared with classification. In this section, we evaluate the local distinguishability of our proposed TAP pre-training method, fine-tuning the pre-trained point cloud model on the ShapeNetPart dataset for the part segmentation task. Quantitative results are shown in Table 3. We implement PointMLP as the backbone model and compare our TAP results with two mainstreams of previous literature. The upper row displays classical architecture-oriented methods that focus on network design and are trained from scratch. The lower row shows members of the generative pre-training family that rely on Transformer-based architectures.

According to results comparisons, our TAP pre-training significantly improves the part segmentation performance of the PointMLP backbone, increasing class  $mIoU$  by 0.6 and instance  $mIoU$  by 0.8. More importantly, our TAP pre-training achieves state-of-the-art performance on both class  $mIoU$  and instance  $mIoU$ , surpassing leading works in both tracks. Specifically, TAP exceeds the performance of Point-M2AE on instance  $mIoU$  by 0.4. This satisfactory perfor-

mance serves as strong evidence to convey that our proposed TAP pre-training is superior to previous uni-model generative pre-training mechanisms in dense prediction tasks. This may be attributed to the factor that our supervision in 2D with MSE loss is more precise than the ambiguous Chamfer Distance in 3D reconstruction. Therefore, models with TAP pre-training obtain more accurate comprehension of local geometry and detail awareness, which contributes to mIoU gain in dense prediction tasks. What’s more, TAP outperforms KPConv on instance mIoU by 0.5, demonstrating that the proposed 3D-to-2D generative pre-training method can fully exploit the potential of the point cloud model and help it have a better perception of objects’ geometric structure. As TAP is adaptable to any architecture, future improvements in architectural design can also benefit from TAP pre-training.

### 4.2.3 Few-shot Classification

Following Point-BERT [58], we conduct few-shot classification with Standard Transformers on ModelNet40 [52] dataset. As shown in Table 4, we report mean overall accuracy and standard deviation (mOA±std) on 10 pre-defined data folders for each few-shot setting. The *way* and *shot* in Table 4 specify the number of categories and the number of training examples per category, respectively.

From the results, TAP achieves the highest mean overall accuracy across all few-shot settings when compared to previous generative pre-training approaches. Furthermore, TAP exhibits significantly lower standard deviations than those reported in the existing literature for the majority of few-shot settings, which signifies its robust performance and consistent superiority. This indicates that TAP is not only capable of achieving high mean overall accuracy but also exhibits reliability and robustness across various few-shot settings. Such stability is crucial in real-world applications, where consistency and predictability are vital for practical deployment.

### 4.3. Scene-level Dense Predictions

To assess the effectiveness of TAP in handling scene-level dense prediction tasks, we carry out experiments on more complicated scene-level object detection and semantic segmentation on the ScanNetV2 [12] dataset. For the object detection task, we adopt 3DETR [29] and pre-train its encoder on the object-level dataset ShapeNet [6] with TAP. Average precision at 0.25 and 0.5 IoU thresholds are reported. Regarding semantic segmentation, we employ the PointTransformerV2 [51](PTv2) model and pre-train it on the ScanNetV2 dataset with TAP. We report mean IoU for evaluation metric. It is worth mentioning that PTv2 represents the current state-of-the-art approach with open-source code availability.

Based on the results presented in Table 5, TAP consis-

Table 4: **Few-shot Classification with Standard Transformers on ModelNet40 dataset.** We report mean overall accuracy and standard deviation on 10 pre-defined data folders for each setting. Best results are marked **bold**.

Method	5-way		10-way	
	10-shot	20-shot	10-shot	20-shot
w/o pre-training	87.8 ± 5.2	93.3 ± 4.3	84.6 ± 5.5	89.4 ± 6.3
Point-BERT [58]	94.6 ± 3.1	96.3 ± 2.7	91.0 ± 5.4	92.7 ± 5.1
MaskPoint [23]	95.0 ± 3.7	97.2 ± <b>1.7</b>	91.4 ± 4.0	93.4 ± 3.5
Point-MAE [30]	96.3 ± 2.5	<b>97.8</b> ± 1.8	92.6 ± 4.1	95.0 ± 3.0
<b>TAP</b>	<b>97.3 ± 1.8</b>	<b>97.8 ± 1.7</b>	<b>93.1 ± 2.6</b>	<b>95.8 ± 1.0</b>

Table 5: **Scene-level object detection and semantic segmentation on ScanNetV2 [12].** Average precision at 0.25 IoU thresholds (AP<sub>0.25</sub>) and 0.5 IoU thresholds (AP<sub>0.5</sub>) of detection and mean Intersection-over-Union (mIoU) of semantic segmentation are reported.

Method	Det (3DETR [29])		Seg (PTv2 [51])
	AP <sub>0.25</sub>	AP <sub>0.5</sub>	mIoU
Baseline	62.1	37.9	72.4
+TAP	63.0 (+0.9)	41.4 (+3.5)	72.6 (+0.2)

tently enhances the performance of all baselines, thereby showcasing its efficacy in tackling more intricate scene-level dense prediction tasks. Remarkably, even with the encoder solely pre-trained on an object-level dataset for scene-level detection task, significant improvements are observed in both AP<sub>0.25</sub> and AP<sub>0.5</sub> metrics. This suggests that the learned representations from TAP effectively capture relevant information and generalize well to complex scenes, even when the pre-training data is limited to object-level collections. Such generalization capabilities are valuable in scenarios where obtaining large-scale fully annotated scene-level datasets may be challenging or expensive.

### 4.4. Ablation Studies

To investigate the architectural design of our proposed *Photograph Module* in TAP pre-training pipeline, we conduct extensive ablation studies on the ScanObjectNN dataset with PointMLP as the backbone model.

**Photograph Module Architectural Designs.** In Photograph Module, we implement cross-attention layers to generate view image feature maps conditioned on pose instruction. We believe that letting the module learn by itself how to rearrange 3D point features in 2D grids will enhance the representation ability of the 3D backbone. Therefore, we conduct ablation studies to verify this hypothesis. As shown in Table 6a, we implement Model A<sub>1</sub> with no attention lay-



Table 6: **Ablation studies on Photograph module in TAP pre-training pipeline.** We choose PointMLP as the backbone model and conduct ablation studies on the ScanObjectNN dataset from two aspects: overall architectural designs and query designs. In Table(a), we first investigate the effectiveness of cross-attention design compared with direct projection (Model  $A_1$ ) and direct project with self-attention (Model  $A_2$ ). Then we analyze the influence of the different number of attention layers and feature channels in Model B and Model C. Finally, we discuss whether pad token in memory builder is beneficial in Model D. In Table(b), we conduct further experiments to compare different approaches for query designing: (1) Using learnable query based on the given viewpoints (Model E) or use mathematical formulations derived in Eq. 6. (2) The information we need when we mathematically encode pose information into init query status: (i) **Origin**: the coordinate of origin point  $O$  that the optical line passes through. (ii) **Direction**: The normalized direction of the optical line. (iii) **PE**: The position embedding for each grid.

(a) Overall Architectural Designs.						(b) Query Designs.					
Model	Attention Type	LayerNum.	Channels	Mem.Pad	Acc.(%)	Model	Query Type	Origin	Direction	PE	Acc.(%)
$A_1$	None	–	256	✗	87.6 (-0.9)	E	Learnable	✗	✗	✗	87.5 (-1.0)
$A_2$	SelfAttn	6 layers	256	✗	87.8 (-0.7)	$F_1$	Formula	✗	✓	✓	86.5 (-2.0)
B	CrossAttn	2 layers	256	✓	87.9 (-0.6)	$F_2$	Formula	✓	✗	✗	87.8 (-0.7)
C	CrossAttn	6 layers	512	✓	87.8 (-0.7)	$F_3$	Formula	✓	✓	✗	88.0 (-0.5)
D	CrossAttn	6 layers	256	✗	88.3 (-0.2)	$F_4$	Formula	✓	✗	✓	88.1 (-0.4)
TAP	CrossAttn	6 layers	256	✓	88.5	TAP	Formula	✓	✓	✓	88.5

ers, directly projecting 3D feature points to 2D grids based on Eq. 3 and Eq. 4. This results in a much simpler pre-training task, as the projection relation has been directly told. Additionally, in Model  $A_2$ , we add self-attention layers after explicit projection to help the model capture longer-range correlations. Pose knowledge is encoded as a pose token that is concatenated to projected grids, similar to the CLS token in classification Transformers. According to quantitative results comparison with TAP that implements cross-attention layers, fine-tuning results of pre-training methods in  $A_1$  and  $A_2$  version show inferiority. Therefore, the cross-attention architecture we designed to entirely LEARN the projection relation is the most suitable choice for the proposed 3D-to-2D generative pre-training.

What’s more, we discuss the number of cross-attention layers, the dimension of feature channels and whether to concatenate pad token in memory builder in Model  $B, C, D$ . According to the results, more cross-attention layers show stronger representation ability, while too large channel number will lead to performance decrease caused by over-fitting. The performance gain from Model  $D$  to TAP also verifies that the pad token design in the memory builder is essential.

**Query Generator Designs.** In the query generator, we derive the mathematical formulation of the optical lines passing through 2D grids. We propose to concatenate the coordinate of origin point  $O$ , normalized direction  $\mathbf{d}^\dagger$  and position embedding  $(u/h, v/w)$  as the initial state of queries. In Table 6b, we first compare this mathematical design with totally learnable queries that takes pose matrix  $R$  as input and implements MLP layers to predict query for each grid. As shown in Model  $E$ , learnable queries cannot satisfactorily encode pose information, while our derived formulation for query construction is both clearer in physical meaning and

more competitive in fine-tuning accuracy.

In ablation  $F_1$  to  $F_4$ , we progressively discuss the three components of query generation. Quantitative comparison with TAP verifies that every component is indispensable for query generation, where coordinates of origin points are of the most importance.

## 5. Conclusions

In this paper, we have proposed a novel 3D-to-2D generative pre-training method TAP that is adaptable to any point cloud model. We implemented the cross-attention mechanism to generate view images of point clouds from instructed camera poses. To better encode pose conditions and generate physically meaningful queries, we derived mathematical formulations of optical lines. The proposed TAP pre-training had higher preciseness in supervision and broader adaptation to different backbones, compared with directly reconstructing point clouds in previous methods. Experimental results conveyed that the TAP pre-training can help the backbone models better capture the structural knowledge and stereoscopic relations. Fine-tuning results of TAP pre-training achieve state-of-the-art performance on ScanObjectNN classification and ShapeNetPart segmentation, among methods that do not include any pre-trained image or text models. We believe the cross-modal generative pre-training paradigm will be a promising direction for future research.

## Acknowledgement

This work was supported in part by the National Key Research and Development Program of China under Grant 2022ZD0114903 and in part by the National Natural Science Foundation of China under Grant 62125603.

## References

- [1] Idan Achituve, Haggai Maron, and Gal Chechik. Self-supervised learning for domain adaptation on point clouds. In *WACV*, 2021.
- [2] Mohamed Afham, Isuru Dissanayake, Dinithi Dissanayake, Amaya Dharmasiri, Kanchana Thilakarathna, and Ranga Rodrigo. Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding. In *CVPR*, pages 9902–9912, June 2022.
- [3] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *NeurIPS*, 2019.
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [6] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [7] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [8] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [9] Silin Cheng, Xiwu Chen, Xinwei He, Zhe Liu, and Xiang Bai. Pra-net: Point relation-aware network for 3d point cloud analysis. *IEEE TIP*, 30:4436–4448, 2021.
- [10] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, 2019.
- [11] Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, 2008.
- [12] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017.
- [13] Runpei Dong, Zekun Qi, Linfeng Zhang, Junbo Zhang, Jianjian Sun, Zheng Ge, Li Yi, and Kaisheng Ma. Autoencoders as cross-modal teachers: Can pretrained 2d image transformers help 3d representation learning? *arXiv preprint arXiv:2212.08320*, 2022.
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [15] Ankit Goyal, Hei Law, Bowei Liu, Alejandro Newell, and Jia Deng. Revisiting point cloud shape classification with a simple and effective baseline. In *ICML*, 2021.
- [16] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, 2018.
- [17] Abdullah Hamdi, Silvio Giancola, and Bernard Ghanem. MVTN: multi-view transformation network for 3d shape recognition. In *ICCV*, 2021.
- [18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021.
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.
- [20] Tianxin Huang, Zhonggan Ding, Jiangning Zhang, Ying Tai, Zhenyu Zhang, Mingang Chen, Chengjie Wang, and Yong Liu. Learning to measure the point cloud reconstruction loss in a representation space. In *CVPR*, 2023.
- [21] Roman Klokov and Victor Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In *ICCV*, 2017.
- [22] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *NeurIPS*, 2018.
- [23] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. In *ECCV*. Springer, 2022.
- [24] Minghua Liu, Lu Sheng, Sheng Yang, Jing Shao, and Shi-Min Hu. Morphing and sampling network for dense point cloud completion. In *AAAI*, 2020.
- [25] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018.
- [27] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point cloud: A simple residual mlp framework. In *ICLR*, 2022.
- [28] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *IROS*. IEEE, 2015.
- [29] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *ICCV*, 2021.
- [30] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *ECCV*. Springer, 2022.
- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 2019.
- [32] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.
- [33] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++ deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017.

- [34] Zekun Qi, Runpei Dong, Guofan Fan, Zheng Ge, Xiangyu Zhang, Kaisheng Ma, and Li Yi. Contrast with reconstruct: Contrastive 3d representation learning guided by generative pretraining. *arXiv preprint arXiv:2302.02318*, 2023.
- [35] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. In *NeurIPS*, 2022.
- [36] Shi Qiu, Saeed Anwar, and Nick Barnes. Dense-resolution network for point cloud classification and segmentation. In *WACV*, 2021.
- [37] Shi Qiu, Saeed Anwar, and Nick Barnes. Geometric back-projection network for point cloud classification. *IEEE TMM*, 2022.
- [38] Haoxi Ran, Jun Liu, and Chengjie Wang. Surface representation for point clouds. In *CVPR*, 2022.
- [39] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *CVPR*, 2017.
- [40] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.
- [41] Jonathan Sauder and Bjarne Sievers. Self-supervised deep learning on point clouds by reconstructing space. *NeurIPS*, 2019.
- [42] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *ICCV*, 2015.
- [43] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NeurIPS*, 2017.
- [44] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. Kpconv: Flexible and deformable convolution for point clouds. *ICCV*, 2019.
- [45] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *ICCV*, 2019.
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [47] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *ICCV*, 2021.
- [48] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *TOG*, 2019.
- [49] Ziyi Wang, Xumin Yu, Yongming Rao, Jie Zhou, and Jiwen Lu. P2p: Tuning pre-trained image models for point cloud analysis with point-to-pixel prompting. In *NeurIPS*, 2022.
- [50] Tong Wu, Liang Pan, Junzhe Zhang, Tai Wang, Ziwei Liu, and Dahua Lin. Balanced chamfer distance as a comprehensive metric for point cloud completion. In *NeurIPS*, 2021.
- [51] Xiaoyang Wu, Yixing Lao, Li Jiang, Xihui Liu, and Hengshuang Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. In *NeurIPS*, 2022.
- [52] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015.
- [53] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *CVPR*, 2020.
- [54] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *ECCV*, 2020.
- [55] Chenfeng Xu, Shijia Yang, Tomer Galanti, Bichen Wu, Xiangyu Yue, Bohan Zhai, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Image2point: 3d point-cloud understanding with 2d image pretrained models. *arXiv preprint arXiv:2106.04180*, 2021.
- [56] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *T&G*, 2016.
- [57] Xumin Yu, Yongming Rao, Ziyi Wang, Zuyan Liu, Jiwen Lu, and Jie Zhou. Pointr: Diverse point cloud completion with geometry-aware transformers. In *ICCV*, 2021.
- [58] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *CVPR*, 2022.
- [59] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *CVPR*, 2022.
- [60] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Point-m2ae: multi-scale masked autoencoders for hierarchical point cloud pre-training. *arXiv preprint arXiv:2205.14401*, 2022.
- [61] Renrui Zhang, Liuhui Wang, Yu Qiao, Peng Gao, and Hongsheng Li. Learning 3d representations from 2d pre-trained models via image-to-point masked autoencoders. In *CVPR*, 2023.