

Learning 3D Representations from 2D Pre-trained Models via Image-to-Point Masked Autoencoders

Renrui Zhang^{1,2}, Liuhui Wang³, Yu Qiao¹, Peng Gao^{†1}, Hongsheng Li^{2,4}

¹Shanghai Artificial Intelligence Laboratory

²CUHK MMLab ³Peking University ⁴CPII under InnoHK

{zhangrenrui, gaopeng}@pjlab.org.cn, hqli@ee.cuhk.edu.cn

Abstract

Pre-training by numerous image data has become de-facto for robust 2D representations. In contrast, due to the expensive data processing, a paucity of 3D datasets severely hinders the learning for high-quality 3D features. In this paper, we propose an alternative to obtain superior 3D representations from 2D pre-trained models via Image-to-Point Masked Autoencoders, named as *I2P-MAE*. By self-supervised pre-training, we leverage the well learned 2D knowledge to guide 3D masked autoencoding, which reconstructs the masked point tokens with an encoder-decoder architecture. Specifically, we first utilize off-the-shelf 2D models to extract the multi-view visual features of the input point cloud, and then conduct two types of image-to-point learning schemes. For one, we introduce a 2D-guided masking strategy that maintains semantically important point tokens to be visible. Compared to random masking, the network can better concentrate on significant 3D structures with key spatial cues. For another, we enforce these visible tokens to reconstruct multi-view 2D features after the decoder. This enables the network to effectively inherit high-level 2D semantics for discriminative 3D modeling. Aided by our image-to-point pre-training, the frozen *I2P-MAE*, without any fine-tuning, achieves **93.4%** accuracy for linear SVM on ModelNet40, competitive to existing fully trained methods. By further fine-tuning on ScanObjectNN’s hardest split, *I2P-MAE* attains the state-of-the-art **90.11%** accuracy, +3.68% to the second-best, demonstrating superior transferable capacity. Code is available at <https://github.com/ZrrSkywalker/I2P-MAE>.

1. Introduction

Driven by huge volumes of image data [11, 40, 56, 61], pre-training for better visual representations has gained

[†] Corresponding author

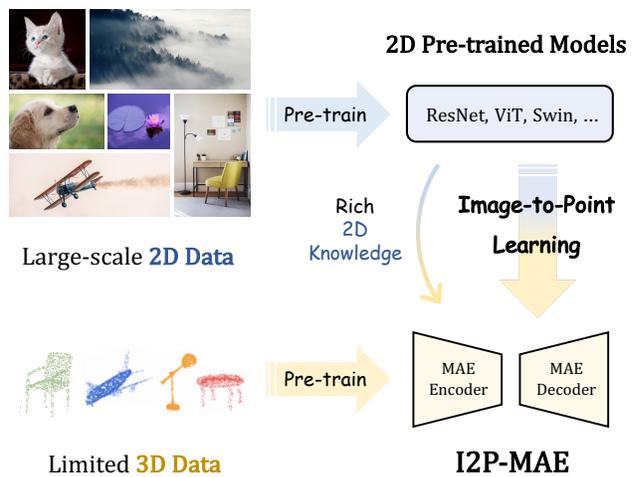


Figure 1. **Image-to-Point Masked Autoencoders.** We leverage the 2D pre-trained models to guide the MAE pre-training in 3D, which alleviates the need of large-scale 3D datasets and learns from 2D knowledge for superior 3D representations.

much attention in computer vision, which benefits a variety of downstream tasks [8, 29, 39, 55, 62]. Besides supervised pre-training with labels, many researches develop advanced self-supervised approaches to fully utilize raw image data via pre-text tasks, e.g., image-image contrast [5, 9, 10, 21, 28], language-image contrast [12, 53], and masked image modeling [3, 4, 18, 19, 27, 31, 42, 70]. Given the popularity of 2D pre-trained models, it is still absent for large-scale 3D datasets in the community, attributed to the expensive data acquisition and labor-intensive annotation. The widely adopted ShapeNet [6] only contains 50k point clouds of 55 object categories, far less than the 14 million ImageNet [11] and 400 million image-text pairs [53] in 2D vision. Though there have been attempts to extract self-supervisory signals for 3D pre-training [33, 41, 47, 64, 69, 76, 78, 83], raw point clouds with sparse structural patterns cannot provide suf-

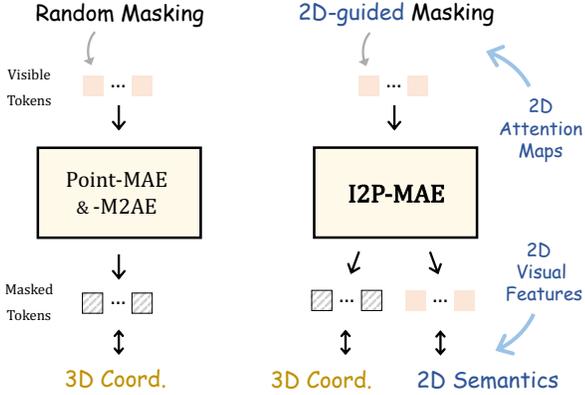


Figure 2. **Comparison of (Left) Existing Methods [47, 78] and (Right) our I2P-MAE.** On top of the general 3D MAE architecture, I2P-MAE introduces two schemes of image-to-point learning: 2D-guided masking and 2D-semantic reconstruction.

efficient and diversified semantics compared to colorful images, which constrain the generalization capacity of pre-training. Considering the homology of images and point clouds, both of which depict certain visual characteristics of objects and are related by 2D-3D geometric mapping, we ask the question: *can off-the-shelf 2D pre-trained models help 3D representation learning by transferring robust 2D knowledge into 3D domains?*

To tackle this challenge, we propose **I2P-MAE**, a Masked Autoencoding framework that conducts **Image-to-Point** knowledge transfer for self-supervised 3D point cloud pre-training. As shown in Figure 1, aided by 2D semantics learned from abundant image data, our I2P-MAE produces high-quality 3D representations and exerts strong transferable capacity to downstream 3D tasks. Specifically, referring to 3D MAE models [47, 78] in Figure 2 (Left), we first adopt an asymmetric encoder-decoder transformer [14] as our fundamental architecture for 3D pre-training, which takes as input a randomly masked point cloud and reconstructs the masked points from the visible ones. Then, to acquire 2D semantics for the 3D shape, we bridge the modal gap by efficiently projecting the point cloud into multi-view depth maps. This requires no time-consuming offline rendering and largely preserves 3D geometries from different perspectives. On top of that, we utilize off-the-shelf 2D models to obtain the multi-view 2D features along with 2D attention maps of the point cloud, and respectively guide the pre-training from two aspects, as shown in Figure 2 (Right).

Firstly, different from existing methods [47, 78] to randomly sample visible tokens, we introduce a 2D-guided masking strategy that reserves point tokens with more spatial semantics to be visible for the MAE encoder. In detail, we back-project the multi-view attention maps into 3D space as a spatial attention cloud. Each element in the

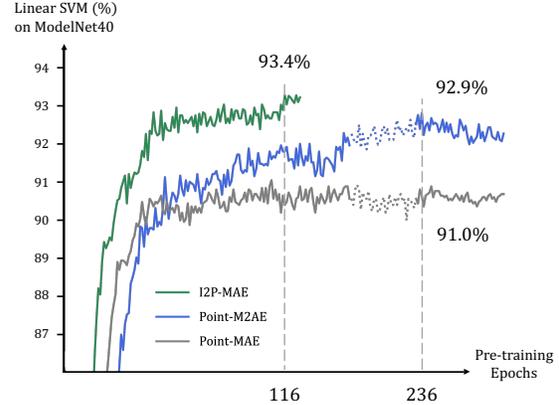


Figure 3. **Pre-training Epochs vs. Linear SVM Accuracy on ModelNet40 [67].** With the image-to-point learning schemes, I2P-MAE exerts superior transferable capability with much faster convergence speed than Point-MAE [47] and Point-M2AE [78].

cloud indicates the semantic significance of the corresponding point token. Guided by this, the 3D network can better focus on the visible critical structures to understand the global 3D shape, and also reconstruct the masked tokens from important spatial cues.

Secondly, in addition to the recovering of masked point tokens, we propose to concurrently reconstruct 2D semantics from the visible point tokens after the MAE decoder. For each visible token, we respectively fetch its projected 2D representations from different views, and integrate them as the 2D-semantic learning target. By simultaneously reconstructing the masked 3D coordinates and visible 2D concepts, I2P-MAE is able to learn both low-level spatial patterns and high-level semantics pre-trained in 2D domains, contributing to superior 3D representations.

With the aforementioned image-to-point guidance, our I2P-MAE significantly accelerates the convergence speed of pre-training and exhibits state-of-the-art performance on 3D downstream tasks, as shown in Figure 3. Learning from 2D ViT [14] pre-trained by CLIP [53], I2P-MAE, without any fine-tuning, achieves **93.4%** classification accuracy by linear SVM on ModelNet40 [67], which has surpassed the fully fine-tuned results of Point-BERT [76] and Point-MAE [33]. After fine-tuning, I2P-MAE further achieves **90.11%** classification accuracy on the hardest split of ScanObjectNN [63], significantly exceeding the second-best Point-M2AE [78] by +3.68%. The experiments fully demonstrate the effectiveness of learning from pre-trained 2D models for superior 3D representations.

Our contributions are summarized as follows:

- We propose Image-to-Point Masked Autoencoders (I2P-MAE), a pre-training framework to leverage 2D pre-trained models for learning 3D representations.

- We introduce two strategies, 2D-guided masking and 2D-semantic reconstruction, to effectively transfer the well learned 2D knowledge into 3D domains.
- Extensive experiments have been conducted to indicate the significance of our image-to-point pre-training.

2. Related Work

3D Point Cloud Pre-training. Supervised learning for point clouds has attained remarkable progress by delicately designed architectures [22, 50, 51, 80, 81] and local operators [2, 37, 65, 68, 72, 82]. However, such methods learned from closed-set datasets [63, 67] are limited to produce general 3D representations. Instead, self-supervised pre-training via unlabelled point clouds [6] have revealed promising transferable ability, which provides a good network initialization for downstream fine-tuning. Mainstream 3D self-supervised approaches adopt encoder-decoder architectures to recover the input point clouds from the transformed representations, including point rearrangement [57], part occlusion [64], rotation [49], downsampling [36], and codeword encoding [74]. Concurrent works also adopt contrastive pre-text tasks between 3D data pairs, such as local-global relation [17, 54], temporal frames [32], and augmented viewpoints [69]. More recent works adopt pre-trained CLIP [53] for zero-shot 3D recognition [24, 79, 86], or introduce masked point modeling [13, 16, 33, 41] as strong 3D self-supervised learners. Therein, Point-BERT [76] utilizes pre-trained tokenizers to indicate discrete point tokens, while Point-MAE [47], Point-M2AE [78], and PiMAE [7] apply Masked Autoencoders (MAE) to directly reconstruct masked 3D coordinates. Our I2P-MAE also adopts MAE as the basic pre-training framework, but is guided by 2D pre-trained models via image-to-point learning schemes, which benefits 3D pre-training with diverse 2D semantics.

2D-to-3D Learning. Except for jointly training 2D-3D networks [1, 23, 38, 46, 73], only a few existing researches focus on 2D-to-3D learning, and can be categorized into two groups. To eliminate the modal gap, the first group either upgrades 2D pre-trained networks into 3D variants for processing point clouds (convolutional kernels inflation [71, 77], modality-agnostic transformer [52]), or projects 3D point clouds into 2D images with parameter-efficient tuning (multi-view adapter [79], point-to-pixel prompting [66]). Different from them, our I2P-MAE benefits from three properties. **1) Pre-training:** We learn from 2D models during the pre-training stage, and then can be flexibly adapted to various 3D scenarios by fine-tuning. However, prior works are required to utilize 2D models on every different downstream task. **2) Self-supervision:** I2P-MAE is pre-trained by raw point clouds in a self-supervised manner and learns more general 3D representations. In con-

trast, prior works are supervised by labelled downstream 3D data, which might constrain the diverse 2D semantics into specific 3D domains. **3) Independence:** Prior works directly inherit 2D network architectures for 3D training, which is memory-consuming and not 3D extensible, while we train an independent 3D network and regard 2D models as semantic teachers. Another group of 2D-to-3D methods utilize paired real-world 2D-3D data of indoor [44] or outdoor [59] scenes, and conduct contrastive learning for knowledge transfer. Our approach also differs from them in two ways. **4) Masked Autoencoding.** I2P-MAE learns 2D semantics via the MAE-style pre-training without any contrastive loss between point-pixel pairs. **5) 3D Data Only.** We require no real-world image dataset during the pre-training and efficiently project the 3D shape into depth maps for 2D features extraction.

3. Method

The overall pipeline of I2P-MAE is shown in Figure 4. In Section 3.1, we first introduce I2P-MAE’s basic 3D architecture for point cloud masked autoencoding without 2D guidance. Then in Section 3.2, we show the details of utilizing 2D pre-trained models to obtain visual representations from 3D point clouds. Finally in Section 3.3, we present how to conduct image-to-point knowledge transfer for 3D representation learning.

3.1. Basic 3D Architecture

As our basic framework for image-to-point learning, I2P-MAE refers to existing works [47, 78] to conduct 3D point cloud masked autoencoding, which consists of a token embedding module, an encoder-decoder transformer, and a head for reconstructing masked 3D coordinates.

Token Embedding. Given a raw point cloud $P \in \mathbb{R}^{N \times 3}$, we first adopt Furthest Point Sampling (FPS) to downsample the point number from N to M , denoted as $P^T \in \mathbb{R}^{M \times 3}$. Then, we utilize k Nearest-Neighbour (k -NN) to search the k neighbors for each downsampled point, and aggregate their features via a mini-PointNet [50] to obtain M point tokens. In this way, each token can represent a local spatial region and interact long-range features with others in the follow-up transformer. We formulate them as $T \in \mathbb{R}^{M \times C}$, where C denotes the feature dimension.

Encoder-Decoder Transformer. To build the pre-text learning targets, we mask the point tokens with a high ratio, e.g., 80%, and only feed the visible ones, $T_{\text{vis}} \in \mathbb{R}^{M_{\text{vis}} \times C}$, into the transformer encoder, where M_{vis} denotes the visible token number. Each encoder block contains a self-attention layer and is pre-trained to understand the global

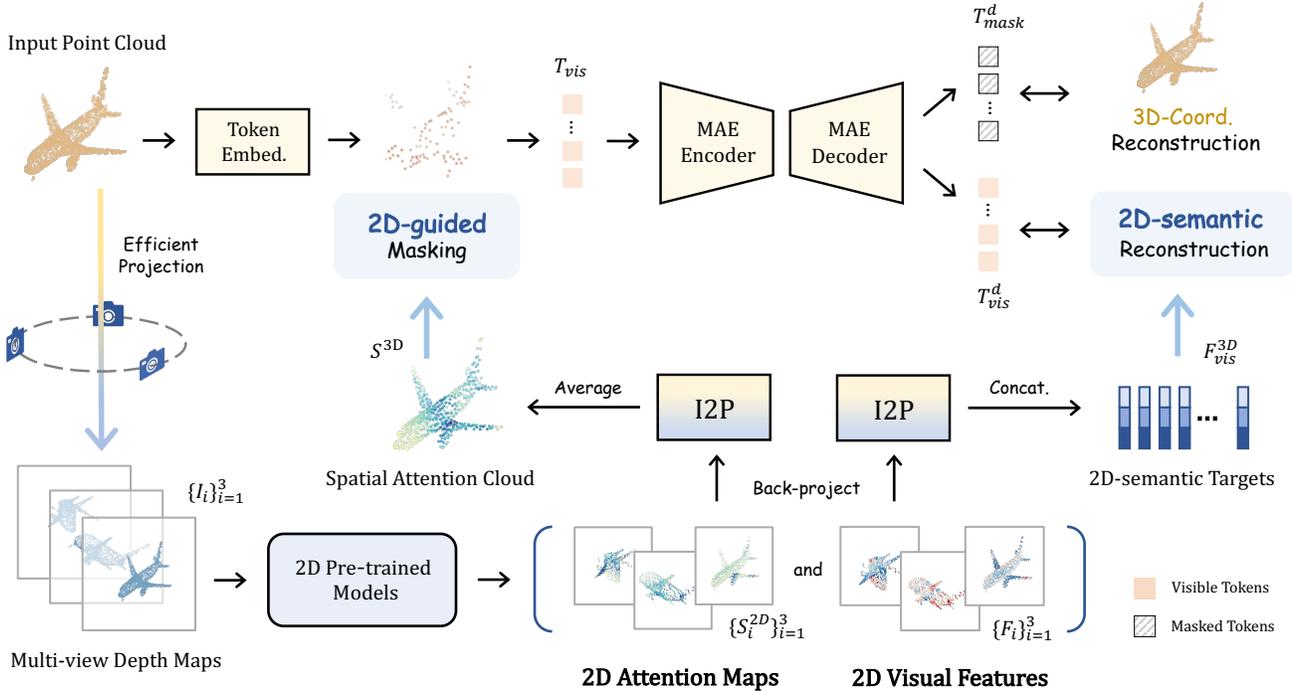


Figure 4. **The Pipeline of I2P-MAE.** Given an input point cloud, we leverage the 2D pre-trained models to generate two guidance signals from the projected depth maps: 2D attention maps and 2D visual features. We respectively conduct 2D-guided masking and 2D-semantic reconstruction to transfer the encoded 2D knowledge for 3D point cloud pre-training.

3D shape from the remaining visible parts. After encoding, we concatenate the encoded T_{vis}^e with a set of shared learnable masked tokens $T_{mask} \in \mathbb{R}^{M_{mask} \times C}$, and input them into a lightweight decoder, where M_{mask} denotes the masked token number and $M = M_{mask} + M_{vis}$. In the transformer decoder, the masked tokens learn to capture informative spatial cues from visible ones and reconstruct the masked 3D coordinates. Importantly, we follow PointM2AE [78] to modify our transformer to be a hierarchical architecture for encoding the multi-scale representations.

3D-coordinate Reconstruction. On top of the decoded point tokens $\{T_{vis}^d, T_{mask}^d\}$, we utilize T_{mask}^d to reconstruct 3D coordinates of the masked tokens along with their k neighboring points. A reconstruction head of a single linear projection layer is adopted to predict $P_{mask} \in \mathbb{R}^{M_{mask} \times k \times 3}$, the ground-truth 3D coordinates of the masked points. Then, we compute the loss by Chamfer Distance [15] and formulate it as

$$\mathcal{L}_{3D} = \frac{1}{M_{mask}k} \text{Chamfer} \left(H_{3D}(T_{mask}^d), P_{mask} \right), \quad (1)$$

where $H_{3D}(\cdot)$ denotes the head to reconstruct the masked 3D coordinates.

3.2. 2D Pre-trained Representations

We can leverage 2D models of different architectures (ResNet [30], ViT [14]) and various pre-training approaches (supervised [30, 45] and self-supervised [5, 53] ones) to assist the 3D representation learning. To align the input modality for 2D models, we project the input point cloud onto multiple image planes to create depth maps, and then encode them into multi-view 2D representations.

Efficient Projection. To ensure the efficiency of pre-training, we simply project the input point cloud P from three orthogonal views respectively along the x, y, z axes. For every point, we directly omit each of its three coordinates and round down the other two to obtain the 2D location on the corresponding map. The projected pixel value is set as the omitted coordinate to reflect relative depth relations of points, which is then repeated by three times to imitate the three-channel RGB. The projection of I2P-MAE is highly time-efficient and involves no offline rendering [48, 60], projective transformation [20, 79], or learnable prompting [66]. We denote the projected multi-view depth maps of P as $\{I_i\}_{i=1}^3$.

2D Visual Features. We then utilize a 2D pre-trained model, e.g., a pre-trained ResNet or ViT, to extract multi-

view features of the point cloud with C channels, formulated as $\{F_i^{2D}\}_{i=1}^3$, where $F_i \in \mathbb{R}^{H \times W \times C}$ and H, W denote the feature map size. Such 2D features contain sufficient high-level semantics learned from large-scale image data. The geometric information loss during projection can also be alleviated by encoding from different views.

2D Attention Maps. In addition to 2D features, we also acquire a semantic attention map for each view via the 2D pre-trained model. The one-channel maps indicate the semantic importance of different image regions, which we denote as $\{S_i^{2D}\}_{i=1}^3$, where $S_i \in \mathbb{R}^{H \times W \times 1}$. For ResNet, we adopt pixel-wise max pooling on $\{F_i^{2D}\}_{i=1}^3$ to reduce the feature dimension to one. For ViT, we adopt attention maps of the class token at the last transformer layer, since the attention weights to the class token reveal how much the features contribute to the final classification.

3.3. Image-to-Point Learning Schemes

On top of the 2D pre-trained representations for the point cloud, I2P-MAE’s pre-training is guided by two image-to-point learning designs: 2D-guided masking before the encoder, and 2D-semantic reconstruction after the decoder.

2D-guided Masking. The conventional masking strategy samples masked tokens randomly following a uniform distribution, which might prevent the encoder from ‘seeing’ important spatial characteristics and disturb the decoder by nonsignificant structures. Therefore, we leverage the 2D attention maps to guide the masking of point tokens, which samples more semantically significant 3D parts for encoding. Specifically, indexed by the coordinates of point tokens $P^T \in \mathbb{R}^{M \times 3}$, we back-project the multi-view attention maps $\{S_i^{2D}\}_{i=1}^3$ into 3D space and aggregate them as a 3D attention cloud $S^{3D} \in \mathbb{R}^{M \times 1}$. For each point in S^{3D} , we assign the semantic score by averaging the corresponding 2D values from multi-view attention maps as

$$S^{3D} = \text{Softmax} \left(\frac{1}{3} \sum_{i=1}^3 \text{I2P}(S_i^{2D}, P^T) \right), \quad (2)$$

where $\text{I2P}(\cdot)$ denotes the 2D-to-3D back-projection operation in Figure 5. We apply a softmax function to normalize the M points within S^{3D} , and regard each element’s magnitude as the visible probability for the corresponding point token. With this 2D semantic prior, the random masking becomes a nonuniform sampling with different probabilities for different tokens, where the tokens covering more critical 3D structures are more likely to be preserved. This boosts the representation learning of the encoder by focusing more on significant 3D geometries, and provides the masked tokens with more informative cues at the decoder for better reconstruction.

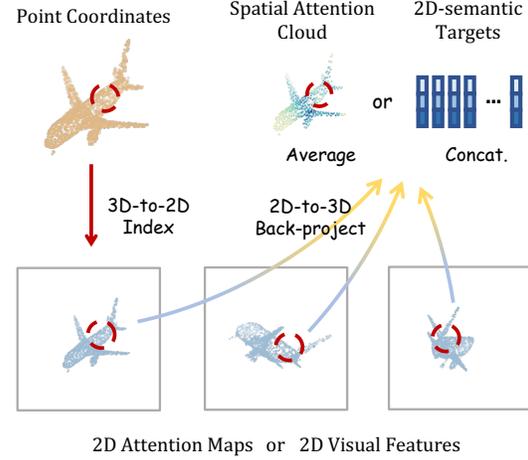


Figure 5. **Image-to-Point Operation (I2P).** Indexed by 3D point coordinates, the corresponding multi-view 2D representations are back-projected into 3D space for aggregation.

2D-semantic Reconstruction. The 3D-coordinate reconstruction of masked point tokens enables the network to explore low-level 3D patterns. On top of that, we further enforce the visible point tokens, T_{vis}^d , to reconstruct the extracted 2D semantics from different views, which effectively transfers the 2D pre-trained knowledge into 3D pre-training. To obtain the 2D-semantic targets, we utilize the coordinates of visible point tokens $P_{\text{vis}}^T \in \mathbb{R}^{M_{\text{vis}} \times 3}$ as indices to aggregate corresponding 2D features from $\{F_i^{2D}\}_{i=1}^3$ by channel-wise concatenation, formulated as

$$F_{\text{vis}}^{3D} = \text{Concat} \left\{ \text{I2P}(F_i^{2D}, P_{\text{vis}}^T) \right\}_{i=1}^3, \quad (3)$$

where $F_{\text{vis}}^{3D} \in \mathbb{R}^{M_{\text{vis}} \times 3C}$. As the multi-view depth maps depict a 3D shape from different perspectives, the concatenation between multi-view 2D features can better integrate the rich semantics inherited from 2D pre-trained models. Then, we also adopt a reconstruction head of a single linear projection layer for T_{vis}^d , and compute the l_2 loss as

$$\mathcal{L}_{2D} = \frac{1}{M_{\text{vis}}} \left(H_{2D}(T_{\text{vis}}^d) - F_{\text{vis}}^{3D} \right)^2, \quad (4)$$

where $H_{2D}(\cdot)$ denotes the head to reconstruct visible 2D semantics, parallel to $H_{3D}(\cdot)$ for masked 3D coordinates. The final pre-training loss of our I2P-MAE is then formulated as $\mathcal{L}_{\text{I2P}} = \mathcal{L}_{3D} + \mathcal{L}_{2D}$. With such image-to-point feature learning, I2P-MAE not only encodes low-level spatial variations from 3D coordinates, but also explores high-level semantic knowledge from 2D representations. As the 2D-guided masking has preserved visible tokens with more spatial significance, the 2D-semantic reconstruction upon them further benefits I2P-MAE to learn more discriminative 3D representations.

Method	ModelNet40	OBJ-BG
SO-Net [35]	87.3	-
FoldingNet [74]	88.4	-
VIP-GAN [26]	90.2	-
DGCNN + Jigsaw [58]	90.6	59.5
DGCNN + OcCo [64]	90.7	78.3
DGCNN + CrossPoint [1]	91.2	81.7
Transformer + OcCo [76]	89.6	-
Point-BERT [76]	87.4	-
Point-MAE [47]	91.0	77.7
3D-OAE [85]	92.3	-
Point-M2AE [78]	<u>92.9</u>	<u>84.1</u>
I2P-MAE	93.4	87.1
<i>Improvement</i>	<i>+0.5</i>	<i>+3.0</i>

Table 1. **Linear SVM Classification on ModelNet40 [67] and ScanObjectNN [63].** We compare the accuracy (%) of existing self-supervised methods, and the second-best one is underlined. We adopt the OBJ-BG split for ScanObjectNN following [1].

4. Experiments

In Section 4.1, we first introduce our pre-training settings and the linear SVM classification performance without fine-tuning. Then in Section 4.2, we present the results by fully fine-tuning I2P-MAE on various 3D downstream tasks. Finally in Section 4.3, we conduct ablation studies to investigate the characteristics of I2P-MAE.

4.1. Image-to-Point Pre-training

Settings. We adopt the popular ShapeNet [6] for self-supervised 3D pre-training, which contains 57,448 synthetic point clouds with 55 object categories. For fair comparison, we follow the same MAE transformer architecture as Point-M2AE [78]: a 3-stage encoder with 5 blocks per stage, a 2-stage decoder with 1 block per stage, 2,048 input point number (N), 512 downsampled number (M), 16 nearest neighbors (k), 384 feature channels (C), and the mask ratio of 80%. For off-the-shelf 2D models, we utilize ViT-Base [14] pre-trained by CLIP [53] as default and keep its weights frozen during 3D pre-training. We project the point cloud into three 224×224 depth maps, and obtain the 2D feature size $H \times W$ of 14×14 . I2P-MAE is pre-trained for 300 epochs with a batch size 64 and learning rate 10^{-3} . We adopt AdamW [34] optimizer with a weight decay 5×10^{-2} and the cosine scheduler with 10-epoch warm-up.

Linear SVM. To evaluate the transfer capacity, we directly utilize the features extracted by I2P-MAE’s encoder for linear SVM on the synthetic ModelNet40 [67] and real-world ScanObjectNN [63] without any fine-tuning or voting. As shown in Table 1, for 3D shape classification in

Method	OBJ-BG	OBJ-ONLY	PB-T50-RS
PointNet [50]	73.3	79.2	68.0
PointNet++ [51]	82.3	84.3	77.9
PointCNN [37]	86.1	85.5	78.5
MVTN [25]	-	-	82.8
PointMLP [2]	-	-	85.2
Transformer [76]	79.86	80.55	77.24
[P] Point-BERT [76]	87.43	88.12	83.07
[P] 3D-OAE [85]	89.16	88.64	83.17
[P] MaskPoint [41]	89.30	88.10	84.30
[P] Point-MAE [47]	90.02	88.29	85.18
[P] MAE3D [33]	-	-	86.20
[P] Point-M2AE [78]	<u>91.22</u>	<u>88.81</u>	<u>86.43</u>
[P] I2P-MAE	94.15	91.57	90.11
<i>Improvement</i>	<i>+2.93</i>	<i>+2.76</i>	<i>+3.68</i>

Table 2. **Real-world 3D Classification on ScanObjectNN [63].** We report the accuracy (%) on the official three splits of ScanObjectNN. [P] denotes to fine-tune the models after self-supervised pre-training.

both domains, I2P-MAE shows superior performance and exceeds the second-best respectively by +0.5% and +3.0% accuracy. Our SVM results (93.4%, 87.1%) can even surpass some existing methods after full downstream training in Table 2 and 3, e.g., PointCNN [37] (92.2%, 86.1%), Transformer [76] (91.4%, 79.86%), and Point-BERT [76] (92.7%, 87.43%). In addition, guided by the 2D pre-trained models, I2P-MAE exhibits much faster pre-training convergence than Point-MAE [47] and Point-M2AE in Figure 3. Therefore, the SVM performance of I2P-MAE demonstrates its learned high-quality 3D representations and the significance of our image-to-point learning schemes.

4.2. Downstream Tasks

After pre-training, I2P-MAE is fine-tuned for real-world and synthetic 3D classification, and part segmentation. Except ModelNet40 [67], we do not use the voting strategy [43] for evaluation.

Real-world 3D Classification. The challenging ScanObjectNN [63] consists of 11,416 training and 2,882 test 3D shapes, which are scanned from the real-world scenes and thus include backgrounds with noises. As shown in Table 2, our I2P-MAE exerts great advantages over other self-supervised methods, surpassing the second-best by +2.93%, +2.76%, and +3.68% respectively for the three splits. This is also the first model reaching 90% accuracy on the hardest PB-T50-RS split. As the pre-training point clouds [6] are synthetic 3D shapes with a large domain gap with the real-world ScanObjectNN, the results well indicate the universality of I2P-MAE inherited from the 2D pre-trained models.

Method	Points	No Voting	Voting
PointNet [50]	1k	89.2	-
PointNet++ [51]	1k	90.7	-
PointCNN [37]	1k	92.2	-
Point Transformer [84]	-	93.7	-
Transformer [76]	1k	91.4	-
[P] Point-BERT [76]	1k	92.7	93.2
[P] 3D-OAE [85]	1k	-	93.4
[P] Point-MAE [47]	1k	93.2	93.8
Point-MAE	8k	-	94.0
[P] Point-M2AE [78]	1k	93.4	94.0
[P] I2P-MAE-SVM	1k	93.4	-
I2P-MAE	1k	93.7	94.1

Table 3. **Synthetic 3D Classification on ModelNet40 [67]**. We report the accuracy (%) before and after the voting [43]. [P] denotes to fine-tune the models after self-supervised pre-training.

Synthetic 3D Classification. The widely adopted ModelNet40 [67] contains 9,843 training and 2,468 test 3D point clouds, which are sampled from the synthetic CAD models of 40 categories. We report the classification accuracy of existing methods before and after the voting in Table 3. As shown, our I2P-MAE achieves leading performance for both settings with only 1k input point number. For linear SVM without any fine-tuning, I2P-MAE-SVM can already attain 93.4% accuracy and exceed most previous works, indicating the powerful transfer ability. By fine-tuning the entire network, the accuracy can be further boosted by +0.3% accuracy, and achieve 94.1% after the offline voting.

Part Segmentation. The synthetic ShapeNetPart [75] is selected from ShapeNet [6] with 16 object categories and 50 part categories, which contains 14,007 and 2,874 samples for training and validation. We utilize the same segmentation head after the pre-trained encoder as previous works [47, 78] for fair comparison. The head only conducts simple upsampling for point tokens at different stages and concatenates them along the feature dimension as the output. Two types of mean IoU scores, $mIoU_C$ and $mIoU_I$ are reported in Table 4. For such highly saturated benchmark, I2P-MAE can still exert leading performance guided by the well learned 2D knowledge, e.g., +0.29% and +0.25% higher than Point-M2AE concerning the two metrics. This demonstrates that the 2D guidance also benefits the understanding for fine-grained point-wise 3D patterns.

4.3. Ablation Study

In this section, we explore the effectiveness of different characteristics in I2P-MAE. We utilize our final solution as the baseline and report the linear SVM accuracy (%).

Method	$mIoU_C$	$mIoU_I$
PointNet [50]	80.39	83.70
PointNet++ [51]	81.85	85.10
DGCNN [65]	82.33	85.20
PointMLP [2]	84.60	86.10
Transformer [76]	83.42	85.10
[P] Transformer + OcCo [76]	83.42	85.10
[P] Point-BERT [76]	84.11	85.60
[P] 3D-OAE [85]	-	85.70
[P] MaskPoint [41]	84.40	86.00
[P] Point-MAE [47]	-	86.10
[P] Point-M2AE [78]	84.86	86.51
[P] I2P-MAE	85.15	86.76

Table 4. **Part Segmentation on ShapeNetPart [75]**. We report the average IoU scores (%) for part categories and instances, respectively denoted as ‘ $mIoU_C$ ’ and ‘ $mIoU_I$ ’.

Method	20%	40%	60%	80%	100%
Point-MAE [47]	89.4	90.2	90.3	90.5	91.0
Point-M2AE [78]	<u>90.8</u>	<u>92.0</u>	<u>92.1</u>	<u>92.4</u>	<u>92.9</u>
I2P-MAE	92.1	92.7	93.1	93.1	93.4
<i>Improvement</i>	+1.3	+0.7	+1.0	+0.7	+0.5

Table 5. **Pre-training with Limited 3D Data.** We utilize randomly sampled 3D data from ShapeNet [6] for pre-training.

Pre-training with Limited 3D Data. In Table 5, we randomly sample the pre-training dataset, ShapeNet [6], with different ratios, and evaluate the performance of I2P-MAE when 3D data is further deficient. Aided by 2D pre-trained models, I2P-MAE still achieves competitive downstream accuracy in low-data regimes, especially for 20% and 60%, which outperforms Point-M2AE [78] by +1.3% and +1.0, respectively. *Importantly, with only 60% of the pre-training, I2P-MAE (93.1%) outperforms Point-MAE (91.0%) and Point-M2AE (92.9%) with full training data.* This indicates that our image-to-point learning scheme effectively can alleviate the need for large-scale 3D training datasets.

2D-guided Masking. In Table 6, we experiment different masking strategies for I2P-MAE. The first row represents 2D-guided masking, which preserves more semantically important tokens to be visible for the encoder. Compared to the second-row random masking, the guidance of 2D attention maps contributes to +0.4% and +0.9% accuracy on two datasets. Then, we reverse the token scores in the spatial attention cloud, and instead mask the most important tokens. As shown in the third row, the results are largely harmed, demonstrating the significance of encoding

2D-guided	Visible	Ratio	ModelNet40	OBJ-BG
✓	Important	0.8	93.4	87.1
-	Random	0.8	93.0	86.2
✓	Unimportant	0.8	92.8	84.5
✓	Important	0.7	92.5	83.8
✓	Important	0.9	92.8	85.6

Table 6. **2D-guided Masking.** ‘Visible’ denotes whether the preserved visible point tokens are more semantically important than the masked ones. ‘Ratio’ denotes the masking ratio.

Tokens		Targets		ModelNet40	OBJ-BG
Masked	Visible	3D	2D		
✓	✓	M	V	93.4	87.1
✓	-	M	-	92.9	84.7
-	✓	-	V	91.9	80.2
✓	-	-	M	91.2	77.8
✓	-	M	M	92.6	84.9

Table 7. **2D-semantic Reconstruction.** We utilize ‘M’ and ‘V’ to denote the reconstruction targets of masked and visible tokens. ‘3D’ and ‘2D’ denote 3D coordinates and 2D semantics.

critical 3D structures. Finally, we modify the masking ratio by ± 0.1 , which controls the proportion between visible and masked tokens. The performance decay indicates that, the 2D-semantic and 3D-coordinate reconstruction are required to be well balanced for a properly challenging pre-text task.

2D-semantic Reconstruction. In Table 7, we investigate which groups of tokens are the best for learning 2D-semantic targets. The comparison of the first two rows reveals the effectiveness of reconstructing 2D semantics from visible tokens, i.e., +0.5% and +2.4% classification accuracy. By only using 2D targets for either the visible or masked tokens (the 3rd and 4th rows), we verify that the 3D-coordinate reconstruction still plays an important role in I2P-MAE, which learns low-level geometric 3D patterns and provides complementary knowledge to the high-level 2D semantics. However, if the 3D and 2D targets are both reconstructed from the masked tokens (the last row), the network is restricted to learn 2D knowledge from the masked nonsignificant 3D geometries, other than the more discriminative parts. Also, assigning two targets on the same tokens might cause 2D-3D semantic conflicts. Therefore, the two targets are best to be reconstructed separately.

5. Visualization

We visualize the input point cloud, random masking, spatial attention cloud, 2D-guided masking, and the reconstructed 3D coordinates in Figure 6. Guided by the semantic scores from 2D pre-trained models (darker points indicate

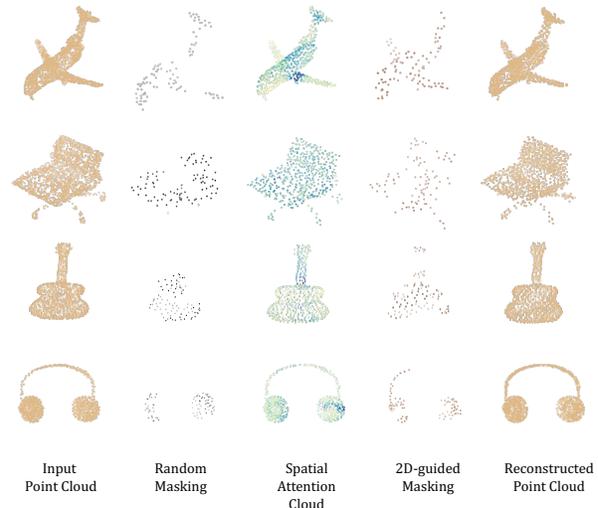


Figure 6. **Visualization of I2P-MAE.** Guided by the spatial attention cloud, I2P-MAE’s masking (the 4th column) preserves more semantically important 3D structures than random masking.

higher scores), the masked point cloud largely preserves significant parts of the 3D shape, e.g., the main body of an airplane, the grip of a guitar, the frame of a chair and headphones. By this, the 3D network can learn more discriminative features by reconstructing these visible 2D semantics.

6. Conclusion

In this paper, we propose I2P-MAE, a masked point modeling framework with effective image-to-point learning schemes. We introduce two approaches to transfer the well learned 2D knowledge into 3D domains: 2D-guided masking and 2D-semantic reconstruction. Aided by the 2D guidance, I2P-MAE learns superior 3D representations and alleviates the demand for large-scale 3D data. For future work, not limited to masking and reconstruction, we will explore more sufficient image-to-point learning approaches, e.g., point token sampling and 2D-3D class-token contrast. Also, we expect our pre-trained models to benefit wider ranges of 3D tasks, e.g., 3D object detection and visual grounding.

Acknowledgement. This project is funded in part by National Key R&D Program of China Project 2022ZD0161100, by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)’s InnoHK, by General Research Fund of Hong Kong RGC Project 14204021, by the National Natural Science Foundation of China (Grant No. 62206272), and by Shanghai Committee of Science and Technology (Grant No. 21DZ1100100). Hongsheng Li is a PI of CPII under the InnoHK.

References

- [1] 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. *arXiv preprint arXiv:2203.00680*, 2022. 3, 6
- [2] Anonymous. Rethinking network design and local geometry in point cloud: A simple residual MLP framework. In *Submitted to The Tenth International Conference on Learning Representations*, 2022. under review. 3, 6, 7
- [3] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatuo Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. *arXiv preprint arXiv:2202.03555*, 2022. 1
- [4] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 1
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. 1, 4
- [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. 1, 3, 6, 7
- [7] Anthony Chen, Kevin Zhang, Renrui Zhang, Zihan Wang, Yuheng Lu, Yandong Guo, and Shanghang Zhang. Pimae: Point cloud and image interactive masked autoencoders for 3d object detection. *arXiv preprint arXiv:2303.08129*, 2023. 3
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 1
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 1
- [10] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 1
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1
- [12] Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11162–11173, 2021. 1
- [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. 3
- [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. *arXiv preprint arXiv:2010.11929*, 2020. 2, 4, 6
- [15] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. 4
- [16] Kexue Fu, Peng Gao, ShaoLei Liu, Renrui Zhang, Yu Qiao, and Manning Wang. Pos-bert: Point cloud one-stage bert pre-training. *arXiv preprint arXiv:2204.00989*, 2022. 3
- [17] Kexue Fu, Peng Gao, Renrui Zhang, Hongsheng Li, Yu Qiao, and Manning Wang. Distillation with contrast is all you need for self-supervised point cloud representation learning. *arXiv preprint arXiv:2202.04241*, 2022. 3
- [18] Peng Gao, Teli Ma, Hongsheng Li, Jifeng Dai, and Yu Qiao. Convmae: Masked convolution meets masked autoencoders. *arXiv preprint arXiv:2205.03892*, 2022. 1
- [19] Peng Gao, Renrui Zhang, Rongyao Fang, Ziyi Lin, Hongyang Li, Hongsheng Li, and Qiao Yu. Mimic before reconstruct: Enhancing masked autoencoders with feature mimicking. *arXiv preprint arXiv:2303.05475*, 2023. 1
- [20] Ankit Goyal, Hei Law, Bowei Liu, Alejandro Newell, and Jia Deng. Revisiting point cloud shape classification with a simple and effective baseline. *arXiv preprint arXiv:2106.05304*, 2021. 4
- [21] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 1
- [22] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7(2):187–199, 2021. 3
- [23] Ziyu Guo, Xianzhi Li, and Pheng Ann Heng. Joint-mae: 2d-3d joint masked autoencoders for 3d point cloud pre-training. *arXiv preprint arXiv:2302.14007*, 2023. 3
- [24] Ziyu Guo, Renrui Zhang, Longtian Qiu, Xianzheng Ma, Xupeng Miao, Xuming He, and Bin Cui. Calip: Zero-shot enhancement of clip with parameter-free attention. *arXiv preprint arXiv:2209.14169*, 2022. 3
- [25] Abdullah Hamdi, Silvio Giancola, and Bernard Ghanem. Mvtn: Multi-view transformation network for 3d shape recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1–11, 2021. 6
- [26] Zhizhong Han, Mingyang Shang, Yu-Shen Liu, and Matthias Zwicker. View inter-prediction gan: Unsupervised representation learning for 3d shapes by learning global shape memories to support local view predictions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8376–8384, 2019. 6

- [27] 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. 1
- [28] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 1
- [29] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 1
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [31] Lang Huang, Shan You, Mingkai Zheng, Fei Wang, Chen Qian, and Toshihiko Yamasaki. Green hierarchical vision transformer for masked image modeling. *arXiv preprint arXiv:2205.13515*, 2022. 1
- [32] Siyuan Huang, Yichen Xie, Song-Chun Zhu, and Yixin Zhu. Spatio-temporal self-supervised representation learning for 3d point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6535–6545, 2021. 3
- [33] Jincen Jiang, Xuequan Lu, Lizhi Zhao, Richard Dazeley, and Meili Wang. Masked autoencoders in 3d point cloud representation learning. *arXiv preprint arXiv:2207.01545*, 2022. 1, 2, 3, 6
- [34] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [35] Jiaxin Li, Ben M Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9397–9406, 2018. 6
- [36] Ruihui Li, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-gan: a point cloud upsampling adversarial network. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7203–7212, 2019. 3
- [37] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *Advances in neural information processing systems*, 31:820–830, 2018. 3, 6, 7
- [38] Zhenyu Li, Zehui Chen, Ang Li, Liangji Fang, Qinrong Jiang, Xianming Liu, Junjun Jiang, Bolei Zhou, and Hang Zhao. Simipu: Simple 2d image and 3d point cloud unsupervised pre-training for spatial-aware visual representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1500–1508, 2022. 3
- [39] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 1
- [40] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1
- [41] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. *arXiv preprint arXiv:2203.11183*, 2022. 1, 3, 6, 7
- [42] Jihao Liu, Xin Huang, Yu Liu, and Hongsheng Li. Mixmim: Mixed and masked image modeling for efficient visual representation learning. *arXiv preprint arXiv:2205.13137*, 2022. 1
- [43] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8895–8904, 2019. 6, 7
- [44] Yueh-Cheng Liu, Yu-Kai Huang, Hung-Yueh Chiang, Hung-Ting Su, Zhe-Yu Liu, Chin-Tang Chen, Ching-Yu Tseng, and Winston H Hsu. Learning from 2d: Contrastive pixel-to-point knowledge transfer for 3d pretraining. *arXiv preprint arXiv:2104.04687*, 2021. 3
- [45] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 4
- [46] Zhengzhe Liu, Xiaojuan Qi, and Chi-Wing Fu. 3d-to-2d distillation for indoor scene parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4464–4474, 2021. 3
- [47] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. *arXiv preprint arXiv:2203.06604*, 2022. 1, 2, 3, 6, 7
- [48] Bui Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975. 4
- [49] Omid Poursaeed, Tianxing Jiang, Han Qiao, Nayun Xu, and Vladimir G Kim. Self-supervised learning of point clouds via orientation estimation. In *2020 International Conference on 3D Vision (3DV)*, pages 1018–1028. IEEE, 2020. 3
- [50] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 3, 6, 7
- [51] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. 3, 6, 7
- [52] Guocheng Qian, Xingdi Zhang, Abdullah Hamdi, and Bernard Ghanem. Pix4point: Image pretrained transformers for 3d point cloud understanding. *arXiv preprint arXiv:2208.12259*, 2022. 3
- [53] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1, 2, 3, 4, 6

- [54] Yongming Rao, Jiwen Lu, and Jie Zhou. Global-local bidirectional reasoning for unsupervised representation learning of 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5376–5385, 2020. 3
- [55] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 1
- [56] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 1
- [57] Jonathan Sauder and Bjarne Sievers. Self-supervised deep learning on point clouds by reconstructing space. *Advances in Neural Information Processing Systems*, 32, 2019. 3
- [58] Jonathan Sauder and Bjarne Sievers. Self-supervised deep learning on point clouds by reconstructing space. *Advances in Neural Information Processing Systems*, 32, 2019. 6
- [59] Corentin Sautier, Gilles Puy, Spyros Gidaris, Alexandre Boulch, Andrei Bursuc, and Renaud Marlet. Image-to-lidar self-supervised distillation for autonomous driving data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9891–9901, 2022. 3
- [60] Jong-Chyi Su, Matheus Gadelha, Rui Wang, and Subhransu Maji. A deeper look at 3d shape classifiers. In *The European Conference on Computer Vision (ECCV) Workshops*, 2018. 4
- [61] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016. 1
- [62] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. 1
- [63] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1588–1597, 2019. 2, 3, 6
- [64] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9782–9792, 2021. 1, 3, 6
- [65] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019. 3, 7
- [66] 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. *arXiv preprint arXiv:2208.02812*, 2022. 3, 4
- [67] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 2, 3, 6, 7
- [68] Tiange Xiang, Chaoyi Zhang, Yang Song, Jianhui Yu, and Weidong Cai. Walk in the cloud: Learning curves for point clouds shape analysis. *arXiv preprint arXiv:2105.01288*, 2021. 3
- [69] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *European conference on computer vision*, pages 574–591. Springer, 2020. 1, 3
- [70] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. *arXiv preprint arXiv:2111.09886*, 2021. 1
- [71] Chenfeng Xu, Shijia Yang, Bohan Zhai, Bichen Wu, Xiangyu Yue, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Image2point: 3d point-cloud understanding with pretrained 2d convnets. *arXiv preprint arXiv:2106.04180*, 2021. 3
- [72] Mutian Xu, Runyu Ding, Hengshuang Zhao, and Xiaojuan Qi. Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3173–3182, 2021. 3
- [73] Xu Yan, Jiantao Gao, Chaoda Zheng, Chao Zheng, Ruimao Zhang, Shuguang Cui, and Zhen Li. 2dprior: 2d priors assisted semantic segmentation on lidar point clouds. In *European Conference on Computer Vision*, pages 677–695. Springer, 2022. 3
- [74] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 206–215, 2018. 3, 6
- [75] 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. *ACM Transactions on Graphics (ToG)*, 35(6):1–12, 2016. 7
- [76] 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. *arXiv preprint arXiv:2111.14819*, 2021. 1, 2, 3, 6, 7
- [77] Haokui Zhang, Ying Li, Yen-an Jiang, Peng Wang, Qiang Shen, and Chunhua Shen. Hyperspectral classification based on lightweight 3-d-cnn with transfer learning. *IEEE Transactions on Geoscience and Remote Sensing*, 57(8):5813–5828, 2019. 3
- [78] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Pointm2ae: Multi-scale masked autoencoders for hierarchical point cloud pre-training. *arXiv preprint arXiv:2205.14401*, 2022. 1, 2, 3, 4, 6, 7
- [79] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. *arXiv preprint arXiv:2112.02413*, 2021. 3, 4

- [80] Renrui Zhang, Liuhui Wang, Ziyu Guo, and Jianbo Shi. Nearest neighbors meet deep neural networks for point cloud analysis. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1246–1255, 2023. 3
- [81] Renrui Zhang, Liuhui Wang, Yali Wang, Peng Gao, Hongsheng Li, and Jianbo Shi. Parameter is not all you need: Starting from non-parametric networks for 3d point cloud analysis. *arXiv preprint arXiv:2303.08134*, 2023. 3
- [82] Renrui Zhang, Ziyao Zeng, Ziyu Guo, Xinben Gao, Kexue Fu, and Jianbo Shi. Dspoint: Dual-scale point cloud recognition with high-frequency fusion. *arXiv preprint arXiv:2111.10332*, 2021. 3
- [83] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-supervised pretraining of 3d features on any point-cloud. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10252–10263, 2021. 1
- [84] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16259–16268, 2021. 7
- [85] Junsheng Zhou, Xin Wen, Yu-Shen Liu, Yi Fang, and Zhizhong Han. Self-supervised point cloud representation learning with occlusion auto-encoder. *arXiv preprint arXiv:2203.14084*, 2022. 6, 7
- [86] Xiangyang Zhu, Renrui Zhang, Bowei He, Ziyao Zeng, Shanghang Zhang, and Peng Gao. Pointclip v2: Adapting clip for powerful 3d open-world learning. *arXiv preprint arXiv:2211.11682*, 2022. 3