

Ponder: Point Cloud Pre-training via Neural Rendering

Di Huang^{1,2} Sida Peng³ Tong He^{2,†} Honghui Yang^{2,3} Xiaowei Zhou³ Wanli Ouyang²
 The University of Sydney¹ Shanghai AI Laboratory² Zhejiang University³

Abstract

We propose a novel approach to self-supervised learning of point cloud representations by differentiable neural rendering. Motivated by the fact that informative point cloud features should be able to encode rich geometry and appearance cues and render realistic images, we train a point-cloud encoder within a devised point-based neural renderer by comparing the rendered images with real images on massive RGB-D data. The learned point-cloud encoder can be easily integrated into various downstream tasks, including not only high-level tasks like 3D detection and segmentation but also low-level tasks like 3D reconstruction and image synthesis. Extensive experiments on various tasks demonstrate the superiority of our approach compared to existing pre-training methods.

1. Introduction

We have witnessed the widespread success of supervised learning in developing vision tasks, such as image classification [17, 10] and object detection [43, 16, 22]. In contrast to the 2D image domain, current 3D point cloud benchmarks only maintain limited annotations, in terms of quantity and diversity, due to the extremely high cost of laborious labeling. Self-supervised learning (SSL) for point cloud [52, 18, 23, 20, 6, 41, 62, 48, 59, 56, 36, 25, 61, 31], consequently, becomes one of the main driving forces and has attracted increasing attention in the 3D research community.

Previous SSL methods for learning effective 3D representation can be roughly categorized into two groups: contrast-based [52, 18, 23, 20, 6, 41, 62] and completion-based [48, 59, 56, 36, 25, 61, 31]. Contrast-based methods are designed to maintain invariant representation under different transformations. To achieve this, informative samples are required. In the 2D image domain, the above challenge is addressed by (1) introducing efficient positive/negative sampling methods, (2) using a large batch size

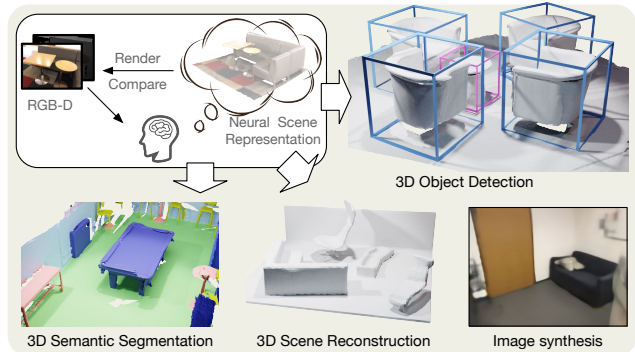


Figure 1. This work proposes a novel point cloud pre-training method via neural rendering, named **Ponder**. **Ponder** is directly trained with RGB-D image supervision, and can be used for various applications, e.g. 3D object detection, 3D semantic segmentation, 3D scene reconstruction, and image synthesis.

and storing representative samples, and (3) applying various data augmentation policies. Inspired by these works, many works [52, 18, 23, 20, 6, 41, 62] are proposed to learn geometry-invariant features on 3D point cloud.

Completion-based methods are another line of research for 3D SSL, which utilizes a pre-training task of reconstructing the masked point cloud based on partial observations. By maintaining a high masking ratio, such a simple task encourages the model to learn a holistic understanding of the input beyond low-level statistics. Although the masked autoencoders have been successfully applied for SSL in images [14] and videos [11, 47], it remains challenging and still in exploration due to the inherent irregularity and sparsity of the point cloud data.

Different from the two groups of methods above, we propose **point cloud pre-training via neural rendering (Ponder)**. Our motivation is that neural rendering, one of the most amazing progress and domain-specific design in 3D vision, can be leveraged to enforce the point cloud features being able to encode rich geometry and appearance cues. As illustrated in Figure 1, we address the task of learning representative 3D features via point cloud rendering. To the best of our knowledge, this is the first exploration of neural rendering for pre-training 3D point cloud models. Specif-

[†]denote corresponding author.

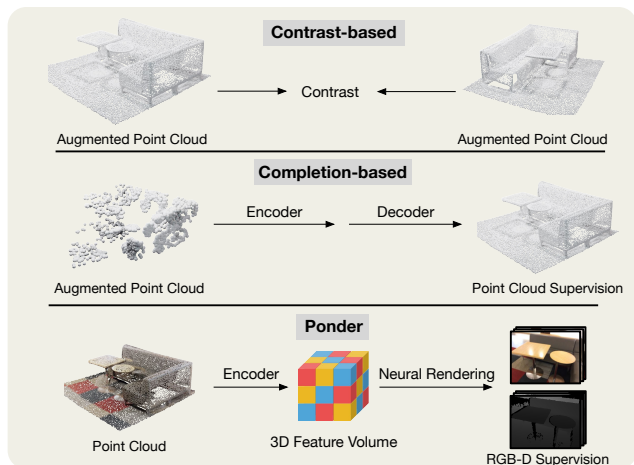


Figure 2. Different types of point cloud pre-training.

ically, given one or a sequence of RGB-D images, we lift them to 3D space and obtain a set of colored points. Points are then forwarded to a 3D encoder to learn the geometry and appearance of the scene via a neural representation. Provided specific parameters of the camera and the neural representation from the encoder, neural rendering is leveraged to render the RGB and depth images in a differentiable way. The network is trained to minimize the difference between rendered and observed 2D images. In doing so, our approach enjoys multiple advantages:

- Our method is able to learn effective point cloud representation, which encodes rich geometry and appearance clues by leveraging neural rendering.
- Our method can be flexibly integrated into various tasks. For the first time, we validate the effectiveness of the proposed pre-training method for low-level tasks like surface reconstruction and image synthesis tasks.
- The proposed method can leverage rich RGB-D images for pre-training. The easier accessibility of the RGB-D data enables the possibility of 3D pre-training on a large amount of data.

Our approach proposes a novel pretext task that can serve as a strong alternative to contrast-based methods and completion-based methods in 3D point cloud pre-training. The proposed framework, **Ponder**, is capable of accommodating a variety of point cloud backbones, both point-based and voxel-based, and has been rigorously evaluated on a range of challenging 3D tasks, including object detection, semantic segmentation, reconstruction, and image synthesis. The consistent improvements demonstrate the effectiveness of our proposed **Ponder**.

2. Related Work

Neural rendering. Neural Rendering is a type of rendering technology that uses neural networks to differentially render images from 3D scene representation. NeRF[30] is one of the representative neural rendering methods, which represents the scene as the neural radiance field and renders the images via volume rendering. Based on NeRF, there are a series of works [33, 57, 50, 34, 58, 51, 60, 42, 3, 53] trying to improve the NeRF representation, including accelerate NeRF training, boost the quality of geometry, and so on. Another type of neural rendering leverages neural point clouds as the scene representation. [2, 40] take points locations and corresponding descriptors as input, rasterize the points with z-buffer, and use a rendering network to get the final image. Later work of PointNeRF[54] renders realistic images from neural point cloud representation using a NeRF-like rendering process. Our work is inspired by the recent progress of neural rendering.

Self-supervised learning in point clouds. Current methods can be roughly categorized into two categories: contrast-based and completion-based. Inspired by the works [15, 5] from the 2D image domain, PointContrast [52] is one of the pioneering works for 3D contrastive learning. Similarly, it encourages the network to learn invariant 3D representation under different transformations. Some works [18, 23, 20, 6, 41, 62] follow the pipeline by either devising new sampling strategies to select informative positive/negative training pairs, or explore various types of data augmentations. Another line of work is completion-based [59, 56, 36, 25, 61, 31] methods, which get inspiration from Masked Autoencoders [14]. PointMAE [36] proposes restoring the masked points via a set-to-set Chamfer Distance. VoxelMAE [31] instead recovers the underlying geometry by distinguishing if the voxel contains points. Another work MaskPoint[25] pre-train point cloud encoder by performing binary classification to check if a sampled point is occupied. Later, IAE [56] proposes to pre-train point cloud encoder by recovering continuous 3D geometry in an implicit manner. Different from the above pipelines, we propose a novel framework for point cloud pre-training via neural rendering.

Multi-modal point cloud pre-training. Some recent works explore the pre-training pipeline with multi-modality data of 2D images and 3D point clouds. Pri3D[19] use 3D point cloud and multi-view images to pre-train the 2D image networks. CrossPoint[1] aligns the 2D image features and 3D point cloud features through a contrastive learning pipeline. [24] proposes a unified framework for exploring the invariances with different input data formats, including 2D images and 3D point clouds. Different from previous methods, most of which attempt to align 2D images and 3D point clouds in the feature space, our method proposes to

connect 2D and 3D in the RGB-D image domain via differentiable rendering.

3. Methods

An overview of our **Ponder** is presented in Figure 3. Provided the camera pose, 3D point clouds are obtained by projecting the RGB-D images back to 3D space (Section 3.1). Then, we extract point-wise feature using a point cloud encoder (Section 3.2) and organize it to a 3D feature volume (Section 3.3), which is used to reconstruct the neural scene representation and render images in a differentiable manner (Section 3.4).

3.1. Constructing point cloud from RGB-D images

The proposed method makes use of sequential RGB-D images $\{(I_i, D_i)\}_{i=1}^N$, the camera intrinsic parameters $\{\mathbf{K}_i\}_{i=1}^N$, and extrinsic poses $\{\xi_i\}_{i=1}^N \in \mathbf{SE}(3)$. N is the input view number. $\mathbf{SE}(3)$ refers to the Special Euclidean Group representing 3D rotations and translations. The camera parameters can be easily obtained from SfM or SLAM.

We construct the point cloud \mathcal{X} by back-projecting RGB-D images to point clouds in a unified coordinate:

$$\mathcal{X} = \bigcup_i^N \pi^{-1}(I_i, D_i, \xi_i, \mathbf{K}_i), \quad (1)$$

where π^{-1} back-projects the RGB-D image to 3D world space using camera poses. Note that different from previous methods which only consider the point location, our method attributes each point with both point location and RGB color. The details of π^{-1} are provided in the supplementary material.

3.2. Point cloud encoder for feature extraction

Given the point cloud \mathcal{X} constructed from RGB-D images, a point cloud encoder f_p is used to extract per-point feature embedding \mathcal{E} :

$$\mathcal{E} = f_p(\mathcal{X}). \quad (2)$$

The encoder f_p pre-trained with the method mentioned in the Section 3.4 serves as a good initialization for various downstream tasks.

3.3. Building feature volume

After completing feature extraction, we use average pooling to convert the point embeddings \mathcal{E} into a 3D feature volume. We then employ a U-Net style 3D CNN to fill in the empty space and aggregate features from the surrounding points to obtain a dense 3D volume, denoted as \mathcal{V} .

3.4. Pre-training with Neural Rendering

This section introduces how to reconstruct the implicit scene representation and render images differentially. We first give a brief introduction to neural scene representation, then illustrate how to integrate it into our point cloud pre-training pipeline. Last, we show the differentiable rendering formulation to render color and depth images from the neural scene representation.

Brief introduction of neural scene representation.

Neural scene representation aims to represent the scene geometry and appearance through a neural network. In this paper, we use the Signed Distance Function (SDF), which measures the distance between a query point and the surface boundary, to represent the scene geometry implicitly. SDF is capable of representing high-quality geometry details. For any query point of the scene, the neural network takes points features as input and outputs the corresponding SDF value and RGB value. In this way, the neural network captures both the geometry and appearance information of a specific scene. Following NeuS[50], the scene can be reconstructed as:

$$s(\mathbf{p}) = \tilde{f}_s(\mathbf{p}), \quad c(\mathbf{p}, \mathbf{d}) = \tilde{f}_c(\mathbf{p}, \mathbf{d}), \quad (3)$$

where \tilde{f}_s is the SDF decoder and \tilde{f}_c is the RGB color decoder. \tilde{f}_s takes point location \mathbf{p} as input, and predicts the SDF value s . \tilde{f}_c takes point location \mathbf{p} and viewing direction \mathbf{d} as input, and outputs the RGB color value c . Both \tilde{f}_s and \tilde{f}_c are implemented by simple MLP networks.

Neural scene representation from point cloud input in Ponder.

To predict a neural scene representation from the input point cloud, we change the scene formulation to take 3D feature volume \mathcal{V} as an additional input. Specifically, given a 3D query point \mathbf{p} and viewing direction \mathbf{d} , the feature embedding $\mathcal{V}(\mathbf{p})$ can be extracted from the processed feature volume \mathcal{V} by trilinear interpolation. The scene is then represented as:

$$s(\mathbf{p}) = f_s(\mathbf{p}, \mathcal{V}(\mathbf{p})), \quad c(\mathbf{p}, \mathbf{d}) = f_c(\mathbf{p}, \mathbf{d}, \mathcal{V}(\mathbf{p})), \quad (4)$$

where \mathcal{V} is predicted by the point cloud encoder f_p and encodes information of each scene. f_s and f_c are SDF and RGB decoders shared for all scenes. Different from Equation (3), which is used for storing single-scene information in the $\{\tilde{f}_s, \tilde{f}_c\}$, the formulation in Equation (4) includes an extra input $\mathcal{V}(\mathbf{p})$ to facilitate representing the information of multiple scenes.

Differentiable rendering. Given the dense 3D volume \mathcal{V} and viewing point, we use differentiable volume rendering to render the projected color images and depth images.

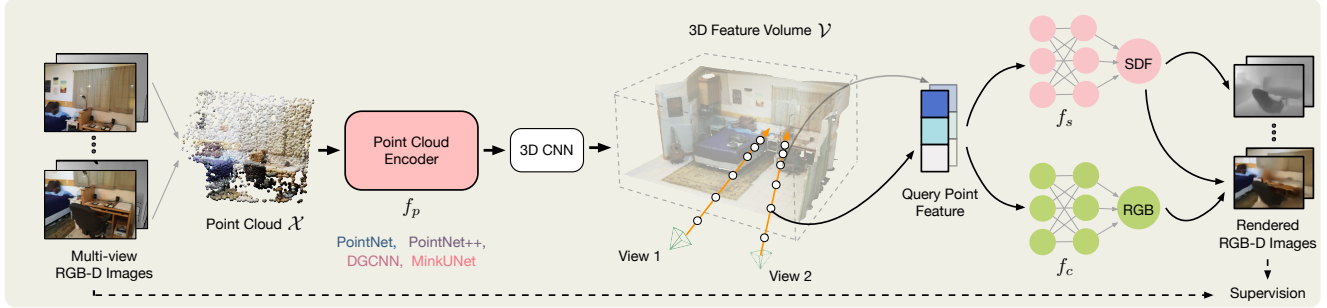


Figure 3. **The pipeline of our point cloud pre-training via neural rendering (Ponder).** Given multi-view RGB-D images, we first construct the point cloud by back-projection, then use a point cloud encoder f_p to extract per-point features \mathcal{E} . \mathcal{E} are organized to a 3D feature volume by average pooling and then processed by the 3D convolution layer. Finally, the 3D feature volume is rendered to multi-view RGB-D images via a differentiable neural rendering, which are compared with the input multi-view RGB-D images as the supervision.

For each rendering ray with camera origin \mathbf{o} and viewing direction \mathbf{d} , we sample a set of ray points $\{\mathbf{p}(z) | \mathbf{p}(z) = \mathbf{o} + z\mathbf{d}, z \in [z_n, z_f]\}$ along the ray, where z denotes the length of the ray. Note that \mathbf{o} and \mathbf{d} can be calculated from paired camera parameters $\{(K_i, \xi_i)\}$. z_n and z_f denote the near and far bounds of the ray. Different from previous methods [30, 50], we automatically determine $\{z_n, z_f\}$ by the ray intersection with the 3D feature volume box, using axis-aligned bounding boxes (AABB) algorithm. Then, the ray color and depth value can be aggregated as:

$$\hat{C} = \int_{z_n}^{z_f} \omega(z) c(\mathbf{p}(z), \mathbf{d}) dz, \quad (5)$$

$$\hat{D} = \int_{z_n}^{z_f} \omega(z) z dz, \quad (6)$$

where the \hat{C} is the ray color and the \hat{D} is the ray depth. We follow NeuS[50] to build an unbiased and occlusion-awareness weight function $w(z)$:

$$w(z) = T(z) \cdot \rho(z). \quad (7)$$

$T(z)$ measures the accumulated transmittance from z_n to z and $\rho(z)$ is the occupied density function which are defined as:

$$T(z) = \exp\left(-\int_{z_n}^z \rho(z) dz\right), \quad (8)$$

$$\rho(z) = \max\left(\frac{-\frac{d\Phi_h}{dz}(s(\mathbf{p}(z)))}{\Phi_h(s(\mathbf{p}(z)))}, 0\right). \quad (9)$$

$\Phi_h(x)$ is the Sigmoid function $\Phi_h(x) = (1 + e^{-hx})^{-1}$ where h^{-1} is treated as a trainable parameter, h^{-1} approaches to zero as the network training converges. In practice, we use a numerically approximated version by quadrature. We make the decode networks $\{f_s, f_c\}$ relatively smaller than [30, 50] to accelerate the training process.

Rendered examples. The rendered color images and depth images are shown in Figure 4. As shown in the figure, even though the input point cloud is pretty sparse, our

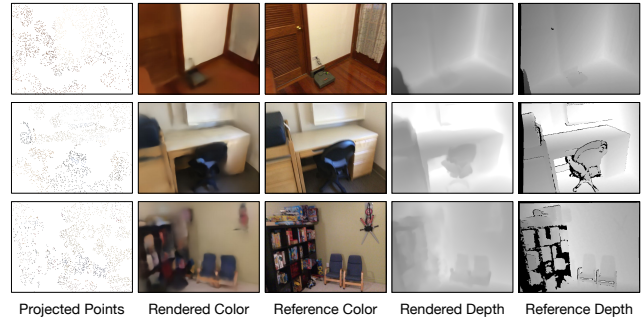


Figure 4. **Rendered images by Ponder** on the ScanNet validation set. The projected point clouds are visualized in the first column. Even though input point clouds are very sparse, our model is still capable of rendering color and depth images similar to the reference images.

method is still capable of rendering color and depth images similar to the reference image.

3.5. Pre-training loss

We leverage the input $\{I_i, D_i\}$ to supervise neural scene representation reconstruction. The total loss function contains five parts,

$$L = \lambda_c L_c + \lambda_d L_d + \lambda_e L_e + \lambda_s L_s + \lambda_f L_f, \quad (10)$$

which are loss functions responsible for color supervision L_c , depth supervision L_d , Eikonal regularization L_e , near-surface SDF supervision L_s , and free space SDF supervision L_f . These loss functions are illustrated in the following section.

Color and depth loss. L_c and L_d are the color loss and depth loss, which measure consistency between the rendered pixels and the ground-truth pixels. Assume that we sample N_r rays for each image and N_p points for each ray,

then the L_c and L_d can be written as:

$$L_c = \frac{1}{N_r} \sum_i^{N_r} \|\hat{C} - C\|_2^2 \quad (11)$$

$$L_d = \frac{1}{N_r} \sum_i^{N_r} \|\hat{D} - D\|_2^2, \quad (12)$$

where C and D are the ground-truth color and depth respectively for each ray, \hat{C} and \hat{D} are their corresponding rendered ones in Eq. (5) and Eq. (6).

Loss for SDF regularization. L_e is the widely used Eikonal loss [13] for SDF regularization:

$$L_e = \frac{1}{N_r N_p} \sum_{i,j}^{N_r, N_p} (|\nabla s(\mathbf{p}_{i,j})| - 1)^2, \quad (13)$$

where $\nabla s(\mathbf{p}_{i,j})$ denotes the gradient of SDF s at location $\mathbf{p}_{i,j}$. Since SDF is a distance measure, L_e encourages this distance to have a unit norm gradient at the query point.

Near-surface and free space loss for SDF. To stabilize the training and improve the reconstruction performance, similar to iSDF [35] and GO-Surf [49], we add additional approximate SDF supervision to help the SDF estimation. Specifically, for near-surface points, the difference between rendered depth and ground-truth depth can be viewed as the pseudo-SDF ground-truth supervision; for points far from the surface, a free space loss is used to regularize the irregular SDF value additionally. To calculate the approximate SDF supervision, we first define an indicator $b(z)$ for each sampled ray point with ray length z and corresponding GT depth D :

$$b(z) = D - z. \quad (14)$$

$b(z)$ can be viewed as the approximate SDF value, which is credible only when $b(z)$ is small. Let t be a human-defined threshold, which is set as 0.05 in this paper. For sampled ray points that satisfy $b(z) \leq t$, we leverage the near-surface SDF loss to constrain the SDF prediction $s(z_{i,j})$:

$$L_s = \frac{1}{N_r N_p} \sum_{i,j}^{N_r, N_p} |s(z_{i,j}) - b(z_{i,j})|. \quad (15)$$

For the remaining sampled ray points, we use a free space loss:

$$L_f = \frac{1}{N_r N_p} \sum_{i,j}^{N_r, N_p} \max(0, e^{-\alpha \cdot s(z_{i,j})} - 1, s(z_{i,j}) - b(z_{i,j})), \quad (16)$$

where α is set as 5 following the same with [35, 49]. Note that due to the noisy depth images, we only apply L_s and L_f on the rays that have valid depth values.

In our experiments, we follow a similar loss of weight with GO-Surf [49], which sets λ_c as 10.0, λ_d as 1.0, λ_s as 10.0, and λ_f as 1.0. We observe that the Eikonal term in our method can easily lead to over-smooth reconstructions, thus we use a small weight of 0.01 for the Eikonal loss.

4. Experiments

4.1. Pre-training

Datasets. We use ScanNet[9] RGB-D images as our pre-training data. ScanNet is a widely used real-world indoor dataset, which contains more than 1500 indoor scenes. Each scene is carefully scanned by an RGB-D camera, leading to about 2.5 million RGB-D frames in total. We follow the same train/val split with VoteNet[38].

Data preparation. During pre-training, a mini-batch of batch size 8 includes point clouds from 8 scenes. The point cloud of a scene, serving as the input of the point cloud encoder in our approach, is back-projected from the 5 RGB-D frames of the video for the scene with an interval of 20. The 5 frames are also used as the supervision of the network. We randomly down-sample the input point cloud to 20,000 points and follow the masking strategy as used in Mask Point [25].

Implementation details. We train the proposed pipeline for 100 epochs using an AdamW optimizer [29] with a weight decay of 0.05. The learning rate is initialized as $1e-4$ with Exponential scheduling. For the rendering process, we randomly choose 128 rays for each image and sample 128 points for each ray. More implementation details can be found in the supplementary materials.

4.2. Transfer Learning

In contrast to previous methods, our approach is able to encode rich geometry and appearance cues into the point cloud representations via neural rendering. These strengths make it flexible to be applied to various tasks, including not only 3D semantic segmentation and 3D detection tasks but also low-level surface reconstruction and image synthesis.

4.2.1 High-level 3D Tasks

3D object detection. We select two representative approaches, VoteNet [38] and H3DNet [63], as the baselines. VoteNet leverages a voting mechanism to obtain object centers, which are used for generating 3D bounding box proposals. By introducing a set of geometric primitives, H3DNet achieves a significant improvement in accuracy compared to previous methods. Two datasets are applied to verify the effectiveness of our method: ScanNet[9] and SUN RGB-D[44]. Different from ScanNet, which contains fully reconstructed 3D scenes, SUN RGB-D is a single-

Method	Detection Model	Pre-training Type	Pre-training Epochs	ScanNet		SUN RGB-D	
				AP ₅₀ ↑	AP ₂₅ ↑	AP ₅₀ ↑	AP ₂₅ ↑
3DETR[32]	3DETR	-	-	37.5	62.7	30.3	58.0
Point-BERT[59]	3DETR	Completion	300	38.3	61.0	-	-
MaskPoint[25]	3DETR	Completion	300	40.6	63.4	-	-
VoteNet [38]	VoteNet	-	-	33.5	58.6	32.9	57.7
STRL[20]	VoteNet	Contrast	100	38.4	59.5	35.0	58.2
RandomRooms[41]	VoteNet	Contrast	300	36.2	61.3	35.4	59.2
PointContrast[52]	VoteNet	Contrast	-	38.0	59.2	34.8	57.5
PC-FractalDB[55]	VoteNet	Contrast	-	38.3	61.9	33.9	59.4
DepthContrast[62]	VoteNet	Contrast	1000	39.1	62.1	35.4	60.4
IAE[56]	VoteNet	Completion	1000	39.8	61.5	36.0	60.4
Ponder	VoteNet	Rendering	100	41.0 (+7.5)	63.6 (+5.0)	36.6 (+3.7)	61.0 (+3.3)

Table 1. **3D object detection** AP₂₅ and AP₅₀ on ScanNet and SUN RGB-D. VoteNet[38] is a baseline model. Purple numbers indicate improvements over the corresponding baseline. The DepthContrast[62] and Point-BERT[59] results are adopted from IAE[56] and MaskPoint[25]. **Ponder** outperforms both VoteNet-based and 3DETR-based point cloud pre-training methods with fewer training epochs.

Method	AP ₅₀ ↑	AP ₂₅ ↑
VoteNet[38]	33.5	58.7
3DETR[32]	37.5	62.7
3DETR-m[32]	47.0	65.0
H3DNet[63]	48.1	67.2
Ponder+H3DNet	50.9 (+2.8)	68.4 (+1.2)

Table 2. **3D object detection**. AP₂₅ and AP₅₀ on ScanNet. **Ponder** significantly boosts the detection accuracy of H3DNet by a margin of +2.8 and +1.2 for AP₅₀ and AP₂₅, respectively.

view RGB-D dataset with 3D bounding box annotations. It has 10,335 RGB-D images for 37 object categories. For pre-training, we use PointNet++ as the point cloud encoder f_p , which is identical to the backbone used in VoteNet and H3DNet. We pre-train the point cloud encoder on the ScanNet dataset and transfer the weight as the downstream initialization. Following [38], we use average precision with 3D detection IoU threshold 0.25 and threshold 0.5 as the evaluation metrics.

The 3D detection results are shown in Table 1. Our method improves the baseline of VoteNet without pre-training by a large margin, boosting AP₅₀ by 7.5% and 3.7% for ScanNet and SUN RGB-D, respectively. IAE [56] is a pre-training method that represents the inherent 3D geometry in a continuous manner. Our learned point cloud representation achieves higher accuracy because it is able to recover both the geometry and appearance of the scene. The AP₅₀ and AP₂₅ of our method are higher than that of IAE by 1.2% and 2.1% on ScanNet, respectively. Besides, we have observed that our method surpasses the recent point cloud pre-training approach, MaskPoint [25], even when using a less sophisticated backbone (PointNet++ vs. 3DETR), as presented in Table 1. To verify the effectiveness of **Pon-**

der, we also apply it for a much stronger baseline, H3DNet. As shown in Table 2, our method surpasses H3DNet by +2.8 and +1.2 for AP₅₀ and AP₂₅, respectively.

3D semantic segmentation. 3D semantic segmentation is another fundamental scene understanding task. We select one of the top-performing backbones, MinkUNet[8], for transfer learning. MinkUNet leverage 3D sparse convolution to extract effective 3D scene features. For pre-training, we use MinkUNet as the point cloud encoder f_p , and pre-train the model on ScanNet. We report the finetuning results on the ScanNet dataset with the mean IoU of the validation set as the evaluation metric. Table 3 shows the quantitative results of **Ponder** with MinkUNet. The results demonstrate that **Ponder** is effective in improving the semantic segmentation performance, achieving a significant improvement of 1.3 mIoU.

4.2.2 Low-level 3D Tasks

Low-level 3D tasks like scene reconstruction and image synthesis are getting increasing attention due to their wide applications. However, most of them are trained from scratch. How to pre-train a model with a good initialization is desperately needed. We are the first pre-training work to demonstrate a strong transferring ability to such low-level 3D tasks.

3D scene reconstruction. 3D scene reconstruction task aims to recover the scene geometry, e.g. mesh, from the point cloud input. We choose ConvONet[37] as the baseline model, whose architecture is widely adopted in [7, 26, 57]. Following the same setting as ConvONet, we conduct experiments on the Synthetic Indoor Scene Dataset

Method	mIoU \uparrow
PointNet++[39]	53.5
KPConv[46]	69.2
SparseConvNet[12]	69.3
PT[64]	70.6
MinkUNet[8]	72.2
Ponder+MinkUNet	73.5 (+1.3)

Table 3. **3D segmentation mIoU** on ScanNet dataset.

Method	Encoder	IoU \uparrow	NC \uparrow	F-Score \uparrow
ConvONet[37]	PointNet++	77.8	88.7	90.6
IAE[56]	PointNet++	75.7	88.7	91.0
Ponder	PointNet++	80.2 (+2.4)	89.3	92.0

Table 4. **3D scene reconstruction IoU, NC, and F-Score** on SISD dataset with PointNet++ model. **Ponder** is able to boost the reconstruction performance.

(SISD)[37], which is a synthetic dataset and contains 5000 scenes with multiple ShapeNet [4] objects. To make a fair comparison with IAE [56], we use the same VoteNet-style PointNet++ as the encoder of ConvONet, which down-samples the original point cloud to 1024 points. Following [37], we use Volumetric IoU, Normal Consistency (NC), and F-Score [45] with the threshold value of 1% as the evaluation metrics.

The results are shown in Table 4. Compared to the baseline ConvONet model with PointNet++, IAE is not able to boost the reconstruction results, while the proposed approach can improve the reconstruction quality (+2.4% for IoU). The results show the effectiveness of **Ponder** for the 3D reconstruction task.

Image synthesis from point clouds. We also validate the effectiveness of our method on another low-level task of image synthesis from point clouds. We use Point-NeRF[54] as the baseline. Point-NeRF uses neural 3D point clouds with associated neural features to render images. It can be used both for a generalizable setting for various scenes and a single-scene fitting setting. In our experiments, we mainly focus on the generalizable setting of Point-NeRF. We replace the 2D image features of Point-NeRF with point features extracted by a DGCNN network. Following the same setting with PointNeRF, we use DTU[21] as the evaluation dataset. DTU dataset is a multiple-view stereo dataset containing 80 scenes with paired images and camera poses. We transfer both the DGCNN encoder and color decoder as the weight initialization of Point-NeRF. We use PSNR as the metric for synthesized image quality evaluation.

The results are shown in Figure 5. By leveraging the pre-trained weights of our method, the image synthesis model is able to converge faster with fewer training steps and achieve better final image quality than training from scratch.

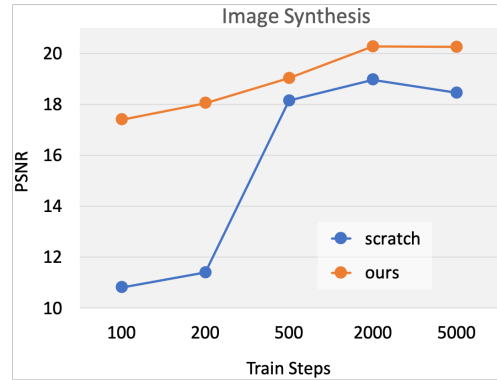


Figure 5. **Comparison of image synthesis from point clouds.** Compared with training from scratch, our **Ponder** model is able to converge faster and achieve better image synthesis results.

4.3. Ablation study

In this section, we conduct a series of ablation experiments to evaluate the effectiveness of our proposed approach. All experiments are conducted on ScanNet and SUN RGB-D datasets. We use 3D object detection for evaluation due to its simplicity.

Influence of Rendering Targets. The rendering part of our method contains two items: RGB color image and depth image. We study the influence of each item with the transferring task of 3D detection. The results are presented in Table 5. Combining depth and color images for reconstruction shows the best detection results. In addition, using depth reconstruction presents better performance than color reconstruction for 3D detection.

Influence of mask ratio. To augment point cloud data, we employ random masking as one of the augmentation methods, which divides the input point cloud into 2048 groups with 64 points. In this ablation study, we evaluate the performance of our method with different mask ratios, ranging from 0% to 90%, on the ScanNet and SUN RGB-D datasets, and report the results in Table 6. Notably, we find that even when no dividing and masking strategy is applied (0%), our method achieves a competitive AP_{50} performance of 40.7 and 37.3 on ScanNet and SUN RGB-D, respectively. Our method achieves the best performance on ScanNet with a mask ratio of 75% and a AP_{50} performance of 41.7. Overall, these results suggest that our method is robust to the hyper-parameter of mask ratio and can still achieve competitive performance without any mask operation.

Influence of 3D feature volume resolution. In our method, **Ponder** constructs a 3D feature volume with a resolution of [16, 32, 64], which is inspired by recent progress in multi-resolution 3D reconstruction. However, building such a high-resolution feature volume can consume significant GPU memory. To investigate the effect of feature volume resolution, we conduct experiments with different resolutions and report the results in Table 7. From the results, we

observe that even with a smaller feature volume resolution of 16, **Ponder** can still achieve competitive performance on downstream tasks.

Supervision	ScanNet	SUN RGB-D
VoteNet	33.5	32.9
+Depth	40.9 (+7.4)	36.1 (+3.2)
+Color	40.5 (+7.0)	35.8 (+2.9)
+Depth+Color	41.0 (+7.5)	36.6 (+3.7)

Table 5. **Ablation study for supervision type.** 3D detection AP_{50} on ScanNet and SUN RGB-D. Combining color supervision and depth supervision can lead to better detection performance than using a single type of supervision.

Mask ratio	ScanNet	SUN RGB-D
VoteNet	33.5	32.9
0%	40.7 (+7.2)	37.3 (+4.4)
25%	40.7 (+7.2)	36.2 (+3.3)
50%	40.3 (+6.8)	36.9 (+4.0)
75%	41.7 (+8.2)	37.0 (+4.1)
90%	41.0 (+7.5)	36.6 (+3.7)

Table 6. **Ablation study for mask ratio.** 3D detection AP_{50} on ScanNet and SUN RGB-D.

Resolution	ScanNet	SUN RGB-D
VoteNet	33.5	32.9
16	40.7 (+7.2)	36.6 (+3.7)
16+32+64	41.0 (+7.5)	36.6 (+3.7)

Table 7. **Ablation study for feature volume resolution.** 3D detection AP_{50} on ScanNet and SUN RGB-D.

View number	ScanNet	SUN RGB-D
VoteNet	33.5	32.9
1 view	40.1 (+6.6)	35.4 (+2.5)
3 views	40.8 (+7.3)	36.0 (+3.1)
5 views	41.0 (+7.5)	36.6 (+3.7)

Table 8. **Ablation study for view number.** 3D detection AP_{50} on ScanNet and SUN RGB-D. Using multi-view supervision for point cloud pre-training can achieve better performance.

Number of input RGB-D view. Our method utilizes N RGB-D images, where N is the input view number. We study the influence of N and conduct experiments on 3D detection, as shown in Table 8. We change the number of input views while keeping the scene number of a batch still 8. Using multi-view supervision helps to reduce single-view ambiguity. Similar observations are also found in the multi-view reconstruction task [27]. Compared with the

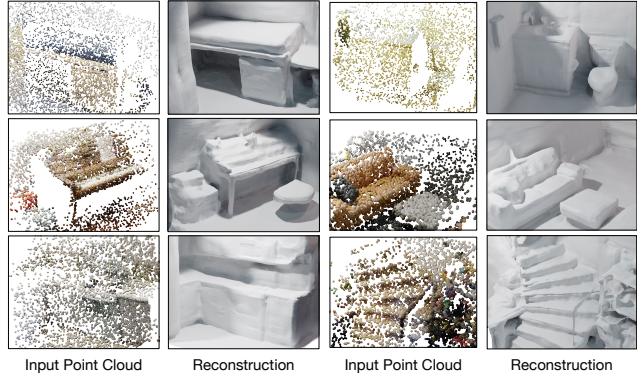


Figure 6. Reconstructed surface by **Ponder**. Our pre-training method can be easily integrated into the task of 3D reconstruction. Despite the sparsity of the input point cloud (only 2% points are used), our method can still recover precise geometric details.

single view, multiple views achieve higher accuracy, boosting AP_{50} by 0.9% and 1.2% for ScanNet and SUN RGB-D datasets, respectively.

4.4. Other applications

The pre-trained model from our pipeline **Ponder** itself can also be directly used for surface reconstruction from sparse point clouds. Specifically, after learning the neural scene representation, we query the SDF value in the 3D space and leverage the Marching Cubes [28] to extract the surface. We show the reconstruction results in Figure 6. The results show that even though the input is sparse point clouds from complex scenes, our method is able to recover high-fidelity meshes. Check the supplementary for more image synthesis and 3D reconstruction results.

5. Conclusion

This paper shows that differentiable neural rendering is a powerful tool for point cloud representation learning. The proposed pre-training pipeline, **Ponder**, is able to encode rich geometry and appearance cues into the point cloud representation via neural rendering. For the first time, our model can be transferred to both high-level 3D perception tasks and 3D low-level tasks, like 3D reconstruction and image synthesis from point clouds. Besides, the learned **Ponder** model can be directly used for 3D reconstruction and image synthesis from sparse point clouds. We also exploratively validate the effectiveness of **Ponder** on **out-door scenario** and other input modalities, where we observe 1.41% mAP improvement for 3D multiview object detection on the NuScene dataset that takes multiview images as input (details in the supplementary material).

Several directions could be explored in future works. First, recent progress in neural representations could help **Ponder** achieve better rendering quality and gain more ac-

curate supervision from 2D images. Second, thanks to the flexible architecture design, **Ponder** can potentially be expanded to other self-supervised learning fields, e.g., pre-training 2D image backbones, and other downstream tasks.

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. In *CVPR*, 2022.
- [2] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *ECCV*. Springer, 2020.
- [3] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, 2021.
- [4] 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.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*. PMLR, 2020.
- [6] Yujin Chen, Matthias Nießner, and Angela Dai. 4dcontrast: Contrastive learning with dynamic correspondences for 3d scene understanding. In *ECCV*, 2022.
- [7] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *CVPR*, 2020.
- [8] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, 2019.
- [9] 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.
- [10] 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.
- [11] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *arXiv preprint arXiv:2205.09113*, 2022.
- [12] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. *CVPR*, 2018.
- [13] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.
- [14] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022.
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [18] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *CVPR*, 2021.
- [19] Ji Hou, Saining Xie, Benjamin Graham, Angela Dai, and Matthias Nießner. Pri3d: Can 3d priors help 2d representation learning? In *ICCV*, 2021.
- [20] Siyuan Huang, Yichen Xie, Song-Chun Zhu, and Yixin Zhu. Spatio-temporal self-supervised representation learning for 3d point clouds. *arXiv preprint arXiv:2109.00179*, 2021.
- [21] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, 2014.
- [22] Ge-Peng Ji, Deng-Ping Fan, Yu-Cheng Chou, Dengxin Dai, Alexander Liniger, and Luc Van Gool. Deep gradient learning for efficient camouflaged object detection. *Machine Intelligence Research*, 20(1):92–108, 2023.
- [23] Li Jiang, Shaoshuai Shi, Zhuotao Tian, Xin Lai, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In *ICCV*, 2021.
- [24] Lanxiao Li and Michael Heizmann. A closer look at invariances in self-supervised pre-training for 3d vision. *arXiv preprint arXiv:2207.04997*, 2022.
- [25] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. *arXiv preprint arXiv:2203.11183*, 2022.
- [26] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 33, 2020.
- [27] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. *arXiv preprint arXiv:2206.05737*, 2022.
- [28] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4), 1987.
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [30] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1), 2021.
- [31] Chen Min, Dawei Zhao, Liang Xiao, Yiming Nie, and Bin Dai. Voxel-mae: Masked autoencoders for pre-training large-scale point clouds. *arXiv preprint arXiv:2206.09900*, 2022.
- [32] Ishan Misra, Rohit Girdhar, and Armand Joulin. An End-to-End Transformer Model for 3D Object Detection. In *ICCV*, 2021.

- [33] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Trans. Graph.*, 41(4), July 2022.
- [34] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *ICCV*, 2021.
- [35] Joseph Ortiz, Alexander Clegg, Jing Dong, Edgar Sucar, David Novotny, Michael Zollhoefer, and Mustafa Mukadam. isdf: Real-time neural signed distance fields for robot perception. *arXiv preprint arXiv:2204.02296*, 2022.
- [36] 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.
- [37] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*. Springer, 2020.
- [38] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *ICCV*, 2019.
- [39] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 30, 2017.
- [40] Ruslan Rakhimov, Andrei-Timotei Ardelean, Victor Lempitsky, and Evgeny Burnaev. Npbg++: Accelerating neural point-based graphics. In *CVPR*, 2022.
- [41] Yongming Rao, Benlin Liu, Yi Wei, Jiwen Lu, Cho-Jui Hsieh, and Jie Zhou. Randomrooms: Unsupervised pre-training from synthetic shapes and randomized layouts for 3d object detection. In *ICCV*, 2021.
- [42] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *ICCV*, pages 14335–14345, 2021.
- [43] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 28, 2015.
- [44] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, 2015.
- [45] Maxim Tatarchenko, Stephan R Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*, 2019.
- [46] 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.
- [47] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022.
- [48] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *ICCV*, 2021.
- [49] Jingwen Wang, Tymoteusz Bleja, and Lourdes Agapito. Go-surf: Neural feature grid optimization for fast, high-fidelity rgb-d surface reconstruction. *arXiv preprint arXiv:2206.14735*, 2022.
- [50] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021.
- [51] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Sriniwasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. *arXiv preprint arXiv:2102.13090*, 2021.
- [52] 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*. Springer, 2020.
- [53] Wenpeng Xing, Jie Chen, and Yike Guo. Robust local light field synthesis via occlusion-aware sampling and deep visual feature fusion. *Machine Intelligence Research*, 20(3):408–420, 2023.
- [54] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *CVPR*, 2022.
- [55] Ryosuke Yamada, Hirokatsu Kataoka, Naoya Chiba, Yukiyasu Domae, and Tetsuya Ogata. Point cloud pre-training with natural 3d structures. In *CVPR*, 2022.
- [56] Siming Yan, Zhenpei Yang, Haoxiang Li, Li Guan, Hao Kang, Gang Hua, and Qixing Huang. Implicit autoencoder for point cloud self-supervised representation learning. *arXiv preprint arXiv:2201.00785*, 2022.
- [57] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021.
- [58] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. <https://arxiv.org/abs/2012.02190>, 2020.
- [59] 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.
- [60] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. NERF++: Analyzing and improving neural radiance fields. <https://arxiv.org/abs/2010.07492>, 2020.
- [61] 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.
- [62] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-supervised pretraining of 3d features on any point-cloud. In *ICCV*, pages 10252–10263, 2021.
- [63] Zaiwei Zhang, Bo Sun, Haitao Yang, and Qixing Huang. H3dnet: 3d object detection using hybrid geometric primitives. In *ECCV*. Springer, 2020.
- [64] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *ICCV*, 2021.