Ponder: Point Cloud Pre-training via Neural Rendering

Di Huang^{1,2} Sida Peng³ Tong He^{2,†} 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 pointcloud 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 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. The code will be released at https://dihuangdh.github.io/ponder.

1. Introduction

We have witnessed the widespread success of supervised learning in developing vision tasks, such as image classification [11, 17] and object detection [16, 42]. 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 [7, 18, 20, 22, 25, 31, 36, 40, 47, 52, 55, 58, 60, 61], 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 [7, 18, 20, 22, 40, 52, 61] and completionbased [25, 31, 36, 47, 55, 58, 60]. 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 and storing representative samples, and (3) applying vari-



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.

ous data augmentation policies. Inspired by these works, many works [7, 18, 20, 22, 40, 52, 61] 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 [12,46], 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. Specifically, given one or a sequence of RGB-D images, we lift them to 3D space and obtain a set of colored points. Points

[†] denote corresponding author.



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

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 to 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.

We conduct comprehensive experiments on a host of tasks. The consistent improvements demonstrate the effectiveness of our proposed Ponder. Our approach can serve as a strong alternative to contrast-based methods and completion-based methods in 3D point cloud pre-training.

2. Related Work

Neural rendering. Neural Rendering is a type of rendering technology that uses neural networks to differentiablely 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 [4,33,34,41,49,50,56,57,59] 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, 39] 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 [53] 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 [6,15] 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 [7, 18, 20, 22, 40, 61] 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 completionbased [25, 31, 36, 55, 58, 60] 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 [55] 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. [23] 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 en-



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 (visualized as an image in this figure) by average pooling. Finally, the 3D feature volume is rendered to multi-view RGB-D images via a differentiable neural rendering, which are compared with the original input multi-view RGB-D images as the supervision. Point cloud encoder f_p and color decoder f_c are used for transfer learning.

coder (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:

$$\mathcal{X} = \bigcup_{i}^{N} \pi^{-1}(I_i, D_i, \boldsymbol{\xi}_i, \mathbf{K}_i), \tag{1}$$

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} :

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

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

Once the feature extraction is done, we map the point embeddings \mathcal{E} to a 3D sparse feature volume. To fill in the

empty space, we perform average pooling, followed by a 3D CNN, to aggregate features from the nearby points. The dense 3D volume is denoted as V.

3.4. Pre-training with Neural Rendering

This section introduces how to reconstruct the implicit scene representation and render images differentiablely. We first give a brief introduction to neural scene representation, then illustrate how to integrate it into our point cloud pretraining 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 [49], the scene can be reconstructed as:

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

where f_s is the SDF decoder and f_c is the RGB color decoder. \tilde{f}_s takes point location **p** as input, and predicts the SDF value s. \tilde{f}_c takes point location **p** and viewing direction **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 **p** and viewing direction **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. For each rendering ray with camera origin **o** and viewing direction **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 **o** and **d** can be calculated from paired camera parameters $\{(\mathbf{K}_i, \boldsymbol{\xi}_i)\}$. z_n and z_f denote the near and far bounds of the ray. Different from previous methods [30, 49], 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:

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

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

where the \hat{C} is the ray color and the \hat{D} is the ray depth. We follow NeuS [49] to build an unbiased and occlusionawareness weight function w(z):

$$w(z) = T(z) \cdot \rho(z). \tag{7}$$

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

$$T(z) = \exp(-\int_{z_n}^{z_f} \rho(z) dz), \tag{8}$$

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

 $\Phi_h(x)$ 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, 49] to accelerate the training process.



Projected Points Rendered Color Reference Color Rendered Depth Reference Depth

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.

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 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,

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

which are loss functions responsible for color supervision L_c , depth supervision L_d , Eikonal regularization L_e , nearsurface 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:

$$L_c = \frac{1}{N_r} \sum_{i}^{N_r} ||\hat{C} - C||_2^2 \tag{11}$$

$$L_d = \frac{1}{N_r} \sum_{i}^{N_r} ||\hat{D} - D||_2^2,$$
(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:

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

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 [48], 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:

$$b(z) = D - z. \tag{14}$$

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 $b(z) \le t$, we leverage the near-surface SDF loss to constrain the SDF prediction $s(z_{i,j})$:

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

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

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

where α is set as 5 following the same with [35,48]. 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 [48], 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 [10] RGB-D images as our pre-training data. ScanNet is a widely used real-world in-door 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.

Data augmentation. We augment the point cloud by random sampling, normalization, and random masking. First, we randomly down-sample the point cloud to 20,000 points. Then, the point cloud is normalized into a 3D unit cube. Finally, we apply the same masking strategy as used in Mask Point [25]. Specifically, we use FPS to split the point cloud into 2,048 groups, each group containing 64 points, then mask the point groups with a mask ratio of 90%.

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. For transfer learning on 3D object detection task, we use VoteNet [38] as the baseline. VoteNet leverage a voting mechanism to generate object centers, which are used for 3D bounding box proposals. Two datasets are applied to verify the effectiveness of our method: ScanNet [10] and SUN RGB-D [44]. Different from ScanNet, which contains fully reconstructed 3D scenes, SUN RGB-D is a single-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 Point-Net++ as the point cloud encoder f_p , which is identical to the backbone used in VoteNet. We pre-train the point cloud encoder on the ScanNet dataset and transfer the weight as the VoteNet initialization. Following [38], we use average precision with 3D detection IoU threshold 0.25 and threshold 0.5 as the evaluation metrics.

Method	Detection	Pre-training	Pre-training	Pre-training	Scar	nNet	SUN RGB-D	
	Model	Туре	Data	Epochs	$AP_{50}\uparrow$	$\mathrm{AP}_{25}\uparrow$	$\mathrm{AP}_{50}\uparrow$	$\mathrm{AP}_{25}\uparrow$
3DETR [32]	3DETR	-	-	-	37.5	62.7	30.3	58.0
Point-BERT [58]	3DETR	Completion	3D Model	300	38.3	61.0	-	-
MaskPoint [25]	3DETR	Completion	Depth	300	40.6	63.4	-	-
VoteNet [38]	VoteNet	-	-	-	33.5	58.6	32.9	57.7
STRL [20]	VoteNet	Contrast	Depth	100	38.4	59.5	35.0	58.2
RandomRooms [40]	VoteNet	Contrast	Synthesis	300	36.2	61.3	35.4	59.2
PointContrast [52]	VoteNet	Contrast	3D Model	-	38.0	59.2	34.8	57.5
PC-FractalDB [54]	VoteNet	Contrast	Synthesis	-	38.3	61.9	33.9	59.4
DepthContrast [61]	VoteNet	Contrast	Depth	1000	39.1	62.1	35.4	60.4
IAE [55]	VoteNet	Completion	3D Model	1000	39.8	61.5	36.0	60.4
Ponder	VoteNet	Rendering	Depth	100	40.9	64.2	36.1	60.3
Ponder	VoteNet	Rendering	Color & Depth	100	41.0	63.6	36.6	61.0

Table 1. **3D** object detection AP_{25} and AP_{50} on ScanNet and SUN RGB-D. VoteNet [38] and 3DETR [32] are two baseline 3D object detection models. The DepthContrast [61] and Point-BERT [58] results are adopted from IAE [55] and MaskPoint [25]. Ponder outperforms both VoteNet-based and 3DETR-based point cloud pre-training methods with fewer training epochs.

The 3D detection results are shown in Table 1. Our method improves the baseline of VoteNet without pretraining by a large margin, boosting AP_{50} by 7.5% and 3.7% for ScanNet and SUN RGB-D, respectively. IAE [55] 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 AP_{50} and AP_{25} of our method are higher than that of IAE by 1.2% and 2.1% on ScanNet, respectively. Mask-Point [25] is another method aiming to learn a continuous surface by classifying if the query point is occupied. However, its performance can be constrained due to the noisy labeling of the query point occupancy value. As presented in Table 1, even with an inferior backbone (PointNet++ vs 3DETR), our method is able to achieve better accuracy with fewer pre-training epochs.

3D semantic segmentation. 3D semantic segmentation is another fundamental scene understanding task. Following [43,47,55], we choose DGCNN [51] as our baseline for a fair comparison. DGCNN applies a dynamic graph CNN as the backbone. For pre-training, we use DGCNN as the point cloud encoder f_p , and pre-train the model on ScanNet. We validate the effectiveness of our method by transferring the weights to Stanford Large-Scale3D Indoor Spaces (S3DIS) [3] dataset, which is an indoor 3D understanding dataset containing 6 large-scale indoor scenes with point semantic annotations. Following the same setting of [51], we use the overall accuracy (OA) mean IoU(mIoU) on points as the evaluation metric, and report the average evaluation results across six folds.

Table 2 shows the quantitative results. Compared with the DGCNN baseline, the proposed method boost the seg-

mentation performance by a large margin, boosting OA and mIoU for 2.1% and 5%, respectively. Jigsaw and OcCo use ShapeNet as the pre-train dataset. Although they get improvements compared with the baseline, the limited scale of training data constrains the transferring ability. IAE achieves significant improvements by leveraging the large-scale dataset and an implicit reconstruction manner. Compared with IAE, the proposed approach achieves a higher semantic segmentation performance with the DGCNN backbone (+0.3% for OA and +0.4% for mIoU). Besides, IAE requires a large amount of 3D mesh for supervision. Our approach, in contrast, only requires RGB-D images as the supervision, which is much cheaper and easy to fetch.

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 are widely adopted in [9, 26, 56]. Following the same setting as ConvONet, we conduct experiments on the Synthetic Indoor Scene Dataset (SISD) [37], which is a synthetic dataset and contains 5000 scenes with multiple ShapeNet [5] objects. We pre-train the PointNet encoder, which is the same as the original Con-

Method	OA↑	mIoU↑
DGCNN [51]	84.1	56.1
Jigsaw [43]	84.4	56.6
OcCo [47]	85.1	58.5
IAE [55]	85.9	60.7
Ponder	86.2	61.1

Table 2. **3D** semantic segmentation *OA* and *mIoU* on S3DIS dataset with DGCNN model. Ponder outperforms previous state-ofthe-art models.

Method	Encoder	IoU↑	Normal Consistency↑	F-Score↑
ConvONet [37]	PointNet	84.9	0.915	0.964
Ponder	PointNet	85.7	0.917	0.965
ConvONet	PointNet++	77.8	0.887	0.906
Ponder	PointNet++	80.2	0.893	0.920

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



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.

vONet implementation, and test the reconstruction quality on the SISD dataset. Additionally, to compare with another self-supervised learning method IAE [55], we add extra experiments using VoteNet-style PointNet++ as the encoder of ConvONet. Following [37], we use Volumetric IoU, Normal Consistency, and F-Score [45] with the threshold value of 1% as the evaluation metrics.

The results are shown in Table 3. Compared to the baseline ConvONet model with PointNet, the proposed approach is able to improve the reconstruction quality (+0.8% for IoU). By replacing the encoder of ConvONet from PointNet to PointNet++, ours is able to achieve more accuracy improvement (+2.4% for IoU and +0.014 for F-Score). Our method also gets better reconstruction results than IAE. Check our supplementary materials for more details.

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 [53] 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 re-

Supervision	ScanNet	SUN RGB-D		
Depth	40.9	36.1		
Color	40.5	35.8		
Color+RGB	41.0	36.6		

Table 4. 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.

place 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 pretrained 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.

4.3. Ablation study

In this section, we do two ablation experiments. First, we show the effectiveness of using different 2D supervision. Then, we test how the view number affects the final performance. For both experiments, we use the 3D object detection task as the transfer learning task.

Influence on 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 4. 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.



Figure 6. **Direct applications of Ponder** on the ScanNet validation set. The proposed Ponder model can be directly used for various applications, such as 3D reconstruction and image synthesis. The input point clouds are drawn as spheres for better clarity.

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 5. 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 single view, multiple views achieve higher accuracy, boosting AP₅₀ by 0.9% and 1.2% for ScanNet and Sun RGB-D datasets, respectively.

4.4. Other applications

In previous sections, we show that the proposed pipeline can be used for transfer learning. In this section, we show that the pre-trained model from our pipeline Ponder itself can also be directly used for surface reconstruction and image synthesis from sparse point clouds.

3D reconstruction from sparse point clouds. The learned model has the capability to recover the scene surface 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.

Image synthesis from sparse point clouds. Another interesting experiment to explore is that our pipeline is able to render realistic images from sparse point cloud input. As shown in Figure 6, our method is able to recover similar color images with the ground truth. Also, the recovered

depth may even look better compared with the ground-truth depth image which has irregular values.

View number	ScanNet	SUN RGB-D		
1 view	40.1	35.4		
3 views	40.8	36.0		
5 views	41.0	36.6		

Table 5. 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.

5. Conclusion

In this paper, we show 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 not only high-level 3D perception tasks but also 3D low-level tasks, like 3D reconstruction and image synthesis from point clouds. Also, the learned Ponder model can be directly used for 3D reconstruction and image synthesis from sparse point clouds.

Several directions could be explored in future works. First, there are various types of neural rendering, which could also be leveraged for point cloud representation learning. Second, other 3D domain-specific designs could be integrated into point cloud pre-training pipelines. Third, exploring the proposed pre-training pipeline Ponder on a larger dataset and more downstream tasks is also a potential research direction.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9902–9912, 2022. 2
- [2] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *European Conference on Computer Vision*, pages 696–712. Springer, 2020. 2
- [3] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1534–1543, 2016. 6
- [4] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 2
- [5] 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. 6
- [6] 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. 2
- [7] Yujin Chen, Matthias Nießner, and Angela Dai. 4dcontrast: Contrastive learning with dynamic correspondences for 3d scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 1, 2
- [8] Zhang Chen, Yinda Zhang, Kyle Genova, Sean Fanello, Sofien Bouaziz, Christian Häne, Ruofei Du, Cem Keskin, Thomas Funkhouser, and Danhang Tang. Multiresolution deep implicit functions for 3d shape representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13087–13096, 2021. 12
- [9] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6970–6981, 2020. 6, 12
- [10] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2017. 5
- [11] 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. 1

- [12] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. arXiv preprint arXiv:2205.09113, 2022. 1
- [13] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. arXiv preprint arXiv:2002.10099, 2020. 5
- [14] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 16000– 16009, 2022. 1, 2
- [15] 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. 2
- [16] 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
- [17] 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. 1
- [18] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15587–15597, 2021. 1, 2
- [19] Ji Hou, Saining Xie, Benjamin Graham, Angela Dai, and Matthias Nießner. Pri3d: Can 3d priors help 2d representation learning? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5693–5702, 2021. 2
- [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. 1, 2, 6
- [21] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 406–413, 2014. 7
- [22] 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 *Proceedings of the IEEE/CVF International Conference* on Computer Vision, pages 6423–6432, 2021. 1, 2
- [23] Lanxiao Li and Michael Heizmann. A closer look at invariances in self-supervised pre-training for 3d vision. arXiv preprint arXiv:2207.04997, 2022. 2
- [24] Lanxiao Li and Michael Heizmann. A closer look at invariances in self-supervised pre-training for 3d vision. In ECCV, 2022. 14
- [25] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. arXiv preprint arXiv:2203.11183, 2022. 1, 2, 5, 6, 13
- [26] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. Advances in Neural Information Processing Systems, 33:15651–15663, 2020. 6

- [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. 8
- [28] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. ACM siggraph computer graphics, 21(4):163–169, 1987. 8
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101, 2017. 5
- [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):99–106, 2021. 2, 4, 12
- [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. 1, 2
- [32] Ishan Misra, Rohit Girdhar, and Armand Joulin. An End-to-End Transformer Model for 3D Object Detection. In *ICCV*, 2021. 6
- [33] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. ACM Trans. Graph., 41(4):102:1– 102:15, July 2022. 2
- [34] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5589–5599, 2021. 2
- [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. 5
- [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. 1, 2
- [37] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision*, pages 523–540. Springer, 2020. 6, 7, 12, 13
- [38] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In Proceedings of the IEEE International Conference on Computer Vision, 2019. 5, 6
- [39] Ruslan Rakhimov, Andrei-Timotei Ardelean, Victor Lempitsky, and Evgeny Burnaev. Npbg++: Accelerating neural point-based graphics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15969–15979, 2022. 2
- [40] Yongming Rao, Benlin Liu, Yi Wei, Jiwen Lu, Cho-Jui Hsieh, and Jie Zhou. Randomrooms: Unsupervised pretraining from synthetic shapes and randomized layouts for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3283–3292, 2021. 1, 2, 6

- [41] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps, 2021. 2
- [42] 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
- [43] Jonathan Sauder and Bjarne Sievers. Self-supervised deep learning on point clouds by reconstructing space. Advances in Neural Information Processing Systems, 32, 2019. 6, 7
- [44] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 567–576, 2015. 5
- [45] Maxim Tatarchenko, Stephan R Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 3405–3414, 2019. 7
- [46] 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. 1
- [47] 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, 6, 7
- [48] Jingwen Wang, Tymoteusz Bleja, and Lourdes Agapito. Go-surf: Neural feature grid optimization for fast, highfidelity rgb-d surface reconstruction. arXiv preprint arXiv:2206.14735, 2022. 5, 12
- [49] 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. 2, 3, 4, 12
- [50] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, 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. 2
- [51] 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), 2019. 6, 7
- [52] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pretraining for 3d point cloud understanding. In *European conference on computer vision*, pages 574–591. Springer, 2020. 1, 2, 6
- [53] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5438–5448, 2022. 2, 7, 12
- [54] Ryosuke Yamada, Hirokatsu Kataoka, Naoya Chiba, Yukiyasu Domae, and Tetsuya Ogata. Point cloud pretraining with natural 3d structures. In *Proceedings of the*

IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 21283–21293, 2022. 6

- [55] 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. 1, 2, 6, 7, 13
- [56] 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. 2, 6
- [57] 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. 2
- [58] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2022. 1, 2, 6
- [59] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. NERF++: Analyzing and improving neural radiance fields. https://arxiv.org/abs/2010.07492, 2020. 2
- [60] 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
- [61] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-supervised pretraining of 3d features on any point-cloud. 2021. 1, 2, 6

Ponder: Point Cloud Pre-training via Neural Rendering Supplementary Material

A. Implementation Details

In this section, we give more implementation details of our **Ponder** model. **Our code will be released upon acceptance.**

A.1. Pre-training Details

3D feature volume. In our experiments, we build a hierarchical feature volume \mathcal{V} with a resolution of [16, 32, 64]. Building a 3D hierarchical feature volume has been wildly used for recovering detailed 3D geometry, e.g. [8,9]. After processing the 3D feature volume with a 3D CNN, we use trilinear interpolation to get the feature of the query point **p**, denoted as $\mathcal{V}(\mathbf{p})$. We use the drop-in replacement of *grid_sampler* from [48] to accelerate the training.

Ray sampling strategy. Similar to [30, 49], we sample twice for each rendering ray. First, we uniformly sample coarse points between the near bound z_n and far bound z_f . Then, we use importance sampling with the coarse probability estimation to sample fine points. Folowing [49], the coarse probability is calculated based on $\Phi_h(s)$. By this sampling strategy, our method can automatically determine sample locations and can collect more points near the surface, which makes the training process more efficient.

Back projection Here we give details of the back projection function π^{-1} to get point clouds from depth images. Let **K** be camera intrinsic parameters, $\xi = [\mathbf{R}|\mathbf{t}]$ be camera extrinsic parameters, where **R** is the rotation matrix and **t** is the translation matrix. X_{uv} is the projected point location and X_w is the point location in the 3D world coordinate. Then, according to the pinhole camera model:

$$s\boldsymbol{X}_{uv} = \mathbf{K}(\mathbf{R}\boldsymbol{X}_w + \mathbf{t}), \tag{17}$$

where s is the depth value. After expanding the X_{uv} and X_w :

$$s \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{K} (\mathbf{R} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} + \mathbf{t}).$$
(18)

Then, the 3D point location can be calculated as follows:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \mathbf{R}^{-1} (\mathbf{K}^{-1} s \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} - \mathbf{t})$$
(19)

The above Equation 19 is the back-projection equation π^{-1} used in this paper.

Training Time. The **Ponder** model is trained with 8 NVIDIA A100 GPUs for 96 hours.

A.2. Transfer Learning Details

3D scene reconstruction. ConvONet [37] reconstructs scene geometry from the point cloud input. It follows a two-step manner, which first encodes the point cloud into a 3D feature volume or multiple feature planes, then decodes the occupancy probability for each query point. To test the transfer learning ability of our point cloud encoder, we directly replace the point cloud encoder of ConvONet, without any other modification. We choose the highest performing configuration of ConvONet as the baseline setting, which uses a 3D feature volume with a resolution of 64. For the training of ConvONet, we follow the same training setting as the released code¹.

Image synthesis from point clouds. Point-NeRF [53] renders images from neural point cloud representation. It first generates neural point clouds from multi-view images, then uses point-based volume rendering to synthesize images. To transfer the learned network weight to the Point-NeRF pipeline, we 1) replace the 2D image feature backbone with a pre-trained point cloud encoder to get the neural point cloud features, 2) replace the color decoder by a pretrained color decoder, 3) keep the other Point-NeRF module untouched. Since a large amount of point cloud is hard to be directly processed by the point cloud encoder, we downsample the point cloud to 1%, which will decrease the rendering quality but help reduce the GPU memory requirements. We report the PSNR results of the unmasked region as the evaluation metric, which is directly adopted from the original codebase². For training Point-NeRF, we follow the same setting as Point-NeRF.

B. Supplementary Experiments

B.1. Ablation Study

¹https://github.com/autonomousvision/convolutional_occupancy_networks ²https://github.com/Xharlie/pointnerf

Mask ratio	ScanNet	SUN RGB-D
0%	40.7	37.3
25%	40.7	36.2
50%	40.3	36.9
75%	41.7	37.0
90%	41.0	36.6

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

Resolution	ScanNet	SUN RGB-D		
16	40.7	36.6		
16+32+64	41.0	36.6		

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

Method	IoU↑	Normal Consistency↑	F-Score↑
ConvOcc [37]	0.778	0.887	0.906
IAE [55]	0.757	0.887	0.910
Ours	0.802	0.893	0.920

Table 8. **3D** scene reconstruction *IoU*, *NC*, and *F*-Score on SISD dataset with PointNet++ model.

Influence on mask ratio. In this paper, we use random masking as one type of point cloud augmentation. We apply the same mask ratio as MaskPoint [25]. Here, we give additional experimental results to show the influence of using different mask ratios in Table 6. For the mask ratio of 0%, we do not apply any mask strategy to the input point cloud.

3D feature volume resolution. As mentioned in Section **A**, **Ponder** build a 3D feature volume with a resolution of [16, 32, 64], which is inspired by the recent progress of multi-resolution in 3D reconstruction. However, building such a 3D feature volume with large resolutions requires heavy GPU memory usage. We conduct experiments in Table 7 to test the performance with a smaller resolution. As shown in the table, even with a small resolution, **Ponder** is still able to achieve comparable accuracy, demonstrating the robustness to the feature volume resolution.

B.2. Transfer Learning

3D scene reconstruction As mentioned in the paper, we transfer the learned PointNet++ model of IAE to the 3D reconstruction task. The results are shown in Table 8. Compared with the ConvONet baseline, the IAE pre-trained model gets a better F-Score with 0.004 but gets worse results on the IoU metric. Our method, on the other hand,



Figure 7. **Label efficiency training.** We show the 3d object detection experiment results using limited downstream data. Our pretrained model is capable of achieving better performance than training from scratch using the same percentage of data or requires fewer data to get the same detection accuracy.

gets a better reconstruction performance than both the ConvONet and IAE.

Label Efficiency Training. We also do experiments to show the performance of our method with limited labeling for the downstream task. Specifically, we test the label efficiency training on the 3D object detection task for ScanNet. Following the same setting with IAE [55], we use 20%, 40%, 60%, and 80% of ground truth annotations. The results are shown in Figure 7. We show constantly improved results over training from scratch, especially when only 20% of the data is available.

Color information for downstream tasks. Different from previous works, since our pre-training model uses a colored point cloud as the input, we also use color information for the downstream tasks. Results are shown in Table 9. Using color as an additional point feature can help the VoteNet baseline achieve better performance on the SUN RGB-D dataset, but get little improvement on the ScanNet dataset. This shows that directly concatenating point positions and colors as point features shows limited robustness to application scenarios. By leveraging the proposed **Ponder** pre-training method, the network is well initialized to handle the point position and color features, and achieve better detection accuracy.

More comparisons on 3D detection. More detection accuracy comparisons are given in Table 9. Even using an inferior backbone, our **Ponder** model is able to achieve similar detection accuracy with 9 in ScanNet and better accuracy in SUN RGB-D.

Method	Detection	Pre-training	Pre-training	Pre-training	Scar	nNet	SUN F	RGB-D
	Model	Туре	Data	Epochs	$AP_{50}\uparrow$	$\mathrm{AP}_{25}\uparrow$	$\mathrm{AP}_{50}\uparrow$	$\mathrm{AP}_{25}\uparrow$
VoteNet*	VoteNet*	-	-	-	37.6	60.0	33.3	58.4
DPCo [24]	VoteNet*	Contrast	Depth	120	41.5	64.2	35.6	59.8
IPCo [24]	VoteNet*	Contrast	Color & Depth	120	40.9	63.9	35.5	60.2
VoteNet (w color)	VoteNet	-	-	-	33.4	58.8	34.3	58.3
Ponder	VoteNet	Rendering	Depth	100	40.9	64.2	36.1	60.3
Ponder	VoteNet	Rendering	Color & Depth	100	41.0	63.6	36.6	61.0

Table 9. **3D object detection** AP_{25} and AP_{50} on ScanNet and SUN RGB-D. * means a different but stronger version of VoteNet.

B.3. More application examples

As mentioned in the paper, the pre-trained **Ponder** model can be directly used for surface reconstruction and image synthesis tasks. We give more application examples in Figure 8 and Figure 9.



Figure 8. More results of application examples of Ponder on the ScanNet validation set (part 1). The input point clouds are drawn as spheres for better clarity.



Figure 9. More results of application examples of Ponder on the ScanNet validation set (part 2). The input point clouds are drawn as spheres for better clarity.