

Instance-aware Dynamic Prompt Tuning for Pre-trained Point Cloud Models

Yaohua Zha^{1*} Jinpeng Wang^{1*} Tao Dai^{2†} Bin Chen³ Zhi Wang¹ Shu-Tao Xia⁴

Tsinghua Shenzhen International Graduate School, Tsinghua University¹
College of Computer Science and Software Engineering, Shenzhen University²

Harbin Institute of Technology, Shenzhen³

Research Center of Artificial Intelligence, Peng Cheng Laboratory⁴

{chayh21, wjp20}@mails.tsinghua.edu.cn

Abstract

Pre-trained point cloud models have found extensive applications in 3D understanding tasks like object classification and part segmentation. However, the prevailing strategy of full fine-tuning in downstream tasks leads to large per-task storage overhead for model parameters, which limits the efficiency when applying large-scale pre-trained models. Inspired by the recent success of visual prompt tuning (VPT), this paper attempts to explore prompt tuning on pre-trained point cloud models, to pursue an elegant balance between performance and parameter efficiency. We find while instance-agnostic static prompting, e.g. VPT, shows some efficacy in downstream transfer, it is vulnerable to the distribution diversity caused by various types of noises in real-world point cloud data. To conquer this limitation, we propose a novel Instance-aware Dynamic Prompt Tuning (IDPT) strategy for pre-trained point cloud models. The essence of IDPT is to develop a dynamic prompt generation module to perceive semantic prior features of each point cloud instance and generate adaptive prompt tokens to enhance the model’s robustness. Notably, extensive experiments demonstrate that IDPT outperforms full fine-tuning in most tasks with a mere 7% of the trainable parameters, providing a promising solution to parameter-efficient learning for pre-trained point cloud models. Code is available at <https://github.com/zyh16143998882/ICCV23-IDPT>.

1. Introduction

With the rapid development of 3D scanning technology, point clouds, as irregular point sets that represent

*These authors contributed equally to this work.

†Corresponding author. (✉ daitao.edu@gmail.com)

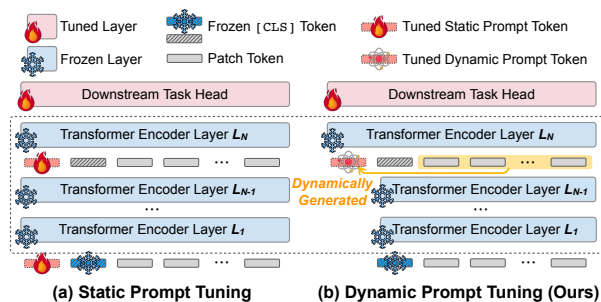


Figure 1. The pipeline of (a) the previous static prompt tuning in VPT [22] and (b) our dynamic prompt tuning. Unlike the static prompt tuning that is instance-agnostic, ours is adaptive to input by concatenating the instance-aware prompt generated by a prompt module into the last Transformer layer input.

3D geometry, have been widely used in various fields and tasks. Deep learning-based point cloud processing techniques [16, 24, 31, 33, 34, 45, 35] have drawn considerable attention as they can directly process raw point cloud data while preserving its rich information. As a classic deep learning paradigm, fine-tuning the foundation model pre-trained [11, 26, 32, 51, 55] on massive raw point clouds in specific downstream tasks has achieved state-of-the-art performance. However, this approach is storage-intensive, as it requires storing and deploying a separate copy of the backbone parameters for each task.

Recently, prompt tuning has surpassed fine-tuning in multiple downstream tasks in the language and image domains, significantly reducing storage requirements by fixing the parameters of a pre-trained model and introducing a small amount of task-specific learnable parameters into the input space. Although some work [21, 20, 46, 54] have attempted to introduce prompt into point cloud processing, they all relied on pre-trained image models [38, 12]. To date, less research has been denoted to prompt tuning in

point cloud pre-trained models.

Inspired by the success of visual prompt tuning (VPT [22]), it is natural to adopt this idea to point clouds. As shown in Figure 1(a), we call VPT a static prompting strategy because it introduces prompt tokens as a few learnable parameters concatenated to the input of a pre-trained point cloud model. The prompt tokens are shared by any input and therefore are instance-agnostic. Although such a strategy performs well on synthetic datasets (*e.g.* ModelNet40 [47]), it causes significant performance degradation on real scanned point cloud datasets (*e.g.* ScanObjectNN [41]). Thus, static prompt tuning is not suitable for real point clouds, where point clouds with different types of missing or noisy points belong to different distributions. These observations motivate us to design a universal prompt-tuning strategy for both synthetic and real point clouds.

To address this issue, we proposed an Instance-aware Dynamic Prompt Tuning (IDPT) for point cloud pre-trained models. As shown in Figure 1(b), IDPT develops a prompt generation module to perceive the semantic prior features of each point cloud instance and produces adaptive prompts for different inputs. The proposed IDPT enables an adjusting effect to mitigate the adverse noises in point cloud instances and thus can enhance the robustness of pre-trained models. We insert IDPT into the last Transformer layer for a more accurate representation of point clouds. Extensive experiments demonstrate the effectiveness of IDPT. Typically, IDPT yields competitive performance compared with full fine-tuning but just requires about 7% of trainable parameters in downstream transfer.

The main contributions can be summarized as follows:

- To our best knowledge, this is the first exploration of prompt tuning on pre-trained point cloud models. We reveal that VPT, the static prompting strategy, suffers from the distributional diversity issue caused by various types of noises in real-world point cloud data.
- To address the shortcoming, we propose an Instance-aware Dynamic Prompt Tuning (IDPT) as an effective solution. We develop a dynamic prompt generation module to capture semantic prior features of each point cloud instance and generate adaptive prompt tokens to mitigate the noises.
- Extensive experiments on a variety of downstream tasks show the competitive performance of IDPT with full fine-tuning in most tasks while requiring much less tunable parameters, *e.g.* about 7%.

2. Related Work

2.1. Pre-training Point Cloud Models

Recently, studies on pre-trained foundational models for 3D point clouds have achieved remarkable success. These

approaches first apply a pretext task to pre-train the foundational model to learn the latent semantic information of the point cloud and then fine-tune the model weights for the target task to achieve higher performance. Existing pre-train pretext tasks can be divided into discriminative tasks [4, 9, 14, 48] and generative tasks [3, 11, 19, 25, 26, 32, 51, 55]. The discriminative approach distinguishes different views of the same point cloud instance from other instances, PointContrast [49] and CrossPoint [1] explore the use of contrast learning of intra-domain and cross-domain features to obtain rich self-supervised information. Generation methods typically rely on an autoencoder to learn the latent features of the data by reconstructing the original input. Point-BERT [51], Point-MAE [32] and PointM2AE [55], based on masked autoencoders, have been very successful. Additionally, Point-DAE [55] explores a more general denoising autoencoder for point cloud learning by investigating more types of corruption beyond masking. ACT [11] achieves a significant improvement on real scanned point clouds by using pre-trained language models and image models as cross-modal teachers to guide the learning of 3D self-supervised networks. However, the above methods all utilize full fine-tuning to adapt pre-trained models to various downstream tasks. Our work further explores how to reduce parameter storage in downstream tasks by utilizing prompt tuning, building upon the aforementioned approach.

2.2. Prompt Learning in Computer Vision

Prompt tuning involves adding specific prompt information to the input of a pre-trained model and adjusting downstream tasks to fit the pre-trained model. This is achieved by fixing the pre-trained model parameters and fine-tuning the prompt. It was first proposed in the language model [7, 13, 23, 27, 28, 29] and gained popularity in the image model [38, 39, 40, 57, 56] later due to its flexibility and high performance. CLIP [38] uses fixed class-specific text labels as prompts for prediction. Later, CoOp [57] learns class-specific continuous prompts, and CoCoOp [56] builds upon CoOp by introducing a lightweight network to learn dynamic prompts for each instance. VPT [22] first introduces the continuous prompt tuning framework into image pre-trained models inspired by P-Tuning [28]. Additionally, P2P [46] achieved the first application of prompts in point clouds by learning color information in the input space of point cloud rendering images as prompts for the 2D backbone network. PointCLIP [54] and CLIP2Point [20] project the point cloud as a depth map and then use the pre-trained CLIP [38] to understand the point cloud. However, all the above work relies on pre-trained image models. Our work discusses the tuning of pre-trained point cloud models with appropriate prompting mechanisms.

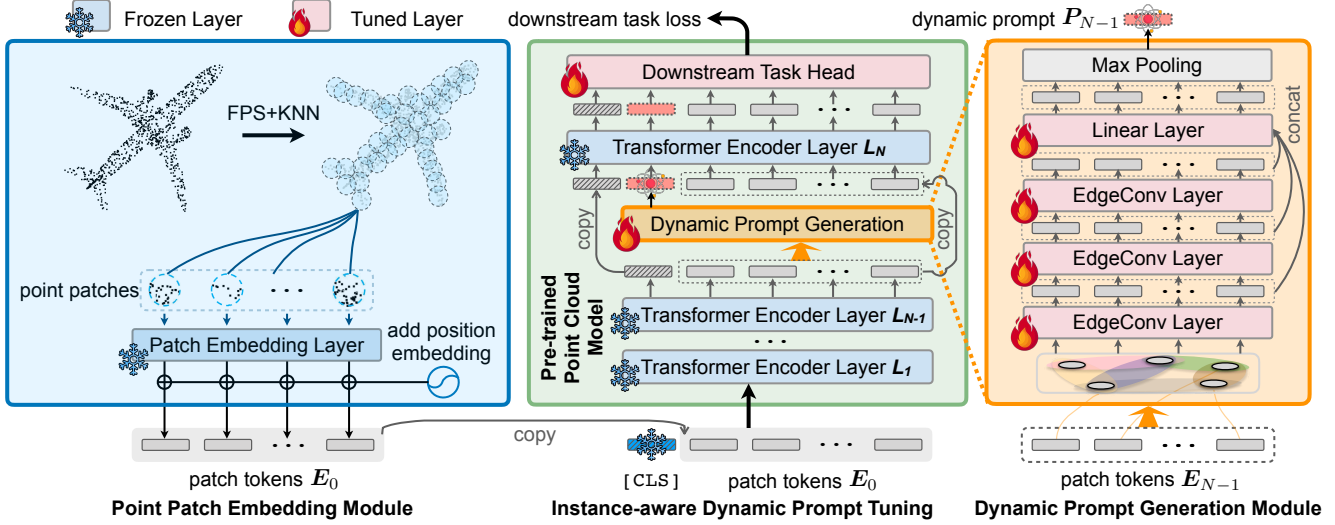


Figure 2. Overall pipeline of Instance-aware Dynamic Prompt Tuning (IDPT) for pre-trained point cloud models, which only updates the parameters of the dynamic prompt generation module and downstream task head during a downstream tuning task. To capture various sub-modes existing in the real-world data and enhance the robustness against noises (e.g., with different types of missing or noisy points), we design a dynamic prompt generation module with graph convolution [45] layers to aggregate multi-scale contextual features and dynamically generate instance-adaptive prompt. Empirically, inserting the dynamic prompt before the last transformer layer yields promising performance and enjoys decent efficiency at the same time.

3. Methodology

In this section, we first introduce the tuning pipeline for a pre-trained point cloud model (§ 3.1). Next, we present the empirical observation of static prompt tuning (e.g. VPT [22]) and discuss its weaknesses (§ 3.2) that highlight our motivations. At last, we describe our Instance-aware Dynamic Prompt Tuning strategy in detail (§ 3.3).

3.1. Preliminaries

When fine-tuning a pre-trained point cloud model (e.g. Point-MAE [32]), a point cloud $\mathbf{X} \in \mathbb{R}^{M \times 3}$ with M points is first divided into m point patches $\mathbf{X}' \in \mathbb{R}^{m \times k \times 3}$ via Farthest Point Sampling (FPS) and K-Nearest Neighborhood (KNN) algorithms, where each patch has k local points. Then, all point patches will be embedded into a series of input tokens $\mathbf{E}_0 \in \mathbb{R}^{m \times d}$ with positional encoding via a point patch embedding module. Next, we insert a classification token (i.e., [CLS]) \mathbf{c}_0 at the head of the patch embeddings and forward the token embeddings to the pre-trained model. Specifically, the forward process of each transformer layer is defined as

$$[\mathbf{c}_i; \mathbf{E}_i] = f_i([\mathbf{c}_{i-1}; \mathbf{E}_{i-1}]), \quad i = 1, 2, \dots, N, \quad (1)$$

where f_i denotes the i -th transformer encoder layer. N is the total transformer layer number of the pre-trained backbone. Finally, the model makes predictions by building a task-specific head g_h upon the output of the pre-trained

backbone:

$$\mathbf{y} = g_h([\mathbf{c}_N; \mathbf{E}_N]). \quad (2)$$

All the parameters of $\{f_i\}_{i=1}^N$ and g_h will be updated in a downstream tuning task, which burdens the storage cost for per-task model weights.

Recently, prompt [57, 22] has shown to be effective for parameter-efficient tuning. The basic idea of prompt tuning is to insert a few learnable prompt tokens into the input token sequence, i.e., we modify Eq.(1) and Eq.(2) by

$$[\mathbf{c}_i; \mathbf{P}_i; \mathbf{E}_i] = f_i([\mathbf{c}_{i-1}; \mathbf{P}_{i-1}; \mathbf{E}_{i-1}]), \quad i = 1, 2, \dots, N, \quad (3)$$

$$\mathbf{y} = g_h([\mathbf{c}_N; \mathbf{P}_N; \mathbf{E}_N]), \quad (4)$$

where \mathbf{P}_i is the inserted prompt tokens at the i -th layer. During the tuning process, we freeze the parameters of $\{f_i\}_{i=1}^N$ and only update prompt $\{\mathbf{P}_i\}_{i=1}^N$ and task-specific head g_h , which can largely reduce per-task storage cost.

3.2. Observation and Discussion

Inspired by the success of Visual Prompt Tuning (VPT) [22], it is natural to extend such a prompting tuning strategy to point clouds, as shown in Figure 1(a). As the prompt tokens are instance-independent and shared by all samples during downstream tuning, we term this kind of strategy *static prompt tuning*. Our empirical study showed that although VPT improves downstream performance compared with tuning the task head only, it underperforms fully

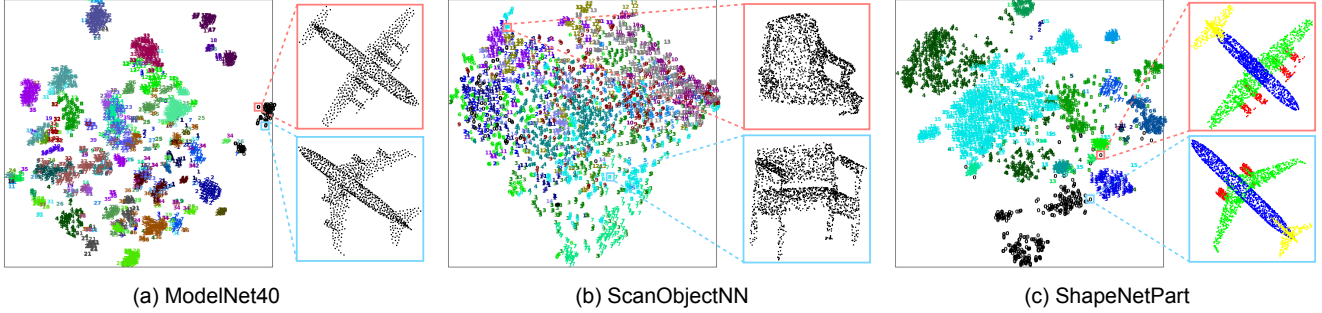


Figure 3. The t-SNE [42] visualization of the point cloud features extracted from with the Point-MAE [32] on three datasets: (a) ModelNet40, (b) ScanObjectNN, and (c) ShapeNetPart. Different from synthetic datasets (*e.g.* ModelNet40 and ShapeNetPart) with clean and compact cluster structures, in real-world datasets (*e.g.* ScanObjectNN), instances within the same category can present various sub-modes (*i.e.* sub-clusters scattered in the embedding space) because real-world point clouds contain varying types of missing or noisy points.

Tuning Strategy	#TP (M)	ModelNet40	ScanObjectNN		
			OBJ_BG	OBJ_ONLY	PB_T50_RS
A. Only Head Tuning	0.27	93.2	87.40	87.13	80.33
B. VPT-Shallow [22]	0.28	93.4	87.61	89.04	80.99
C. VPT-Deep [22]	0.36	93.6	89.98	90.19	83.96
D. IDPT (Ours)	1.69	94.4	93.63	93.12	88.51
E. Full Fine-tuning	22.10	93.8	92.94	92.08	88.41

Table 1. Classification accuracy (%) for different tuning strategies is reported. All experiments were conducted based on a pre-trained Point-MAE [32] model, and a simple rotation augmentation in ACT [11] was employed on the ScanObjectNN [41] dataset. ‘#TP’ denotes the number of trainable parameters.

fine-tuning by considerable margins. In particular, Table 1 presents the results of classification on several benchmark datasets. We can learn that static prompt tuning significantly reduces the number of trainable parameters (*i.e.*, about 1% to 2% of backbone parameters) compared with fully fine-tuning. In terms of accuracy, although VPT-Deep performs well on synthetic datasets (*e.g.* ModelNet40 [47]), it presents significant performance degradation on real scanned point cloud datasets. For example, on the PB_T50_RS variant of ScanObjectNN [41] dataset, fully fine-tuning outperforms VPT-Deep by 4.5% accuracy.

Here we briefly discuss why static prompt tuning performs well on synthetic datasets but poorly on real scanned datasets. We adopt a perspective from Domain Adaptation (DA) [2, 5, 6, 17] and consider the transferring from pre-trained models to downstream tasks. Our goal is to bridge the source and target domains with different distributions so as to enhance the prediction robustness of the pre-trained model. By definition, the source domain $p_s(\mathbf{x}_s)$ refers to the distribution of pre-training data, and the target domain $p_t(\mathbf{x}_t)$ refers to the distribution of downstream task data. Empirically, we found that

$$p_t(\mathbf{x}_t) \neq p_s(\mathbf{x}_s), \quad (5)$$

while we ask for a robust model such that

$$p_t(\mathbf{y}_t|\mathbf{x}_t) = p_s(\mathbf{y}_s|\mathbf{x}_s), \quad (6)$$

where \mathbf{x} denotes the input and \mathbf{y} denotes the output. Therefore, the goal of domain adaptation is to find a transformation $\Phi(\cdot)$, such that

$$p_t(\Phi(\mathbf{x}_t)) = p_s(\Phi(\mathbf{x}_s)). \quad (7)$$

Both fine-tuning and prompt tuning can be regarded as approaches to approximate the transformation $\Phi(\cdot)$. Fine-tuning adjusts all parameters of the pre-trained model to fit $\Phi(\cdot)$, whereas prompt tuning introduces additional prompt parameters to fit $\Phi(\cdot)$. From Table 1, we can learn that full fine-tuning achieves satisfactory performance by fitting all model parameters. In contrast, the static prompt shows inflexibility in mitigating the domain gap, suffering from the noises in the target domain.

We also provide an intuitive analysis of such a phenomenon. Specifically, Figure 3 shows the t-SNE [42] visualization of point cloud features extracted with Point-MAE [32] on the test sets of three datasets: (a) ModelNet40 [47], (b) ScanObjectNN [41], and (c) ShapeNetPart [50]). It reflects the downstream task data distribution (*i.e.*, target domain) to some degree. More details about the visualization of t-SNE are in the supplemental experiments. As shown in the figure, on synthetic datasets like ModelNet40 and ShapeNetPart, instances from the same class tend to distribute in relatively clear and tight clusters. Quite differently on the real scanned dataset ScanObjectNN, instances from the same class scatter to many sub-clusters in the feature space, indicating the mixing of various sub-distributions corresponding to different sub-modes. Our intuition is that synthetic datasets like ModelNet40 and ShapeNetPart contain complete, uniform, and relatively clean point clouds, such as the airplanes shown in Figures 3(a) and 3(c). In contrast, ScanObjectNN consists of real scanned point clouds with varying types of missing or noisy points, making up different sub-modes within the same class. For example, Figure 3(b) shows two kinds of chairs: one is mostly missing, while the other is more complete. As static prompt tuning fails to capture various sub-modes in real-world data

distribution, it highlights the necessity of an adaptive or dynamic prompting strategy.

3.3. Instance-aware Dynamic Prompt Tuning

To address the aforementioned issues of static prompt tuning, we propose **Instance-aware Dynamic Prompt Tuning (IDPT)**. Figure 2 shows the pipeline of IDPT.

3.3.1 Dynamic Prompt Generation Module

To capture various sub-modes existing in the real-world data and enhance the robustness against noises (*e.g.*, with different types of missing or noisy points), we utilized EdgeConv [45] at the patch level to perceive local point cloud shapes at a larger scale. Specifically, as shown in Figure 2, point patch tokens E_{N-1} are processed by three EdgeConvs to generate three patch features at different scales. Then, the multi-scale patch features are concatenated and fed into a linear layer, followed by max pooling to generate an instance-aware dynamic prompt P_{N-1} :

$$P_{N-1} = \varphi_P(E_{N-1}). \quad (8)$$

$\varphi_P(\cdot)$ denotes the dynamic prompt generation module.

Next, we forward P_{N-1} accompany with E_{N-1} to the last transformer layer f_N :

$$[c_N; P_N; E_N] = f_N([c_{N-1}; P_{N-1}; E_{N-1}]). \quad (9)$$

Finally, we concatenate [CLS] token c_N , prompt token P_N , and patch tokens E_N together before feeding them into a task head. The final prediction is made by:

$$y = g_h([c_N; P_N; E_N]). \quad (10)$$

3.3.2 Prompt Insert Position

Perceiving various sub-modes in real-world point cloud data relies on high-level semantic information. As higher (or deeper) transformer layers grasp global semantic information (*e.g.* density, shape, or categorical information) better, IDPT prefers to insert prompts at deeper transformer layers to ensure an accurate perception of point cloud semantics. In particular, we found that inserting the dynamic prompt before the last transformer layer yields robust performance and also enjoys decent efficiency. We provide detailed quantitative analysis in § 4.3.3.

4. Experiments

We evaluated the performance of the proposed approach on classification, few-shot learning, and segmentation tasks. We used three widely used pre-trained models, Point-Bert [51], Point-MAE [32], and ACT [11], as our baseline. Notably, our IDPT is a universal paradigm that can be applied to any pre-trained point cloud model.

4.1. Experiment Settings

To ensure comparison fairness, we have used the same experimental settings as the default fine-tuning method for each baseline. This involves freezing the weights of the pre-trained point cloud model and only updating the parameters of the Prompt Model and Head during downstream task training. All experiments were conducted on a single GeForce RTX 3090 24GB. We have explored the performance of the simple rotation augmentation from ACT [11] on the ScanObjectNN [41] dataset, which is denoted as † in our table. In the downstream task experiments, we report the best result of 10 repeated experiments with different random seeds. In the ablation study, we report average results.

4.2. Prompt Tuning in Downstream Tasks

4.2.1 Object Classification on Real-World Dataset

In the study of point cloud pre-training models, it is common practice to conduct pre-training on the ShapeNet [8] dataset, which typically only contains clean point clouds and assumes that all point clouds are identically distributed. However, in reality, point clouds often suffer from issues such as noise and missing points, resulting in a diverse distribution. We first assess our IDPT performance on the ScanObjectNN [41] dataset, which consists of about 15K point cloud samples by 15 categories. These objects are scanned indoor scene data, which are usually cluttered with background and occluded by other objects.

We conducted experiments on three variants of ScanObjectNN [41] (OBJ-BG, OBJ-ONLY, and PB-T50-RS). The results are shown in Table 2, † indicates that the pre-trained model used simple rotational augmentation of ACT during fine-tuning or prompt tuning, and without † indicates the default augmentation method. We observed that: (i) We achieved state-of-the-art (SOTA) performance with IDPT on Point-MAE†. In comparison to the current state-of-the-art method ACT, we have achieved gains of 0.34%, 1.21%, and 0.3% respectively in the three variants of ScanObjectNN, while utilizing only 7% of its trainable parameters. (ii) Our IDPT outperforms full fine-tuning in most cases with fewer trainable parameters. These results demonstrate the excellent performance of our method on real scanned point clouds with various data distributions. We believe this is due to the introduction of a semantic prior of real point cloud data on the one hand, and fewer trainable parameters to mitigate overfitting on the other.

4.2.2 Object Classification on Synthetic Dataset

We evaluate IDPT on the ModelNet40 [47] dataset for object classification. ModelNet40 [47] includes 12,311 clean 3D CAD models for 40 categories. Each point cloud is complete, uniform, and noise-free, and all point clouds in the

Method	#TP (M)	OBJ_BG(†)	OBJ_ONLY(†)	PB_T50_RS(†)
<i>Supervised Learning Only</i>				
PointNet [33]	3.5	73.3	79.2	68.0
PointNet++ [34]	1.5	82.3	84.3	77.9
DGCNN [45]	1.8	82.8	86.2	78.1
PointCNN [24]	0.6	86.1	85.5	78.5
BGA-DGCNN [41]	1.8	-	-	79.7
BGA-PN++ [41]	1.5	-	-	80.2
DRNet [36]	-	-	-	80.3
GBNet [37]	8.8	-	-	80.5
SimpleView [15]	-	-	-	80.8
PRANet [10]	2.3	-	-	81.0
MVTN [18]	-	-	-	82.8
PointMLP [31]	-	-	-	85.7
<i>with Self-Supervised Representation Learning (Full Fine-tuning)</i>				
Transformer [51]	22.1	79.86	80.55	77.24
OcCo [51]	22.1	84.85	85.54	78.79
Point-BERT [51]	22.1	87.43	88.12	83.07
MaskPoint [26]	22.1	89.70	89.30	84.60
Point-MAE [32]	22.1	90.02	88.29	85.18
Point-M2AE [53]	-	91.22	88.81	86.43
ACT [†] [11]	22.1	93.29	91.91	88.21
Point-MAE [†] [32]	22.1	92.94	92.08	88.41
<i>with Self-Supervised Representation Learning (IDPT)</i>				
Point-BERT w/ IDPT	1.7	88.12 († 0.69)	88.30 († 0.18)	83.69 († 0.62)
Point-MAE w/ IDPT	1.7	91.22 († 1.20)	90.02 († 1.73)	84.94 (↓ 0.24)
ACT[†] w/ IDPT	1.7	93.12 (↓ 0.17)	92.26 († 0.35)	87.65 (↓ 0.56)
Point-MAE[†] w/ IDPT	1.7	93.63 († 0.69)	93.12 († 1.04)	88.51 († 0.10)

Table 2. Classification results on three variants of ScanObjectNN dataset, and we report the number of trainable parameters (#TP) and classification accuracy(%). † indicates that the pre-trained model used simple rotational augmentation of ACT [11] during fine-tuning or prompt tuning. We achieve state-of-the-art performance with IDPT on Point-MAE[†] and our IDPT outperforms full fine-tuning in most cases with fewer trainable parameters.

dataset are independently and identically distributed. We follow standard protocols to split ModelNet40 into 9843 instances for the training set and 2468 for the testing set. Standard random scaling and random translation are applied for data augmentation during training. For fair comparisons, following previous work [51, 32, 11], we also use the standard voting method [30] during testing.

As shown in Table 3, Point-MAE with IDPT achieves state-of-the-art performance with an accuracy of 94.4%. This represents a 0.6% improvement compared to fine-tuning. Additionally, other pre-trained models with IDPT, such as Point-BERT and ACT, demonstrate certain improvements compared to full fine-tuning. These results suggest that the incorporation of semantic priors of each instance can yield significant improvements.

4.2.3 Few-shot Learning

We conducted few-shot experiments on ModelNet40, using the n-way, m-shot setting, following previous works [32, 51, 55]. The results for the settings of $n \in 5, 10$ and $m \in 10, 20$ are presented in Table 4. Our IDPT achieved performance gains in most cases compared to full fine-tuning, demonstrating its efficacy in few-shot learning.

Method	ST?	#TP (M)	Data Type	Accuracy (%)
<i>Supervised Learning Only</i>				
PointNet [33]	-	3.5	1k Points	89.2
PointNet++ [34]	-	1.5	1k Points	90.7
DGCNN [45]	-	1.8	1k Points	92.9
PCT [16]	N	2.9	1k Points	93.2
PVT [52]	N	-	1k Points	93.6
PointTransformer [52]	N	-	1k Points	93.7
MVTN [18]	-	11.2	12 Images	93.8
SimpleView [15]	-	-	6 Images	93.9
PointMLP [31]	-	14.9	1k Points	94.5
<i>with Self-Supervised Representation Learning (Full Fine-tuning)</i>				
Transformer [51]	Y	22.1	1k Points	91.4
OcCo [51]	Y	22.1	1k Points	92.1
EPCL [21]	-	-	1k Points	92.9
Point-BERT [51]	Y	22.1	1k Points	93.2
ACT [11]	Y	22.1	1k Points	93.7
Point-MAE [32]	Y	22.1	1k Points	93.8
MaskPoint [26]	Y	22.1	1k Points	93.8
Point-M2AE [53]	N	15.3	1k Points	94.0
CLIP2Point [20]	-	-	10 Images	94.0
P2P [46]	-	1.2	1 Images	94.0
<i>with Self-Supervised Representation Learning (IDPT)</i>				
Point-BERT w/ IDPT	Y	1.7	1k Points	93.4 († 0.2)
ACT w/ IDPT	Y	1.7	1k Points	94.0 († 0.3)
Point-MAE w/ IDPT	Y	1.7	1k Points	94.4 († 0.6)

Table 3. Classification results on ModelNet40 [47] dataset. ‘ST’ indicates whether the backbone is a standard Transformer [43] without any special design or inductive bias. ‘1k Points’ indicates that the input data is 1k points and ‘n Images’ indicates that the input data is n images. Our IDPT outperforms **full fine-tuning** in each baseline with fewer trainable parameters.

	5-way		10-way	
	10-shot	20-shot	10-shot	20-shot
<i>with Self-Supervised Representation Learning (Full Fine-tuning)</i>				
DGCNN-OcCo [44]	90.6±2.8	92.5±1.9	82.9±1.3	86.5±2.2
Transformer-OcCo [51]	94.0±3.6	95.9±2.3	89.4±5.1	92.4±4.6
Point-BERT [51]	94.6±3.1	96.3±2.7	91.0±5.4	92.7±5.1
MaskPoint [26]	95.0±3.7	97.2±1.7	91.4±4.0	93.4±3.5
EPCL [21]	95.1±2.7	97.3±1.6	91.1±4.2	93.5±3.8
Point-MAE [32]	96.3±2.5	97.8±1.8	92.6±4.1	95.0±3.0
Point-M2AE [53]	96.8±1.8	98.3±1.4	92.3±4.5	95.0±3.0
ACT [11]	96.8±2.3	98.0±1.4	93.3±4.0	95.6±2.8
<i>with Self-Supervised Representation Learning (IDPT)</i>				
Point-BERT w/ IDPT	96.0±1.7 (†)	97.2±2.6 (†)	91.9±4.4 (†)	93.6±3.5 (†)
Point-MAE w/ IDPT	97.3±2.1 (†)	97.9±1.1 (†)	92.8±4.1 (†)	95.4±2.9 (†)
ACT w/ IDPT	96.7±2.5 (↓)	98.2±0.9 (†)	92.4±4.5 (↓)	95.5±3.0 (↓)

Table 4. Few-shot learning on ModelNet40. We report the average classification accuracy (%) with the standard deviation (%) of 10 independent experiments. Our IDPT achieved performance gains in most cases compared to full fine-tuning in few-shot learning.

4.2.4 Part Segmentation

For part segmentation, we follow previous work [11, 32] to add prompts to the input of 3-rd, 7-th and 11-th layers and the task head. Since we empirically observed that using a single-layer MLP achieves comparable performance to the three-layer EdgeConv architecture in the segmentation task, we adopt a simple single-layer MLP as the dynamic prompt

Methods	#TP (M)	mIoU _c	mIoU _I
<i>Supervised Learning Only</i>			
PointNet [33]	-	80.39	83.7
PointNet++ [34]	-	81.85	85.1
DGCNN [45]	-	82.33	85.2
<i>with Self-Supervised Representation Learning (Full Fine-tuning)</i>			
Transformer [51]	27.09	83.42	85.1
OcCo [51]	27.09	83.42	85.1
MaskPoint [26]	-	84.60	86.0
Point-BERT [51]	27.09	84.11	85.6
Point-MAE [32]	27.06	84.19	86.1
ACT [11]	27.06	84.66	86.1
<i>with Self-Supervised Representation Learning (Prompt Tuning)</i>			
Point-MAE w/ VPT	5.35	83.64	85.4
Point-MAE w/ IDPT	5.69	83.79	85.7
ACT w/ VPT	5.35	83.48	85.4
ACT w/ IDPT	5.69	83.82	85.9

Table 5. Part segmentation results on the ShapeNetPart dataset. The mean IoU across all categories, *i.e.*, mIoU_c (%), and the mean IoU across all instances, *i.e.*, mIoU_I (%) are reported.

generation module at each layer to reduce the number of trainable parameters.

According to experimental results in Table 5, our IDPT outperforms the static prompting strategy, VPT [22]. It verifies the effectiveness of our dynamic prompting strategy in part segmentation. But they are still inferior to full fine-tuning. We attribute the performance gap to the difficulty of fine-grained understanding of point clouds, which makes it challenging to transfer pre-trained backbones with limited tunable parameters to the segmentation task. Fortunately, the design of instance-aware dynamics in IDPT helps to mitigate such a gap. We believe developing effective structure modeling mechanisms in the parameter-efficient tuning strategy is a promising direction for fine-grained point cloud tasks.

4.3. Ablation Study

To investigate the architecture design and tuning settings of our proposed strategy, we conducted extensive ablation studies on classification tasks in 2 variants of ScanObjectNN [41] (OBJ_BG and OBJ_ONLY).

4.3.1 Compare IDPT with Other Tuning Strategies

In order to demonstrate the superiority of our proposed instance-aware dynamic prompt tuning over other tuning strategies (Head tuning, VPT [22] and full fine-tuning), we conducted extensive ablation studies as shown in Table 1. The specific tuning strategies used are as follows: (A) Head tuning (where freezes the backbone and serves as a *true* reference for “whether prompt tuning improves performance”), (B) VPT-Shallow [22], (C) VPT-Deep [22], (D)

Prompt Strategy	Trainable Parameters Type	Tr. Param.	OBJ_BG	OBJ_ONLY
w/o prompt	Head	0.27	87.40	87.13
VPT-Deep	10 prompts + Head	0.36	89.98	90.19
VPT-Deep	28 prompts + Head	0.52	90.02	90.53
VPT-Deep	156 prompts + Head	1.70	90.19	90.53
IDPT	PM (1-layer MLP) + Head	0.52	91.43	90.98
IDPT	PM (3-layer MLPs) + Head	1.49	91.64	91.34
IDPT	PM (1 EdgeConv) + Head	0.81	91.77	91.67
IDPT	PM (2 EdgeConvs) + Head	1.25	91.95	91.67
IDPT	PM (3 EdgeConvs) + Head	1.70	92.48	92.19
IDPT	PM (1 Transformer layer) + Head	2.14	92.03	91.22

Table 6. Effects of the number of trainable parameters and the structure of prompt generation module.

our IDPT, and (E) Fine-tuning.

Our IDPT tuning strategy (D) has outperformed other tuning strategies by achieving the highest level of performance while utilizing fewer trainable parameters. When compared to the baseline (A), we observed a **5.08%**, **5.06%**, and **8.18%** improvement on the three variants of ScanObjectNN, respectively. Our strategy also significantly outperformed the traditional static prompt (B and C). Additionally, our method demonstrated a significant improvement over fine-tuning (E), as it greatly reduced the number of trainable parameters while still improving performance. Overall, these experiments clearly demonstrate that our IDPT tuning strategy is highly effective.

4.3.2 Amount of Trainable Parameters and Structure

We conducted experiments with varying numbers of trainable parameters to evaluate the effectiveness of our models. In VPT-Deep, we adjusted the trainable parameters by controlling the number of prompts in each layer input. Meanwhile, in IDPT, we experimented with various network structures, including MLP, graph convolution (EdgeConv), and Transformer layers. Our experimental results, as shown in Table 6, demonstrate the following: (i) Using a single-layer MLP in IDPT resulted in significant improvements of 4.03% and 3.85% in OBJ_BG and OBJ_ONLY, respectively, when compared to the baseline without prompts. These results suggest that introducing semantic priors from the downstream task data can be highly effective. (ii) Increasing the number of parameters in VPT resulted in only limited performance improvements. It shows that merely increasing the number of parameters without introducing a semantic prior has limited performance improvement. (iii) For IDPT, EdgeConv proved to be an effective prompt module. Compared to MLP and Transformer, EdgeConv focuses more on local neighborhood information, allowing it to better perceive the specific shape of the point cloud and provide a better representation for downstream data semantic priors.

4.3.3 Prompt Insert Position

We conducted an analysis to observe the impact of integrating our dynamic prompt in various depths of Transformer




Figure 4. Effect of prompt insert position. E_i means to insert our dynamic prompt into the input of i -th layer of the pre-trained Transformer. 'All Layers' means to insert our dynamic prompt in each layer of the Transformer using a prompt module with shared parameters.

layers. The experimental results, as demonstrated in Figure 4 (a) and (b), indicated that the addition of a prompt to deeper layers resulted in better performance. This can be attributed to the fact that our prompt generation module utilizes patch tokens to perceive the semantic priors of each point cloud, and deeper patch tokens comprise more comprehensive semantic information.

Additionally, we found that applying the prompt to all layers using a prompt module with shared parameters in the "all_layer" setting did not produce satisfactory results. This is due to the fact that the shared prompt module can cause degradation in the semantic prior representation, and using independent parameters for each layer would result in an unacceptable increase in the number of parameters. As a result, we decided to add the prompt to the input of the final layer of the Transformer.

4.3.4 Qualitative Analysis of the Ability to Approximate $\Phi(\cdot)$ in Downstream Adaptation

We analyzed the effectiveness of various tuning strategies for the transformation function $\Phi(\cdot)$ by conducting a qualitative analysis of their fitting capability. The strategies we evaluated included (a) the pre-trained model, (b) VPT-Deep, and (c) IDPT. To visualize the input (E_{N-1}) and output (E_N) of all point patches in the N -th layer of the Transformer on ScanObjectNN, we utilized a pre-trained model based on Point-MAE. The results of our visualization are presented in Figure 5, which displays the visualization outcomes of the three tuning strategies.

The performance of different strategies for approximating the function $\Phi(\cdot)$ with the Transformer model varies. The pre-trained model (a) uses the Transformer with fixed parameters and shows the worst performance. VPT-Deep (b), which adds trainable static prompt parameters to all layers' inputs to approximate $\Phi(\cdot)$, resulting in a better input feature distribution E_{N-1} in the N -th layer than (a). However, the output features E_N are still scattered, indicating a poorer $\Phi(\cdot)$ approximate. IDPT takes a different approach by introducing the semantic prior of each instance and adding dynamic prompts to the input space of the N -

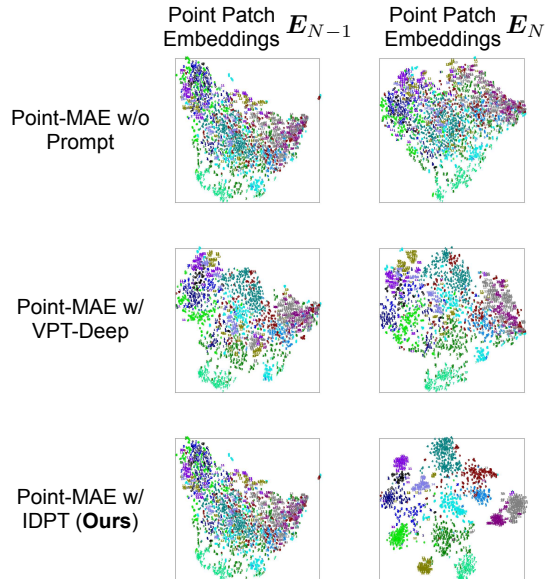


Figure 5. The t-SNE visualization of the point patch features extracted E_{N-1} and E_N from the test sets of ScanObjectNN (PB_T50_RS) using a pre-trained Point-MAE with different tuning strategies. This visualization partly reflects the approximation of the transformation function $\Phi(\cdot)$ by different tuning strategies.

th layer to approximate $\Phi(\cdot)$. Although the input features E_{N-1} at the N -th layer are as scattered as (a), the output features E_N are tightly clustered for the same category after concatenating E_{N-1} with our prompt through the same network as (a). It indicates that our strategy can effectively align the different distributions and is the best approximate strategy for $\Phi(\cdot)$.

5. Conclusion

In this paper, we investigate prompt tuning on pre-trained point cloud models to pursue the balance between performance and parameter efficiency. We found that the popular visual prompt tuning strategy cannot work well in real point clouds due to the distribution diversity. Therefore, we proposed instance-aware dynamic prompt tuning (IDPT) with a prompt generation mechanism to enhance the model's robustness in downstream transfer. Extensive experiments validated IDPT as a universal and effective solution.

6. Acknowledge

This work is supported in part by the National Key Research and Development Program of China, under Grant 2022YFF1202104, National Natural Science Foundation of China, under Grant (6230070671,62171248), Shenzhen Science and Technology Program (Grant No. RCYX20200714114523079, JCYJ20220818101014030, JCYJ20220818101012025), and the PCNL KEY project (PCL2021A07).

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
- [2] Armen Aghajanyan, Luke Zettlemoyer, and Sonal Gupta. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. *arXiv preprint arXiv:2012.13255*, 2020. 4
- [3] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. *arXiv preprint arXiv:2202.03555*, 2022. 2
- [4] Suzanna Becker and Geoffrey E Hinton. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(6356), 1992. 2
- [5] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79, 2010. 4
- [6] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *NeurIPS*, 19, 2006. 4
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33, 2020. 2
- [8] 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
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2
- [10] Silin Cheng, Xiwu Chen, Xinwei He, Zhe Liu, and Xiang Bai. Pra-net: Point relation-aware network for 3d point cloud analysis. *IEEE TIP*, 30, 2021. 6
- [11] Runpei Dong, Zekun Qi, Linfeng Zhang, Junbo Zhang, Jianjian Sun, Zheng Ge, Li Yi, and Kaisheng Ma. Autoencoders as cross-modal teachers: Can pretrained 2d image transformers help 3d representation learning? *arXiv preprint arXiv:2212.08320*, 2022. 1, 2, 4, 5, 6, 7
- [12] 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
- [13] Tianyu Gao, Adam Fisch, and Danqi Chen. Making pretrained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*, 2020. 2
- [14] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. 2
- [15] Ankit Goyal, Hei Law, Bowei Liu, Alejandro Newell, and Jia Deng. Revisiting point cloud shape classification with a simple and effective baseline. In *ICML*, 2021. 6
- [16] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7(2), 2021. 1, 6
- [17] Xu Guo and Han Yu. On the domain adaptation and generalization of pretrained language models: A survey. *arXiv preprint arXiv:2211.03154*, 2022. 4
- [18] Abdullah Hamdi, Silvio Giancola, and Bernard Ghanem. Mvtn: Multi-view transformation network for 3d shape recognition. In *ICCV*, 2021. 6
- [19] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 2
- [20] Tianyu Huang, Bowen Dong, Yunhan Yang, Xiaoshui Huang, Rynson WH Lau, Wanli Ouyang, and Wangmeng Zuo. Clip2point: Transfer clip to point cloud classification with image-depth pre-training. *arXiv preprint arXiv:2210.01055*, 2022. 1, 2, 6
- [21] Xiaoshui Huang, Sheng Li, Wentao Qu, Tong He, Yifan Zuo, and Wanli Ouyang. Frozen clip model is efficient point cloud backbone. *arXiv preprint arXiv:2212.04098*, 2022. 1, 6
- [22] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, 2022. 1, 2, 3, 4, 7
- [23] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 2
- [24] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *NeurIPS*, 31, 2018. 1, 6
- [25] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *CVPR*, 2021. 2
- [26] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. In *ECCV*, 2022. 1, 2, 6, 7
- [27] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM CSUR*, 55(9), 2023. 2
- [28] Xiao Liu, Kaixuan Ji, Yicheng Fu, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*, 2021. 2
- [29] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. Gpt understands, too. *arXiv preprint arXiv:2103.10385*, 2021. 2
- [30] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *CVPR*, 2019. 6
- [31] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point cloud: A simple residual mlp framework. *arXiv preprint arXiv:2202.07123*, 2022. 1, 6

- [32] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. *arXiv preprint arXiv:2203.06604*, 2022. 1, 2, 3, 4, 5, 6, 7
- [33] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 1, 6, 7
- [34] 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. 1, 6, 7
- [35] Zekun Qi, Runpei Dong, Guofan Fan, Zheng Ge, Xiangyu Zhang, Kaisheng Ma, and Li Yi. Contrast with reconstruct: Contrastive 3d representation learning guided by generative pretraining. In *ICML*, 2023. 1
- [36] Shi Qiu, Saeed Anwar, and Nick Barnes. Dense-resolution network for point cloud classification and segmentation. In *WACV*, 2021. 6
- [37] Shi Qiu, Saeed Anwar, and Nick Barnes. Geometric back-projection network for point cloud classification. *IEEE TMM*, 24, 2021. 6
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 2
- [39] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *CVPR*, 2022. 2
- [40] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *NeurIPS*, 34, 2021. 2
- [41] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *ICCV*, 2019. 2, 4, 5, 6, 7
- [42] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 9(11), 2008. 4
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 6
- [44] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *ICCV*, 2021. 6
- [45] 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 ToG*, 38(5), 2019. 1, 3, 5, 6, 7
- [46] Ziyi Wang, Xumin Yu, Yongming Rao, Jie Zhou, and Jiwen Lu. P2p: Tuning pre-trained image models for point cloud analysis with point-to-pixel prompting. *arXiv preprint arXiv:2208.02812*, 2022. 1, 2, 6
- [47] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015. 2, 4, 5, 6
- [48] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 2
- [49] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *ECCV*, 2020. 2
- [50] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM ToG*, 35(6), 2016. 4
- [51] 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. 1, 2, 5, 6, 7
- [52] Cheng Zhang, Haocheng Wan, Shengqiang Liu, Xinyi Shen, and Zizhao Wu. Pvt: Point-voxel transformer for 3d deep learning. *arXiv preprint arXiv:2108.06076*, 2, 2021. 6
- [53] 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. 6
- [54] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *CVPR*, 2022. 1, 2
- [55] Yabin Zhang, Jiehong Lin, Ruihuang Li, Kui Jia, and Lei Zhang. Point-dae: Denoising autoencoders for self-supervised point cloud learning. *arXiv preprint arXiv:2211.06841*, 2022. 1, 2, 6
- [56] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, 2022. 2
- [57] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 130(9), 2022. 2, 3