

# Dynamic Adapter Meets Prompt Tuning: Parameter-Efficient Transfer Learning for Point Cloud Analysis

Xin Zhou<sup>\*1</sup>, Dingkang Liang<sup>\*1</sup>, Wei Xu<sup>1</sup>, Xingkui Zhu<sup>1</sup>, Yihan Xu<sup>1</sup>, Zhikang Zou<sup>2</sup>, Xiang Bai<sup>†1</sup>  
<sup>1</sup>Huazhong University of Science and Technology, {xzhou03, dkliang, xbai}@hust.edu.cn  
<sup>2</sup>Baidu Inc., China

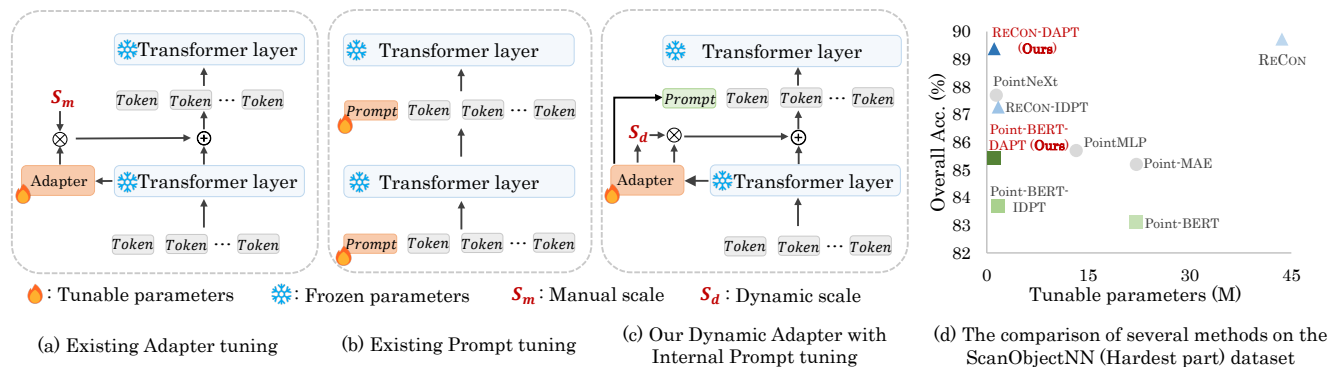


Figure 1. (a) Adapter tuning utilizes additional residual blocks with manual scale. (b) Prompt tuning usually introduces extra randomly initialized prompts into the input space. (c) The proposed Dynamic Adapter generates a dynamic scale for each token and seamlessly integrates with Prompt tuning. (d) The proposed method achieves an ideal trade-off between tunable parameters and performance.

## Abstract

Point cloud analysis has achieved outstanding performance by transferring point cloud pre-trained models. However, existing methods for model adaptation usually update all model parameters, i.e., full fine-tuning paradigm, which is inefficient as it relies on high computational costs (e.g., training GPU memory) and massive storage space. In this paper, we aim to study parameter-efficient transfer learning for point cloud analysis with an ideal trade-off between task performance and parameter efficiency. To achieve this goal, we freeze the parameters of the default pre-trained models and then propose the Dynamic Adapter, which generates a dynamic scale for each token, considering the token significance to the downstream task. We further seamlessly integrate Dynamic Adapter with Prompt Tuning (DAPT) by constructing Internal Prompts, capturing the instance-specific features for interaction. Extensive experiments conducted on five challenging datasets demonstrate that the proposed DAPT achieves superior performance compared to the full fine-tuning counterparts while significantly reducing the trainable parameters and training GPU memory by 95% and 35%, respectively. Code is available at <https://github.com/LMD0311/DAPT>.

## 1. Introduction

3D point cloud analysis, a crucial computer vision task, has a broad range of applications, such as autonomous driving [6, 24, 30] and 3D reconstruction [3, 10, 34]. The intrinsic irregularity and sparsity of point clouds pose significant challenges. In the past few years, various deep learning methods [7, 22, 38] have appeared, yielding substantial advancements in this domain.

To generate robust 3D general representations, using self-supervised learning with vision transformers [9] to pre-train models is a recent hot topic, such as mask modeling [13, 35] and contrast learning [39, 40]. Further fine-tuning the entire parameters of the pre-trained models on downstream datasets usually leads to faster convergence and significant performance improvements. However, such a fine-tuning strategy is sub-optimal for point cloud analysis, as reflected in the following aspects: 1) Fine-tuning the entire backbone might suffer catastrophic forgetting and break the rich prior knowledge embedded in the pre-training. 2) Fine-tuning for each point cloud analysis dataset requires a separate weights copy, and the storage space overhead may become a burden as the number of datasets increases. 3) The computational cost requirements escalate dramatically, especially for larger batch sizes, leading to a substantial increase in GPU memory usage, limiting

\* Equal contribution. † Corresponding author.

its accessibility for researchers with weak hardware.

To alleviate these problems, we turn our attention to Parameter-Efficient Transfer Learning (PETL), an alternative fine-tuning strategy that has achieved marked progress in the Natural Language Processing (NLP) domain. Current PETL methods seek to exploit the representational prior in large language models (LLM) by fixing most of the parameters and adjusting only a selected few. Two pioneering PETL paradigms of NLP are Adapter tuning [12, 16, 17] and Prompt tuning [23, 25], where the former usually introduces a lightweight network for each Attention/FFN block (Fig. 1(a)) and the latter inserts some external learnable parameters called prompts to the input of each transformer layer (Fig. 1(b)). Both paradigms can obtain performance on par with or surpass the full fine-tuning counterpart while significantly reducing the trainable parameters.

However, we empirically find that directly using the Adapter/Prompt tuning cannot achieve satisfactory results in point cloud analysis. We argue the main reasons are twofold: 1) One key aspect of Adapter tuning is the scale, which is used to adjust the features to match the downstream tasks better, as shown in Fig. 1(a). Existing methods [4, 12, 19, 31] often require manual scale setting as a crucial hyper-parameter, while the value remains constant during inference. This static scale may struggle to adapt to point clouds with complex geometry structures and non-uniform distribution. As a result, such challenges demand a more dynamic and adaptable tuning approach. Simply applying a static scale overlooks the unique characters of each point cloud. 2) The Prompt tuning usually adds external random initialized prompts that often lack relevance to the point clouds, as shown in Fig. 1(b). This misalignment makes the network optimization process challenging.

Considering these reasons, in this paper, we propose the **Dynamic Adapter** and seamlessly integrate it with **Prompt Tuning**, namely ‘DAPT’. As shown in Fig. 1(c), we introduce an extremely lightweight module to generate a unique scale for each token for Adapter tuning. Such an operator can make sure the Adapter dynamically adjusts each token by considering the significance score and matches the complex 3D samples better. Besides, we propose **Internal Prompt tuning**, which utilizes the Dynamic Adapter to produce a series of prompts instead of using extra random initialized prompts. Through the Internal Prompt tuning, each prompt will better capture the instance-specific features and be easy to optimize, helping the model capture the global perspective of point clouds.

Extensive experiments on diverse point cloud datasets in different settings demonstrate the effectiveness of our method. In particular, as shown in Fig. 1(d), our DAPT can be simply merged into the state-of-the-art pre-trained models (e.g., Point-BERT [52], RECON [39]) and can significantly reduce the trainable parameters by 95% and save

GPU memory by up to 35% compared with the full fine-tuning paradigm, while achieving similar or even higher performance, e.g., 2.36% accuracy improvement on Point-BERT under ScanObjectNN PB\_50\_RS variant.

Our major contributions can be summarized as follows:

- We reveal the limitation of existing parameter-efficient transfer learning (Adapter tuning and Prompt tuning) of NLP in point cloud analysis.
- We propose the Dynamic Adapter, which generates a unique scale for each token to dynamically adjust features by considering the significance score.
- Based on the Dynamic Adapter, we introduce the Internal Prompt tuning, where we use the Dynamic Adapter to construct the prompt instead of using extra random initialized prompts.

## 2. Related Work

### 2.1. Pre-training on Point Cloud

Transformers have been widely used in 3D vision [2, 15, 49, 56, 58], and pre-training on 3D datasets is receiving significant interest. There are two main pretext tasks paradigms for 3D pre-training, including contrastive learning [1, 48] and mask modeling [27, 28, 52]. Contrastive learning-based methods usually contrast between the different views or instances, such as PointContrast [48] and CrossPoint [1]. Besides, Poursaeed et al. [36] defined the estimation of rotation angles as an auxiliary task for contrastive learning. The mask modeling [35, 55, 57] typically relies on autoencoders to learn the latent features of the data by reconstructing the original input. For example, Point-BERT [52] generates discrete point tokens containing meaningful local information from masked tokens. Point-MAE [35] uses an autoencoder to learn high-level latent features from unmasked patches, aiming to reconstruct the masked point patches. ACT [8] employs an autoencoder as the cross-modal teacher to guide the 3D representation learning. RECON and RECON++ [39, 40] learns through ensemble distillation from both generative modeling teachers and single/cross-modal contrastive teachers.

The above methods usually fine-tune the entire backbone to adopt the pre-trained model to the downstream 3D tasks effectively, achieving promising performance. However, full fine-tuning is inefficient and may break the well-trained knowledge. In this paper, our focus is on efficiently transferring pre-trained 3D models to downstream tasks.

### 2.2. Parameter-Efficient Transfer Learning

The continuous expansion of pre-trained models demands significant computational resources and consumes considerable storage during fine-tuning. To address these challenges, researchers in the NLP and 2D computer vision domain have explored PETL methods:

Prompt tuning [18, 23] usually adds extra information to the model by introducing latent tokens (prompts) to the task input, enhancing the model’s behavior during fine-tuning. Building on prompt tuning, Prefix tuning [25] concatenates tunable prefix vectors to the keys and values of the Attention at every layer. Adapter tuning methods [4, 16] insert lightweight modules named Adapter into the FFN in a complementary approach, allowing for incremental model refinements. LoRA [17] further adopts a low-rank approximation to update weight matrices in the Attention. Building on these foundations, several variations have been proposed, including varying the placement of Adapters [12, 26, 59], empowering the network to select the best method combination [33], and implementing strategies to minimize the trainable parameters [19–21, 51].

So far, only one method, IDPT [54] focuses on parameter-efficient transfer learning on point cloud analysis. IDPT extends the Prompt tuning with a DGCNN [45] to extract instance-aware prompts for model fine-tuning instead of using static prompts. Unlike IDPT, we propose the Dynamic Adapter and seamlessly integrate it with Prompt Tuning, which significantly reduces the tunable parameters and achieves impressive performance.

### 3. Preliminary

In this section, we revisit the vision transformer and introduce Adapter tuning and Prompt tuning as preliminary.

#### 3.1. Vision Transformer

In a vision transformer (ViT) [9], an image is divided into  $N$  non-overlapping patches and transferred to a 1D sequence. The sequence is further processed by  $L$ -layer transformer blocks, along with a classification token. The input  $\mathbf{x} \in \mathbb{R}^{(N+1) \times d}$ , where  $d$  is the embedding dimension, is first transformed to queries  $\mathbf{Q} \in \mathbb{R}^{(N+1) \times d}$ , keys  $\mathbf{K} \in \mathbb{R}^{(N+1) \times d}$  and values  $\mathbf{V} \in \mathbb{R}^{(N+1) \times d}$ . After that, we can calculate the self-attention with an attention layer.

The output of the attention layer will be sent to an FFN module with residual bypasses to extract information in channels. The transformer block can be written as below:

$$\begin{aligned} \mathbf{x}'_i &= \text{Attention}(\text{LN}(\mathbf{x}_{i-1})) + \mathbf{x}_{i-1}, \\ \mathbf{x}_i &= \text{FFN}(\text{LN}(\mathbf{x}'_i)) + \mathbf{x}'_i, \end{aligned} \quad (1)$$

where  $\mathbf{x}_i$  is the output of  $i$ -th block, LN is layer norm. Recently, some work [35, 52] has adapted the ViT for point cloud analysis and utilized self-supervision for pre-training. In ViT, the Attention and FFN components contain most parameters. It would be reasonable to fix them for PETL, but doing so would damage performance.

#### 3.2. Parameter-Efficient Fine-tuning

**Adapter tuning** [4, 12, 16] is a module that acts as a bottleneck and inserts a few parameters into the transformer. It

Table 1. The overall accuracy (%) for Adapter and Prompt tuning strategies on three variants of ScanObjectNN [43] is reported. ‘#TP’ means the number of tunable parameters. Linear probing indicates head-tuned only.

Tuning Strategy	#TP(M)	OBJ_BG	OBJ_ONLY	PB_T50_RS
Point-MAE [35]	22.1	90.02	88.29	85.18
Linear probing	0.3	87.26(-2.76)	84.85(-3.44)	75.99(-9.19)
+ Adapter [16]	0.9	89.50(-0.52)	88.64(+0.35)	83.93(-1.25)
+ VPT [18]	0.4	87.26(-2.76)	87.09(-1.20)	81.09(-4.09)

includes a downward projection to decrease the feature dimension, a non-linear function, and an upward projection to project back to the original dimension. Specifically, given the input as  $\mathbf{x} \in \mathbb{R}^{(N+1) \times d}$ , the output is calculated by

$$\text{Output} = S_m \times [(\phi(\mathbf{x}W_d^T))W_u^T], \quad (2)$$

where  $W_d \in \mathbb{R}^{r \times d}$ ,  $\phi(\cdot)$ ,  $W_u \in \mathbb{R}^{d \times r}$  ( $r \ll d$ ), and  $S_m$  represent the down-projection weight, non-linear function, up-projection weight, and scale, respectively. Note that  $S_m$  is a key hyper-parameter that needs to be manually set in existing methods [4, 12, 19, 31].

**Prompt tuning** [18, 23, 25] creates random initialized tokens as prompts. These prompts are then incorporated into the input of a transformer block or attention layer, where they interact with the original tokens through self-attention. In the fine-tuning process, the weights of the backbone network remain frozen, and only the prompts are updated. Given classification token  $\mathbf{T}_{cls} \in \mathbb{R}^{1 \times d}$  and patch tokens  $\mathbf{T} \in \mathbb{R}^{N \times d}$ , we can represent the inserted prompts as  $\mathbf{P}_i = \{\mathbf{P}_i^k \in \mathbb{R}^{1 \times d} \mid k \in \mathbb{N}, 1 \leq k \leq n\}$ , where  $n$  is the number of prompts. The output of  $i$ -th layer ( $L_i(\cdot)$ )  $\mathbf{x}_i \in \mathbb{R}^{(N+n+1) \times d}$  are obtained by

$$\mathbf{x}_i = L_i([\mathbf{T}_{cls}; \mathbf{P}_i; \mathbf{T}]). \quad (3)$$

However, while these tuning strategies achieve promising results in NLP and 2D vision, they lack targeted design for 3D. As shown in Tab. 1, when it comes to the hardest (*i.e.*, PB\_T50\_RS) variant of ScanObjectNN [43], Adapter and VPT have an accuracy gap of 1.73% and 4.09% when compared to full fine-tuning, respectively. Thus, designing a parameter-efficient fine-tuning method that performs well on 3D data is meaningful.

## 4. Methodology

Fig. 2(a) illustrates an overview of the proposed DAPT, which consists of three key components: a Task-agnostic Feature Transform Strategy (TFTS) for feature regulation, a Dynamic Adapter for feature modification, and an Internal Prompt selected from the output of Dynamic Adapter. The specifics of our DAPT will be discussed below.

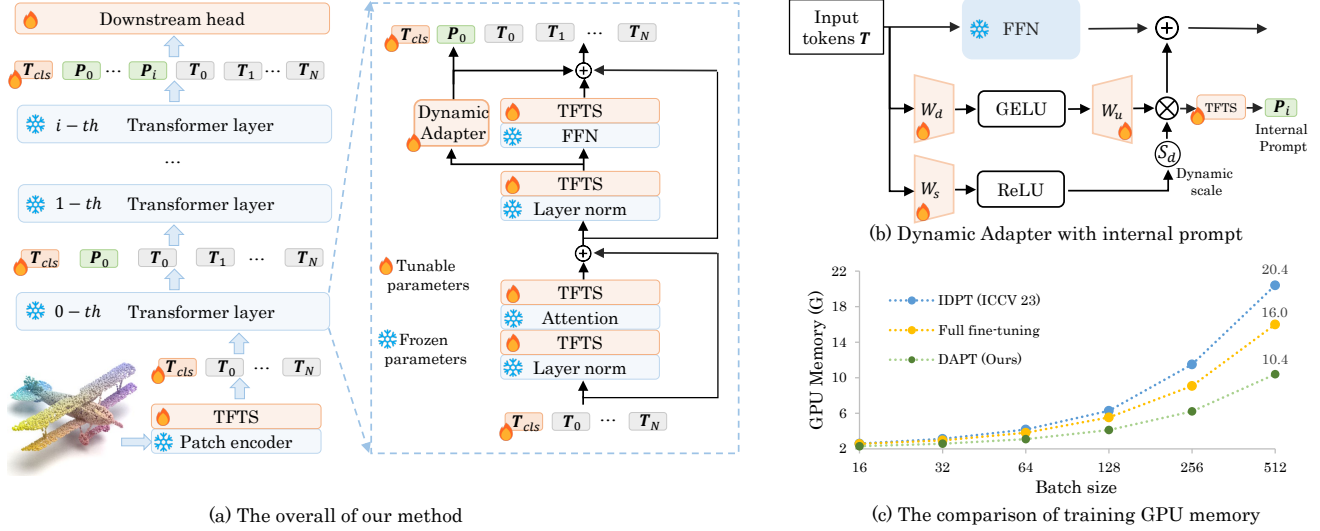


Figure 2. (a) The pipeline of our DAPT. During training, we fix the entire backbone, only fine-tuning the newly added parameters. (b) The detail of the Dynamic Adapter with Internal Prompt. (c) Comparison of GPU memory usage among full fine-tuning, IDPT [54] and ours.

#### 4.1. Task-agnostic Feature Transform Strategy

3D pre-training models, such as masked autoencoders [35, 57] or contrastive learning [1, 39], have demonstrated their effectiveness in acquiring a generic feature representation space. However, these representations are often task-agnostic, thus requiring fine-tuning for downstream tasks. Hence, we introduce a frustratingly simple Task-agnostic Feature Transform Strategy (TFTS), adding two small learnable factors to transform the model’s features [26]. Given the input  $\mathbf{x} \in \mathbb{R}^{(N+1) \times d}$ , the output  $\mathbf{y} \in \mathbb{R}^{(N+1) \times d}$  is calculated as follows:

$$\mathbf{y} = \gamma \times \mathbf{x} + \beta, \quad (4)$$

where  $\gamma, \beta \in \mathbb{R}^{1 \times d}$  are learnable factors. Such a simple strategy effectively converts features to match downstream representations of 3D data, resulting in considerable performance improvements while only introducing an additional 0.5% parameters to the baselines.

#### 4.2. Dynamic Adapter

Adapter tuning [4, 12, 16] from NLP and 2D vision usually involves manual scale, significantly affecting the adapted model’s performance. Furthermore, this scale must be adjusted manually as a hyper-parameter for each specific task, which is time-consuming and inconvenient. To address this problem, we proposed the Dynamic Adapter, comprising two projection layers: a downward projection with parameters  $\mathbf{W}_d \in \mathbb{R}^{r \times d}$  and an upward projection with parameters  $\mathbf{W}_u \in \mathbb{R}^{d \times r}$ , as shown in Fig. 2(b).

The two projection layers are connected by a GELU function. We adopt a parallel MLP to generate dynamic scale  $S_d$  based on variant point cloud features. Specifically,

we utilize a scoring weight matrix  $\mathbf{W}_s \in \mathbb{R}^{1 \times d}$  to calculate a scale  $S_d$  for each token. With  $\mathbf{x}$  as the input, our dynamic scale can be expressed as

$$S_d = \text{ReLU}(\mathbf{x}\mathbf{W}_s^T). \quad (5)$$

Note that we utilize the ReLU to select the positive scale and set the rest as zero, as we think only significant tokens need to be adjusted, and it should depend on the unique characters of each point cloud. In general, the output of our Dynamic Adapter  $\mathbf{x}'$  can be described as

$$\mathbf{x}' = S_d \times [(\text{GELU}(\mathbf{x}\mathbf{W}_d^T))\mathbf{W}_u^T]. \quad (6)$$

To mitigate the influence of Adapter outputs during the initial stages of model training, we initialize  $\mathbf{W}_u$  to zero.

#### 4.3. Internal Prompt Tuning

In this subsection, we introduce how to combine our Dynamic Adapter with Prompt tuning. The existing Prompt tuning methods [18, 23, 25] usually adopt external random initialized embedding (prompts) to each layer input of the Transformer. Such random initialization cannot match well to the point cloud. In addition, during the testing phase, the prompts are static, i.e., they cannot generate instance-specific representations for each point cloud.

To this end, we leverage the Dynamic Adapter to generate the prompt derived from the original model’s internal output using the proposed Dynamic Adapter, called Internal Prompt tuning. Specifically, Dynamic Adapter operates on raw tokens, allowing our prompt to incorporate the model’s pre-training knowledge effectively. In practice, we empirically average pool the output of Dynamic Adapter after applying a GELU activation. Then, we utilized TFTS to obtain a Prompt and concatenated it between the classification



---

**Algorithm 1** Dynamic Adapter and Internal Prompt.

---

```
# An example of Dynamic Adapter and Internal Prompt
import torch.nn as nn
import torch.nn.functional as F

class DynamicAdapter():
    def __init__(self, dim, r):
        super().__init__()
        self.norm = nn.LayerNorm(dim)
        self.scale = nn.Linear(dim, 1)
        self.down_proj = nn.Linear(dim, r)
        self.up_proj = nn.Linear(r, dim)
        self.active = nn.GELU()
        with torch.no_grad():
            nn.init.zeros_(self.up_proj.weight)
            nn.init.zeros_(self.up_proj.bias)
    def forward(self, x):
        # Dynamic Adapter
        x = self.norm(x)
        dynamic_scale = F.relu(self.scale(x))
        out = self.down_proj(x)
        out = self.active(out)
        out = self.up_proj(out)
        out = out * dynamic_scale
        # Generate Internal Prompt
        internal_prompt = self.active(out).mean(dim=1)
        return out, internal_prompt
```

---

token  $T_{cls} \in \mathbb{R}^{1 \times d}$  and patch tokens  $T \in \mathbb{R}^{N \times d}$ . The output of the  $i$ -th transformer layer  $x_i$  can be expressed as

$$x_i = L_i([T_{cls}; P_0, \dots, P_{i-1}; T]). \quad (7)$$

Our Dynamic Adapter adds a residual embedding ( $\Delta x_k$ ) for each token ( $x_k$ ), acting as a local perspective without modeling other point tokens (i.e.,  $\Delta x_k$  only comes from  $x_k$ ). By pooling all outputs of the Dynamic Adapter, our internal prompts ( $\sum \Delta x_k$ ) can serve as a global perspective and interact directly with each point tokens  $x_k + \Delta x_k$ , enabling access to the cumulative information of the entire sequence. Besides, the bi-directional interaction in self-attention strengthens the representation of the Dynamic Adapter, further improving the performance.

Please refer to Algorithm 1 for more detailed information on Dynamic Adapter and Internal Prompt.

#### 4.4. Analysis

**Tunable parameters.** The proposed DAPT keeps the entire backbone frozen and only tunes the newly inserted parameters, effectively reducing the number of tunable parameters. Specifically, the TFTS introduces only two parameters in  $\mathbb{R}^{1 \times d}$  for each operation. Each extra Dynamic Adapter module adds  $2d$  from layer norm,  $d$  from  $W_s$ ,  $2 \times r \times d$  from  $W_d$  and  $W_u$  (bias ignored) as tunable parameters, where the intermediate dimension  $r$  is relatively small compared to  $d$ . Our Internal Prompt tuning can propose a prompt by pooling and TFTS without introducing other parameters. In summary, we only need to fine-tune fewer than 5% of the parameters to achieve comparable results.

**Training memory.** The proposed approach can consistently reduce training memory as the batch size increases.

As shown in Fig. 2(c), with 512 batch size, our DAPT significantly reduces GPU memory<sup>1</sup> usage by 35% and 49% compared to full fine-tuning and IDPT [54] (the recent SOTA PETL method for point cloud), respectively.

**Storage.** The proposed DAPT significantly reduces the storage cost. After being fine-tuned on  $M$  point cloud datasets, only one copy of the original pre-trained model’s full parameters and  $M$  dataset-specific trainable parameters need to be stored, saving up to 95% of storage space. Thus, even if storage is not a current concern, DAPT presents a promising approach for adapting the expansion of increasingly larger models in 3D vision.

## 5. Experiments

### 5.1. Implement Details

To ensure a fair comparison, we employed identical experimental settings to the default fine-tuning method for each baseline. This entails freezing the weights of the pre-trained point cloud backbone and solely updating the newly inserted parameters during training. All experiments are conducted on a single GeForce RTX 4090.

### 5.2. 3D Classification

**Real-World Object Classification.** ScanObjectNN [43] is a highly challenging 3D dataset covering  $\sim 15K$  real-world objects across 15 categories. These objects consist of indoor scene data obtained by scanning, often characterized by cluttered backgrounds and occlusion caused by other objects. As shown in Tab. 2, we conducted experiments on three variants of ScanObjectNN (i.e., OBJ\_BG, OBJ\_ONLY, and PB\_T50\_RS), each with increasing complexity. We utilize various baselines such as Point-BERT [52], Point-MAE [35], and RECON [39] to demonstrate the effectiveness and generalization of our approach. It is worth noting that our DAPT meets or exceeds the performance of full fine-tuning in most cases, while only tuning less than 5% of the parameters. Especially, DAPT achieved increases of 3.62%, 1.55%, and 2.36% in the three ScanObjectNN variants on Point-BERT. Furthermore, we surpass IDPT [54], the most recent SOTA method for 3D efficient fine-tuning, in almost every variant for different baselines, with 35% fewer tunable parameters. A similar phenomenon is also observed in Point-MAE and RECON. Moreover, our DAPT has only a slight increase in FLOPs compared to the full fine-tuning and is quite competitive with IDPT.

**Synthetic Object Classification.** ModelNet40 [47] contains 12,311 pristine 3D CAD models across 40 categories. Every point cloud is complete, uniform, and noise-free. Due to the voting strategy [29] being time-consuming, we prioritize reporting overall Accuracy without voting. As shown in Tab. 2, although the performance in ModelNet40

<sup>1</sup>The gradient checkpoint [5] is already being used to save memory.

Table 2. Classification on three variants of the ScanObjectNN [43] and the ModelNet40 [47], including the number of trainable parameters and overall accuracy (OA). All methods utilize the default data argumentation as the baseline. \* denotes reproduced results. We report ScanObjectNN results without voting. ModelNet40 results are without and with voting, referred to (-/-).

Method	Reference	Tunable params. (M)	FLOPs (G)	ScanObjectNN			ModelNet40	
				OBJ_BG	OBJ_ONLY	PB_T50_RS	Points Num.	OA (%)
<i>Supervised Learning Only</i>								
PointNet [37]	CVPR 17	3.5	0.5	73.3	79.2	68.0	1k	- / 89.2
PointNet++ [38]	NeurIPS 17	1.5	1.7	82.3	84.3	77.9	1k	- / 90.7
DGCNN [45]	TOG 19	1.8	2.4	82.8	86.2	78.1	1k	- / 92.9
MVTN [11]	ICCV 21	11.2	43.7	-	-	82.8	1k	- / 93.8
PointNetXt [41]	NeurIPS 22	1.4	1.6	-	-	87.7	1k	- / 94.0
PointMLP [32]	ICLR 22	13.2	31.4	-	-	85.4	1k	- / 94.5
RepSurf-U [42]	CVPR 22	1.5	0.8	-	-	84.3	1k	- / 94.4
ADS [14]	ICCV 23	-	-	-	-	87.5	1k	- / 95.1
<i>Self-Supervised Representation Learning (Full fine-tuning)</i>								
OcCo [44]	ICCV 21	22.1	4.8	84.85	85.54	78.79	1k	- / 92.1
Point-BERT [52]	CVPR 22	22.1	4.8	87.43	88.12	83.07	1k	- / 93.2
MaskPoint [28]	ECCV 22	22.1	-	89.70	89.30	84.60	1k	- / 93.8
Point-MAE [35]	ECCV 22	22.1	4.8	90.02	88.29	85.18	1k	- / 93.8
Point-M2AE [57]	NeurIPS 22	15.3	3.6	91.22	88.81	86.43	1k	- / 94.0
ACT [8]	ICLR 23	22.1	4.8	93.29	91.91	88.21	1k	- / 93.7
RECON [39]	ICML 23	43.6	5.3	94.15	93.12	89.73	1k	- / 93.9
<i>Self-Supervised Representation Learning (Efficient fine-tuning)</i>								
Point-BERT [52] (baseline)	CVPR 22	22.1 (100%)	4.8	87.43	88.12	83.07	1k	92.7 / 93.2
+ IDPT [54]	ICCV 23	1.7 (7.69%)	7.2	88.12(+0.69)	88.30(+0.18)	83.69(+0.62)	1k	92.6(-0.1) / 93.4(+0.2)
+ DAPT (ours)	-	<b>1.1 (4.97%)</b>	5.0	91.05(+3.62)	89.67(+1.55)	85.43(+2.36)	1k	93.1(+0.4) / 93.6(+0.4)
Point-MAE [35] (baseline)	ECCV 22	22.1 (100%)	4.8	90.02	88.29	85.18	1k	93.2 / 93.8
+ IDPT [54]	ICCV 23	1.7 (7.69%)	7.2	91.22(+1.20)	90.02(+1.73)	84.94(-0.24)	1k	93.3(+0.1) / 94.4(+0.6)
+ DAPT (ours)	-	<b>1.1 (4.97%)</b>	5.0	90.88(+0.86)	90.19(+1.90)	85.08(-0.10)	1k	93.5(+0.3) / 94.0(+0.2)
RECON [39] (baseline <sup>2</sup> )	ICML 23	22.1 (100%)	4.8	94.32	92.77	90.01	1k	92.5 / 93.0
+ IDPT* [54]	ICCV 23	1.7 (7.69%)	7.2	93.29(-1.03)	91.57(-1.20)	87.27(-2.74)	1k	93.4(+0.9) / 93.5(+0.5)
+ DAPT (ours)	-	<b>1.1 (4.97%)</b>	5.0	94.32(0.00)	92.43(-0.34)	89.38(-0.63)	1k	93.5(+1.0) / 94.1(+1.1)

is quite saturated, our method consistently achieves performance gains compared to different baselines and IDPT [54]. Specifically, RECON combined with our DAPT achieves 1.0% improvement and significantly reduces tunable parameters. Even after voting, our DAPT remains the leader.

**Few-shot Learning.** We further conduct few-shot experiments on ModelNet40 [47] to prove our few-shot transfer learning ability. Consistent with previous works [35, 52], we employ the n-way, m-shot configuration. As shown in Tab. 3, it can be observed that DAPT achieved performance enhancements in most cases compared to full fine-tuning and IDPT, indicating its effectiveness in few-shot learning.

### 5.3. Part Segmentation

Part segmentation is challenging to predict a more detailed class label for every point. We evaluate the effectiveness of DAPT on ShapeNetPart [50], which includes 16,881 samples from 16 categories. Following previous works [8, 35], we incorporate Dynamic Adapter into the 2nd, 6th, and 10th layers, generating an Internal Prompt to feed into the input of the 3rd, 7th, and 11th layers and employing them as

global features. As shown in Tab. 4, our DAPT achieves competitive results on Inst. mIoU and notable improvements on Cls. mIoU compared with IDPT [54]. Keep in mind that the increase in tunable parameters mainly comes from the head part, but we still reduce the tunable parameters and gain improvements compared with IDPT.

### 5.4. Compared with Other PETL Methods

To further prove the effectiveness of our proposed method, we compare DAPT with several PETL approaches from NLP and 2D vision. As shown in Tab. 5, we select the most challenging PB\_T50\_RS variant with Point-MAE as the baseline. Note that we have carefully tuned the hyperparameters of these methods to attain optimal performance. While they show different performance improvements compared to linear probing, there is still a considerable gap compared to full fine-tuning. Although these methods excel in NLP or 2D, they still struggle to adapt to 3D data due to

<sup>2</sup>We modify the structure by loading a pre-trained RECON weight into Point-MAE as a single-modal fine-tuning approach to reproduce similar results as the original paper.

Table 3. Few-shot learning on ModelNet40 [47]. Overall accuracy (%)±the standard deviation (%) without voting is reported.

Methods	Reference	5-way		10-way	
		10-shot	20-shot	10-shot	20-shot
<i>with Self-Supervised Representation Learning (Full fine-tuning)</i>					
OcCo [44]	ICCV 21	94.0±3.6	95.9±2.3	89.4±5.1	92.4±4.6
Point-BERT [52]	CVPR 22	94.6±3.1	96.3±2.7	91.0±5.4	92.7±5.1
MaskPoint [28]	ECCV 22	95.0±3.7	97.2±1.7	91.4±4.0	93.4±3.5
Point-MAE [35]	ECCV 22	96.3±2.5	97.8±1.8	92.6±4.1	95.0±3.0
Point-M2AE [57]	NeurIPS 22	96.8±1.8	98.3±1.4	92.3±4.5	95.0±3.0
ACT [8]	ICLR 23	96.8±2.3	98.0±1.4	93.3±4.0	95.6±2.8
RECON [39]	ICML 23	97.3±1.9	98.9±3.9	93.3±3.9	95.8±3.0
<i>with Self-Supervised Representation Learning (Efficient fine-tuning)</i>					
Point-BERT [52] (baseline)	CVPR 22	94.6±3.1	96.3±2.7	91.0±5.4	92.7±5.1
+ IDPT [54]	ICCV 23	<b>96.0±1.7</b>	97.2±2.6	91.9±4.4	93.6±3.5
+ DAPT (ours)	-	95.8±2.1	<b>97.3±1.3</b>	<b>92.2±4.3</b>	<b>94.2±3.4</b>
Point-MAE [35] (baseline)	ECCV 22	96.3±2.5	97.8±1.8	92.6±4.1	95.0±3.0
+ IDPT [54]	ICCV 23	<b>97.3±2.1</b>	97.9±1.1	92.8±4.1	95.4±2.9
+ DAPT (ours)	-	96.8±1.8	<b>98.0±1.0</b>	<b>93.0±3.5</b>	<b>95.5±3.2</b>

Table 4. Part segmentation on the ShapeNetPart [50]. The mIoU for all classes (Cls.) and for all instances (Inst.) are reported. #TP represents the tunable parameters. \* denotes reproduced results.

Methods	Reference	#TP (M)	Cls. mIoU (%)	Inst. mIoU (%)
<i>Supervised Learning Only</i>				
PointNet [37]	CVPR 17	-	80.39	83.7
PointNet++ [38]	NeurIPS 17	-	81.85	85.1
DGCNN [45]	TOG 19	-	82.33	85.2
APES [46]	CVPR 23	-	83.67	85.8
<i>Self-Supervised Representation Learning (Full fine-tuning)</i>				
OcCo [44]	ICCV 21	27.09	83.42	85.1
MaskPoint [28]	ECCV 22	-	84.60	86.0
Point-BERT [52]	CVPR 22	27.09	84.11	85.6
Point-MAE [35]	ECCV 22	27.06	84.19	86.1
ACT [8]	ICLR 23	27.06	84.66	86.1
<i>Self-Supervised Representation Learning (Efficient fine-tuning)</i>				
Point-BERT [52] (baseline)	CVPR 22	27.06	84.11	85.6
+ IDPT* [54]	ICCV 23	5.69	83.50	85.3
+ DAPT (ours)	-	<b>5.65</b>	<b>83.83</b>	<b>85.5</b>
Point-MAE [35] (baseline)	ECCV 22	27.06	84.19	86.1
+ IDPT [54]	ICCV 23	5.69	83.79	85.7
+ DAPT (ours)	-	<b>5.65</b>	<b>84.01</b>	<b>85.7</b>
RECON [39] (baseline <sup>2</sup> )	ICML 23	27.06	84.52	86.1
+ IDPT* [54]	ICCV 23	5.69	83.66	85.7
+ DAPT (ours)	-	<b>5.65</b>	<b>83.87</b>	<b>85.7</b>

missing point clouds and other complexities. It is worth noting that our method surpasses the IDPT while reducing the number of tunable parameters by 35% and achieving a better trade-off between tunable parameters and performance.

## 5.5. Ablation Study

We conducted ablation studies on the most challenging PB\_T50\_RS variant to investigate the rationalization and effectiveness of the design of our proposed DAPT and then fine-tuned on Point-MAE as the baseline.

**Analysis on each component.** We conduct experiments to prove the effectiveness of the proposed components of our DAPT and display the results in Tab. 6. When no additional inserted components exist, the method is the same

Table 5. Comparisons of parameter efficient transfer learning methods from NLP and 2D Vision on the hardest variant of ScanObjectNN [43]. Overall accuracy (%) without voting is reported. #TP represents the tunable parameters.

Method	Reference	#TP (M)	PB_T50_RS
Point-MAE [28]	ECCV 22	22.1	<b>85.18</b>
Linear probing	-	0.3	75.99
+ Adapter [16]	ICML 19	0.9	83.93
+ Prefix tuning [25]	ACL 21	0.7	77.72
+ BitFit [53]	ACL 21	0.3	82.62
+ LoRA [17]	ICLR 22	0.9	81.74
+ VPT-Deep [18]	ECCV 22	0.4	81.09
+ AdaptFormer [4]	NeurIPS 22	0.9	83.45
+ SSF [26]	NeurIPS 22	0.4	82.58
+ IDPT [54]	ICCV 23	1.7	84.94
+ DAPT (ours)	-	1.1	<b>85.08</b>

Table 6. The effect of each component of our DAPT. The tunable parameters (#TP) and the overall accuracy (%) on the hardest variant of ScanObjectNN [43] are reported. In. and Ex. indicate Internal and External respectively.

TFTS	DA	In. Prompt	Ex. Prompt	#TP (M)	PB_T50_RS
		Full fine-tuning		22.1	85.18
		Linear Probing		0.27	75.99
✓				0.37	83.17
	✓			0.88	84.25
✓	✓			0.99	84.80
✓	✓		✓	1.08	84.70
✓	✓	✓		1.09	<b>85.08</b>

as Linear Probing, which only achieves 75.99% overall accuracy. By adopting our TFTS, we can achieve a substantial 7.18% performance increase while only adding 0.1M tunable parameters. To examine the effectiveness of our proposed Dynamic Adapter, we exclusively use Dynamic Adapter and achieve an 8.26% performance gain. These results indicate that TFTS and Dynamic Adapter are powerful parameter-efficient transfer learning modules. We further combine the above two components to gain an 84.80% performance. We also externally replaced the Internal Prompt, using a Xavier uniform initializer [18] mentioned in VPT [18]. Compared to our Internal Prompt, the External Prompt shows a 0.38% performance fall and even a 0.1% performance drop compared with no Prompt tuning, demonstrating the effectiveness of Internal Prompts.

**Analysis on token selections for head inputs.** We examined the influence of the input tokens of the downstream task head, including the classification token, the pooling of our prompt tokens, and the pooling of point patch tokens. As depicted in Fig. 3, we selected three standard token selection methods. We observe that the best performance is attained when all three features are included:  $T_{cls}$ , prompt tokens, and point patch tokens. Comparing the use of  $T_{cls}$  only and both  $T_{cls}$  and pooling of Patch tokens, using all the features resulted in a 4.13% and 0.66% performance gain.

**Analysis on dimension  $r$  in Dynamic Adapter.** One of our keys to reducing tunable parameters is selecting dimension  $r$  in Dynamic Adapter. As shown in Tab. 7a, we

Table 7. Ablation on hyper-parameters and settings of DAPT, including dimension  $r$ , settings of dynamic scale and inserted layers. The tunable parameters (#TP) and the overall accuracy (%) on the hardest variant of ScanObjectNN [43] are reported.

(a) Ablation on $r$ of Dynamic Adapter.			(b) Ablation on scale $S_d$ in Dynamic Adapter.			(c) Ablation on the inserted layers.		
Dimension $r$	#TP (M)	PB_T50_RS	Scale	#TP (M)	PB_T50_RS	Layers	#TP (M)	PB_T50_RS
8	0.57	84.39	1.0	1.08	84.59	1→3	0.62	80.67
16	0.64	84.52	5.0	1.08	84.42	1→6	0.78	82.48
32	0.79	84.18	10.0	1.08	84.11	1→9	0.93	83.97
64	1.09	<b>85.08</b>	Dynamic scale w/o ReLU	1.09	84.73	4→12	0.93	81.75
72	1.16	84.46	Dynamic scale	1.09	<b>85.08</b>	1→12	1.09	<b>85.08</b>

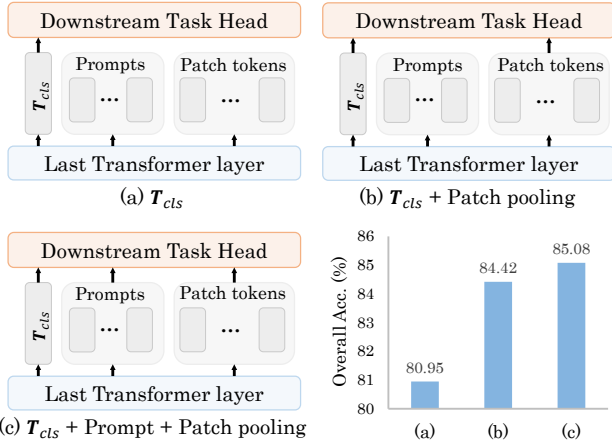


Figure 3. Effect of different inputs for downstream task head. We conduct experiments on the hardest variant (i.e., PB\_T50\_RS) of ScanObjectNN [43] with Point-MAE [35] baseline.

conduct experiments on various values of  $r$  and find that when  $r = 64$ , we can achieve the best trade-off between tunable parameters and performance.

**Analysis on scale  $S_d$  in Dynamic Adapter.** It is crucial to obtain a dynamic scale during inference due to the complexity of point clouds. As indicated in Tab. 7b, fixing the scale to 1.0 yields the best result but decreases by 0.49% compared to our dynamic scale. We also perform a variation of dynamic scale, removing the ReLU, and observe a 0.35% decrease, which is still superior to the fixed scale. In addition, we displayed the average value of  $S_d$  and the ratio of adjusted tokens in each layer by choosing three samples in Fig. 4. While their trends are similar, various samples differ across different layers, demonstrating that our method can dynamically generate scale for different samples to overcome the complex geometry structures of point clouds.

**Analysis on inserted layers for Dynamic Adapter and Internal Prompts.** One way to further reduce the tunable parameters is to select only a portion of the layer to insert our components. However, as shown in Tab. 7c, we observe that fine-tuning only shallow blocks and fine-tuning only deep blocks leads to different performance decline levels. The potential reason may be the difficulty in determining which layer is more important when the network processes the diverse point cloud. To address this issue, we incorpo-




Figure 4. A visual results of the average value of  $S_d$  and the ratio of adjusted tokens for each layer across three samples in OBJ\_ONLY variant of ScanObjectNN [43].

rate our proposed components for each layer and utilize the dynamic scale to regulate the extent of feature adjustment.

## 6. Conclusion and Limitation

In this study, we propose a simple yet effective parameter-efficient transfer learning strategy named DAPT for point cloud analysis. The proposed DAPT freezes the parameters of the pre-trained backbones and leverages the Dynamic Adapter to adjust each token dynamically by considering the significance. In addition, we further use Dynamic Adapter to implement Internal Prompt tuning, better capturing the instance-specific feature. Our method provides a practical solution to reduce storage cost requirements without compromising performance. One limitation is that whether our method could perform well in more complex tasks such as 3D object detection and generation is still unclear, which is our future work.

**Acknowledgement.** This work was supported in part by the NSFC (No.62225603), in part by the Hubei Key R&D Program (No.2022BAA078), in part by the Taihu Lake Innovation Fund for Future Technology (HUST:2023-A-1), and in part by the National Undergraduate Training Projects for Innovation and Entrepreneurship (202310487020).



## 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 *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2022. 2, 4
- [2] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2022. 2
- [3] Guangyan Chen, Meiling Wang, Li Yuan, Yi Yang, and Yufeng Yue. Rethinking point cloud registration as masking and reconstruction. In *Proc. of IEEE Intl. Conf. on Computer Vision*, 2023. 1
- [4] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. In *Proc. of Advances in Neural Information Processing Systems*, 2022. 2, 3, 4, 7
- [5] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016. 5
- [6] Yukang Chen, Jianhui Liu, Xiangyu Zhang, Xiaojuan Qi, and Jiaya Jia. Voxelnex: Fully sparse voxelnet for 3d object detection and tracking. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2023. 1
- [7] Silin Cheng, Xiwu Chen, Xinwei He, Zhe Liu, and Xiang Bai. Pra-net: Point relation-aware network for 3d point cloud analysis. *IEEE Transactions on Image Processing*, 2021. 1
- [8] 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? In *Proc. of Intl. Conf. on Learning Representations*, 2022. 2, 6, 7
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. of Intl. Conf. on Learning Representations*, 2021. 1, 3
- [10] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2017. 1
- [11] Abdullah Hamdi, Silvio Giancola, and Bernard Ghanem. Mvtn: Multi-view transformation network for 3d shape recognition. In *Proc. of IEEE Intl. Conf. on Computer Vision*, 2021. 6
- [12] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *Proc. of Intl. Conf. on Learning Representations*, 2021. 2, 3, 4
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2022. 1
- [14] Cheng-Yao Hong, Yu-Ying Chou, and Tyng-Luh Liu. Attention discriminant sampling for point clouds. In *Proc. of IEEE Intl. Conf. on Computer Vision*, 2023. 6
- [15] Jinghua Hou, Zhe Liu, Dingkan Liang, Zhikang Zou, Xiaoqing Ye, and Xiang Bai. Query-based temporal fusion with explicit motion for 3d object detection. In *Proc. of Advances in Neural Information Processing Systems*, 2024. 2
- [16] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *Proc. of Intl. Conf. on Machine Learning*, 2019. 2, 3, 4, 7
- [17] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *Proc. of Intl. Conf. on Learning Representations*, 2021. 2, 3, 7
- [18] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *Proc. of European Conference on Computer Vision*, 2022. 3, 4, 7
- [19] Shibo Jie and Zhi-Hong Deng. Fact: Factor-tuning for lightweight adaptation on vision transformer. In *Proc. of the AAAI Conf. on Artificial Intelligence*, 2023. 2, 3
- [20] Shibo Jie, Haoqing Wang, and Zhi-Hong Deng. Revisiting the parameter efficiency of adapters from the perspective of precision redundancy. In *Proc. of IEEE Intl. Conf. on Computer Vision*, 2023.
- [21] Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. In *Proc. of Advances in Neural Information Processing Systems*, 2021. 3
- [22] Xin Lai, Jianhui Liu, Li Jiang, Liwei Wang, Hengshuang Zhao, Shu Liu, Xiaojuan Qi, and Jiaya Jia. Stratified transformer for 3d point cloud segmentation. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2022. 1
- [23] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proc. of Conf. on Empirical Methods in Natural Language Processing*, 2021. 2, 3, 4
- [24] Jinyu Li, Chenxu Luo, and Xiaodong Yang. Pillarnext: Rethinking network designs for 3d object detection in lidar point clouds. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2023. 1
- [25] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Annual Meeting of the Association for Computational Linguistics*, 2021. 2, 3, 4, 7
- [26] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. In *Proc. of Advances in Neural Information Processing Systems*, 2022. 3, 4, 7
- [27] Dingkan Liang, Xin Zhou, Xinyu Wang, Xingkui Zhu, Wei Xu, Zhikang Zou, Xiaoqing Ye, and Xiang Bai. Pointmamba: A simple state space model for point cloud analysis. *arXiv preprint arXiv:2402.10739*, 2024. 2

- [28] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. In *Proc. of European Conference on Computer Vision*, 2022. 2, 6, 7
- [29] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2019. 5
- [30] Zhe Liu, Tengpeng Huang, Bingling Li, Xiwu Chen, Xi Wang, and Xiang Bai. Epnnet++: Cascade bi-directional fusion for multi-modal 3d object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1
- [31] Gen Luo, Yiyi Zhou, Tianhe Ren, Shengxin Chen, Xiaoshuai Sun, and Rongrong Ji. Cheap and quick: Efficient vision-language instruction tuning for large language models. In *Proc. of Advances in Neural Information Processing Systems*, 2023. 2, 3
- [32] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point cloud: A simple residual mlp framework. In *Proc. of Intl. Conf. on Learning Representations*, 2022. 6
- [33] Yuning Mao, Lambert Mathias, Rui Hou, Amjad Almahairi, Hao Ma, Jiawei Han, Scott Yih, and Madian Khabsa. Unipelt: A unified framework for parameter-efficient language model tuning. In *Annual Meeting of the Association for Computational Linguistics*, 2022. 3
- [34] Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. Autosdf: Shape priors for 3d completion, reconstruction and generation. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2022. 1
- [35] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Proc. of European Conference on Computer Vision*, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [36] Omid Poursaeed, Tianxing Jiang, Han Qiao, Nayun Xu, and Vladimir G Kim. Self-supervised learning of point clouds via orientation estimation. In *Intl. Conf. on 3D Vision*. IEEE, 2020. 2
- [37] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2017. 6, 7
- [38] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Proc. of Advances in Neural Information Processing Systems*, 2017. 1, 6, 7
- [39] 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 *Proc. of Intl. Conf. on Machine Learning*, 2023. 1, 2, 4, 5, 6, 7
- [40] Zekun Qi, Runpei Dong, Shaochen Zhang, Haoran Geng, Chunrui Han, Zheng Ge, Li Yi, and Kaisheng Ma. Shapellm: Universal 3d object understanding for embodied interaction. *arXiv preprint arXiv:2402.17766*, 2024. 1, 2
- [41] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. In *Proc. of Advances in Neural Information Processing Systems*, 2022. 6
- [42] Haoxi Ran, Jun Liu, and Chengjie Wang. Surface representation for point clouds. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2022. 6
- [43] 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 *Proc. of IEEE Intl. Conf. on Computer Vision*, 2019. 3, 5, 6, 7, 8
- [44] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *Proc. of IEEE Intl. Conf. on Computer Vision*, 2021. 6, 7
- [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 Transactions ON Graphics*, 2019. 3, 6, 7
- [46] Chengzhi Wu, Junwei Zheng, Julius Pfommer, and Jürgen Beyerer. Attention-based point cloud edge sampling. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2023. 7
- [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 *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2015. 5, 6, 7
- [48] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *Proc. of European Conference on Computer Vision*, 2020. 2
- [49] Kaixin Xiong, Shi Gong, Xiaoqing Ye, Xiao Tan, Ji Wan, Errui Ding, Jingdong Wang, and Xiang Bai. Cape: Camera view position embedding for multi-view 3d object detection. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2023. 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 Transactions ON Graphics*, 2016. 6, 7
- [51] Dongshuo Yin, Yiran Yang, Zhechao Wang, Hongfeng Yu, Kaiwen Wei, and Xian Sun. 1% vs 100%: Parameter-efficient low rank adapter for dense predictions. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2023. 3
- [52] 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 *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, 2022. 2, 3, 5, 6, 7
- [53] Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In *Annual Meeting of the Association for Computational Linguistics*, 2022. 7
- [54] Yaohua Zha, Jinpeng Wang, Tao Dai, Bin Chen, Zhi Wang, and Shu-Tao Xia. Instance-aware dynamic prompt tuning for pre-trained point cloud models. In *Proc. of IEEE Intl. Conf. on Computer Vision*, 2023. 3, 4, 5, 6, 7

- [55] Yaohua Zha, Huizhen Ji, Jinmin Li, Rongsheng Li, Tao Dai, Bin Chen, Zhi Wang, and Shu-Tao Xia. Towards compact 3d representations via point feature enhancement masked autoencoders. In *Proc. of the AAAI Conf. on Artificial Intelligence*, 2024. [2](#)
- [56] Dingyuan Zhang, Ding kang Liang, Zhikang Zou, Jingyu Li, Xiaoqing Ye, Zhe Liu, Xiao Tan, and Xiang Bai. A simple vision transformer for weakly semi-supervised 3d object detection. In *Proc. of IEEE Intl. Conf. on Computer Vision*, 2023. [2](#)
- [57] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Point-m2ae: multi-scale masked autoencoders for hierarchical point cloud pre-training. In *Proc. of Advances in Neural Information Processing Systems*, 2022. [2](#), [4](#), [6](#), [7](#)
- [58] Zixiang Zhou, Xiangchen Zhao, Yu Wang, Panqu Wang, and Hassan Foroosh. Centerformer: Center-based transformer for 3d object detection. In *Proc. of European Conference on Computer Vision*, 2022. [2](#)
- [59] Yaoming Zhu, Jiangtao Feng, Chengqi Zhao, Mingxuan Wang, and Lei Li. Counter-interference adapter for multilingual machine translation. In *Proc. of Conf. on Empirical Methods in Natural Language Processing*, 2021. [3](#)