

# Conv-Adapter: Exploring Parameter Efficient Transfer Learning for ConvNets

Hao Chen<sup>1\*</sup>, Ran Tao<sup>1</sup>, Han Zhang<sup>1</sup>, Yidong Wang<sup>2</sup>, Xiang Li<sup>1</sup>,  
Wei Ye<sup>2</sup>, Jindong Wang<sup>3</sup>, Guosheng Hu<sup>4</sup>, Marios Savvides<sup>1</sup>

<sup>1</sup> Carnegie Mellon University, <sup>2</sup> Peking University, <sup>3</sup> Microsoft Research Asia, <sup>4</sup> Oosto

## Abstract

While parameter efficient tuning (PET) methods have shown great potential with transformer architecture on Natural Language Processing (NLP) tasks, their effectiveness with large-scale ConvNets is still under-studied on Computer Vision (CV) tasks. This paper proposes Conv-Adapter, a PET module designed for ConvNets. Conv-Adapter is light-weight, domain-transferable, and architecture-agnostic with generalized performance on different tasks. When transferring on downstream tasks, Conv-Adapter learns task-specific feature modulation to the intermediate representations of backbones while keeping the pre-trained parameters frozen. By introducing only a tiny amount of learnable parameters, e.g., only 3.5% full fine-tuning parameters of ResNet50. It can also be applied for transformer-based backbones. Conv-Adapter outperforms previous PET baseline methods and achieves comparable or surpasses the performance of full fine-tuning on 23 classification tasks of various domains. It also presents superior performance on the few-shot classification with an average margin of 3.39%. Beyond classification, Conv-Adapter can generalize to detection and segmentation tasks with more than 50% reduction of parameters but comparable performance to the traditional full fine-tuning<sup>1</sup>

## 1. Introduction

As transfer learning [54] thrives, large-scale foundation models gradually dominate deep learning over the last few years [3]. Fine-tuning has become the de-facto paradigm adapting a foundation model pre-trained on a pretext task to various downstream tasks for both Computer Vision (CV) and Natural Language Processing (NLP). Albeit its simplicity and prominence, fine-tuning has been posing challenges to development and deployment of the large-scale foundation models on downstream tasks with the drastic growth of computations and storage costs, as the param-




Figure 1. Performance of Conv-Adapter compared to other transfer learning methods on ResNet-50 BiT-M. We compute the relative performance gain w.r.t to fine-tuning and percentage of trainable parameters of the backbone (w/o linear head) on 23 image classification datasets from various domains to compute the results, with mean and standard deviation highlighted. Conv-Adapter achieves a superior trade-off between transfer accuracy and parameter efficiency.

ter size increases from millions [19, 23, 46, 52] to billions [5, 13, 14, 16, 34–36, 45].

Parameter efficient tuning (PET), as an alternative to traditional fine-tuning, has become prevalent in NLP [18, 22, 24, 30, 31] for its efficiency and effectiveness. PET introduces a small number of learnable parameters to a pre-trained network, whose parameters are frozen, and learns the extra introduced parameters only. While attaining promising performance, especially for tasks of low-data regimes [25, 62, 63], PET modules for Convolutional Neural Networks (ConvNets), the popular architectures for CV tasks, are still largely unstudied.

Prior arts on fine-tuning ConvNets to multiple visual domains are restrictive in generalization and parameter efficiency. Bias Tuning [2], which tunes only the bias terms of the backbone, might fail on domains with significant distribution shifts from the pre-training tasks. Residual Adapter [48] and TinyTL [7] are mainly designed for small net-

\*haoc3@andrew.cmu.edu

<sup>1</sup>Code is available at: <https://github.com/Hhhhhhao/Conv-Adapter/tree/main>

works such as ResNet-26 [19] and MobileNet [6, 23]. It is prohibitive to scale these previous designs to larger ConvNets [36] or more diverse domains [60]. Besides, previous PET methods [18, 21, 24, 30, 31] are mainly designed with Transformer [56] architecture for NLP tasks [5, 13]. However, it is not straightforward to apply Transformer-based PET to ConvNets because Transformers tokenize and sequentialize the input and features, while ConvNets do not. Recent works [1, 10, 25] that attempt to use Prompt Tuning [30] and Adapters [21] on CV tasks are also designed for Vision Transformers rather than ConvNets. Furthermore, the downstream CV tasks are usually more diverse with a larger domain gap compared with NLP [45]. These challenges motivate us to design the architecture and adapting scheme of PET for ConvNets, which could make it transferable to various CV tasks, including image classification, object detection, and semantic segmentation.

In this work, we narrow the gap of PET between NLP and CV with the proposal of **Conv-Adapter** – an adaption module that is light-weight, domain-transferable, and architecture-agnostic. Conv-Adapter learns task-specific knowledge on downstream tasks and adapts the intermediate features of each residual block in the pre-trained ConvNets. It has a bottleneck structure consisting of depth-wise separable convolutions [23] and non-linearity. Due to the variety of CV network architectures and tasks, we explore four adapting schemes of Conv-Adapter combining two design perspectives - adapted representations and insertion form to verify the optimal tuning paradigm on ConvNets. We find it is essential for Conv-Adapter to maintain the locality relationship when adapting intermediate feature maps for transferability. More importantly, Conv-Adapter can be formulated under the same mathematical framework as the PET modules used in the NLP field [18]. Conv-Adapter outperforms previous PET baselines and achieves similar or even better performance to the traditional *full* fine-tuning on 23 cross-domain classification datasets with an average of 3.5% of the backbone parameters using ResNet-50 BiT-M [27], as shown in Fig. 1. Conv-Adapter also well generalizes to object detection and semantic segmentation tasks with same-level performance to fully fine-tuning. To further understand Conv-Adapter, in addition, we empirically analyze the performance of Conv-Adapter with both the domain shifting of datasets and the network weights shifting brought by fine-tuning. The core contributions of this work can be summarized as:

- To our knowledge, we are the first to *systematically* investigate the feasible solutions of general parameter-efficient tuning (PET) for ConvNets. This investigation can narrow the gap between NLP and CV for PET.
- We propose Conv-Adapter, a light-weight and plug-and-play PET module, along with four adapting variants following two design dimensions - transferability and pa-

parameter efficiency. Meanwhile, we empirically justify several essential design choices to make Conv-Adapter effectively transferred to different CV tasks.

- Extensive experiments demonstrate the effectiveness and efficiency of Conv-Adapter. It achieves comparable or even better performance to full fine-tuning with only around 5% backbone parameters. Conv-Adapter also well generalizes to detection and segmentation tasks that require dense predictions.

## 2. Related Work

### 2.1. Parameter Efficient Tuning for Transformers

Pre-trained Transformer models in NLP are usually of the size of billions of parameters [5, 13, 16], which makes fine-tuning inefficient as one needs to train and maintain a separate copy of the backbone parameters on each downstream task. Adapter [21] is first proposal to conduct transfer with light-weight adapter modules. It learns the task-specific knowledge and composes it into the pre-trained backbone [43, 44] when adapting to a new task. Similarly, LoRA introduces trainable low-rank matrices to each layer of the backbone model to approximate parameter updates. Different from inserting adaption modules to intermediate layers, Prefix Tuning [31] and Prompt Tuning [30], inspired by the success of textual prompts [5, 33, 45], prepend learnable prompt tokens to input and only train these tokens when transferring to a new task. More recently, a unified formulation of Adapter, LoRA, and Prefix Tuning is proposed in [18], where their core function is to adapt the intermediate representation of the pre-trained model by residual task-specific representation learned by tuning modules.

Visual Prompt Tuning [25] is a recent method adapting Prompt Tuning from NLP to Vision Transformers [25]. Bahng et. al. [1] also explores visual prompts in input pixel space for adapting CLIP models [45] and makes connection with [15]. While showing promising results on Transformers, visual prompts on ConvNets presents much worse transfer results [1, 25], possibly due to the limited capacity of input space visual prompts. Conv-Adapter can adapt the intermediate features thus has larger capacity.

### 2.2. Transfer Learning for ConvNets

While there is no straightforward approach to applying previous PET methods designed for Transformers directly on ConvNets, several attempts have been made in prior research. BatchNorm Tuning [40] and Bias Tuning [2] only tune the batchnorm related terms or the bias terms of the pre-trained backbone. Piggyback [39] instead learns weight masks for downstream tasks while keeping the pre-trained backbone unchanged. They all have limited transferability and update partial parameters of the backbone.

More related to our work, Residual Adapter [48] ex-

plores inserting an extra convolutional layer of kernel size 1 to each convolutional layer in pre-trained ResNet-26 [19], either in parallel or in sequential, to conduct the multi-domain transfer. Similarly, TinyTL introduces extra residual blocks to MobileNet [6, 23] for memory efficient on-device learning. Guo et. al. [17] proposes re-composing a ResNet with depth-wise and point-wise convolutions, and re-training only the depth-wise part during fine-tuning. RepNet [59] exploits a dedicated designed side network to re-program the intermediate features of pre-trained ConvNets. Conv-Adapter differs from previous methods with a design that considers parameter efficiency and transferability from the internal architectures and adapting schemes. Besides, the proposed Conv-Adapter does not require tuning any backbone parameters to achieve comparable performance to fine-tuning.

### 3. Method

#### 3.1. Preliminaries

Parameter efficient tuning (PET) methods [21, 24, 25, 30, 31] introduce learnable adapting modules plugged into the backbone that is frozen during tuning. From a unified point of view, the core function of the adaption modules is to learn task-specific feature modulations on originally hidden representations in the pre-trained backbone [18]. Specifically, considering an intermediate hidden representation  $\mathbf{h}$  generated by a layer or a series of layers with input  $\mathbf{x}$  in a pre-trained network, the PET adaption module learns  $\Delta\mathbf{h}$  and updates  $\mathbf{h}$  as:

$$\mathbf{h} \leftarrow \mathbf{h} + \alpha \cdot \Delta\mathbf{h}, \quad (1)$$

where  $\alpha$  could be a scalar [24] or a gating function [31]. Previous PET methods in NLP mainly follow a similar functional form for constructing  $\Delta\mathbf{h}$  – down-sampling projection, non-linearity, and up-sampling projection. However, they differ in 1) implementation (architecture) - the form of the projections and non-linearity, and 2) the adapting scheme - which  $\mathbf{h}$  in the model to adapt and compute  $\Delta\mathbf{h}$  from which representation. These differences characterize the adaptation to new tasks and robustness to out-of-distribution evaluation [31].

It is non-trivial to design effective PET methods for ConvNets because previous PET modules are mainly developed on Transformers rather than ConvNets. Besides, the components of the architecture and computation dynamics of ConvNets and Transformers are inherently different. Following the unified formulation of PET methods in Eq. (1), we propose **Conv-Adapter**. We construct the  $\Delta\mathbf{h}$  of Conv-Adapter similarly to previous PET methods and design the adaption architecture and scheme on ConvNets from the perspective of transferability and parameter efficiency.

#### 3.2. Motivation

Before delving into the details of our design, we identify the essential difficulty that prevents utilizing prior arts directly on ConvNets as an adaption module and thus inspires us to propose Conv-Adapter. Conventionally, for ConvNets,  $\mathbf{h}$  and  $\Delta\mathbf{h}$  are usually 3-dimensional structural features maps belonging to  $\mathbb{R}^{C \times H \times W}$  with  $C$  being the channel dimension and  $H \times W$  being the spatial size of the feature maps.

The difference in intermediate feature and processing dynamics poses obstacles to transferability. For Transformers,  $\mathbf{h}$  is whereas 2-dimensional sequential features in  $\mathbb{R}^{L \times D}$  where  $L$  is the sequence length and  $D$  is the feature dimension. Previous PET modules for Transformers compute  $\Delta\mathbf{h}$  in various forms, e.g., linear layers over  $\mathbf{h}$  [21] and self-attention over additional input prompts [25, 30, 31]. They can all process the sequential features globally with long-range dependencies as the computing blocks in Transformers. Although it is possible to apply linear layers, or equivalently  $1 \times 1$  convolutional layers [48], to adapt the feature maps of ConvNets, it is yet intuitive that this might produce inferior transfer performance due to the *loss of locality*, which is encoded in the structural features maps by convolutions of kernel size larger than 1. The *loss of locality* results in a radical mismatch of the receptive field in  $\Delta\mathbf{h}$  and  $\mathbf{h}$ , which might be destructive when adapting ConvNets on tasks with significant domain shifts. Apart from the receptive field mismatch, the spatial size of feature maps in ConvNets also significantly affects the transferability of adaption. Earlier attempts to use adapters to transfer ConvNets usually downsample the feature’s spatial size for memory and parameter efficiency. However, for CV tasks beyond image classification like segmentation, the spatial size matters for achieving good results [9, 49].

In summary, it is crucial to design the architecture and adapting scheme of the PET module computing  $\Delta\mathbf{h}$  for ConvNets to have the same spatial size of feature maps and the same receptive field of convolutions for transferability.

#### 3.3. Architecture of Conv-Adapter

Given the above challenges, we design our **Conv-Adapter** as a bottleneck structure, which is also widely used by PET methods of NLP tasks [19, 21]. However, our **Conv-Adapter** designs the bottleneck, particularly for ConvNets. Precisely, it consists of two convolutional layers with a non-linearity function in-between. The first convolution conducts channel dimension down-sampling with a kernel size similar to that of the adapted blocks, whereas the second convolution projects the channel dimension back. For simplicity, we adopt the same activation function used in the backbone as the non-linearity at the middle of the bottleneck. The effective receptive field of the modulated feature maps produced by Conv-Adapter is thus similar to that of the adapted blocks in the backbone. We do not change the




Figure 2. Architecture of Conv-Adapter, which has a bottleneck composed of depth-wise separable convolutions with non-linearity activation.  $C_{in}$ ,  $C_{out}$ ,  $H$ ,  $W$  is set to keep the same as in backbone.  $\alpha$  and  $\gamma$  are hyper-parameters to tune.

spatial size of the feature maps for better transferability on dense prediction tasks. We adopt the depth-wise separable convolutions [23] for Conv-Adapter to reduce the parameter size further.

Figure 2 illustrates our Conv-Adapter architecture. Formally, let the input feature map to the adapted blocks of the ConvNets be  $\mathbf{z} \in \mathbb{R}^{C_{in} \times H \times W}$  and the output feature maps be  $\mathbf{h} \in \mathbb{R}^{C_{out} \times H \times W}$ , where  $C_{in}$  and  $C_{out}$  are the channel dimension of the input and output to the adapted blocks respectively. Assuming the spatial size  $H \times W$  of the feature maps does not change along these blocks, we set the learnable weight as  $\mathbf{W}_{down} \in \mathbb{R}^{\frac{C_{in}}{\gamma} \times \gamma \times K \times K}$  for the depth-wise convolution and  $\mathbf{W}_{up} \in \mathbb{R}^{C_{out} \times \frac{C_{in}}{\gamma} \times 1 \times 1}$  for the point-wise convolution in Conv-Adapter, with the non-linearity denoted as  $f$ . We use a compression factor of  $\gamma$  to denote the down-sampling in the channel dimension, where  $\gamma$  is a hyper-parameter tuned for each task. Mathematically, Conv-Adapter computes  $\Delta\mathbf{h} \in \mathbb{R}^{C_{out} \times H \times W}$  as:

$$\Delta\mathbf{h} = (\mathbf{W}_{up} \otimes f(\mathbf{W}_{down} \hat{\otimes} \mathbf{z})), \quad (2)$$

where  $\otimes$  and  $\hat{\otimes}$  denotes point-wise and depth-wise convolution, respectively. To allow the modulation  $\Delta\mathbf{h}$  to be more flexibly composed into  $\mathbf{h}$ , we set  $\alpha$  in Eq. (1) as a learnable scaling vector in  $\mathbb{R}^{C_{out}}$ , which is initialized as ones. The ablation study on design choices is presented in Sec. 4.5.

### 3.4. Adapting ConvNets with Conv-Adapter

After setting the architecture of Conv-Adapter, we discuss the scheme to adapt a variety of ConvNets. Previous PET methods insert the adapting modules to Self-Attention blocks, Feed-Forward blocks, or both [18] of Transformers, which have a relatively unified architecture. In contrast, modern ConvNets usually stacks either residual blocks [19, 51, 61] or inverted residual blocks [23, 36, 52, 53],

which consists of a series of convolutional layers (and sometimes pooling layers) and a residual identity branch, making it more difficult to use a single adapting scheme to various architectures.

To explore the effective adapting schemes of using Conv-Adapter to tune a ConvNet, we study it mainly from two perspectives, similar to [18], 1) the location of adaptation in pre-trained ConvNets – which intermediate representation  $\mathbf{h}$  to adapt, and 2) the insertion form of Conv-Adapter – how to set the input  $\mathbf{z}$  to Conv-Adapter to compute  $\Delta\mathbf{h}$ . From the location perspective, we study plugging Conv-Adapter to each (inverted) residual block [7] or to each functioning  $K \times K$  convolutional layer within a residual block [17]. From the insertion perspective, Conv-Adapter can be inserted either in parallel or in sequential to the modified components, with the input to Conv-Adapter being  $\mathbf{x}$ , the input to the modified components, or being  $\mathbf{h}$  itself, respectively. Combining the design dimension from these two perspectives, we propose 4 variants of adapting schemes with Conv-Adapter: **Convolution Parallel**, **Convolution Sequential**, **Residual Parallel**, and **Residual Sequential**.

Taking the bottleneck residual block of ResNet-50 [19] as an example, we demonstrate the proposed designs in Fig. 3. As  $1 \times 1$  convolution layer can only transfer channel-wise information, we thus design the adapting of functional convolutions, i.e., intermediate  $K \times K$  convolutions, to keep locality sensitive. On the contrary, adapting the whole residual block considers the transferring of pre-trained knowledge carried by  $1 \times 1$  convolutions. Intuitively, adapting the whole residual blocks has a larger capacity for modulating task-specific features than adapting only  $K \times K$  convolution but may introduce more parameters. Plugging Conv-Adapter stage-wisely is not considered as it is impractical to make the receptive field of Conv-Adapter similar to the adapted stage with only two convolutions. It needs a more sophisticated design on not only the Conv-Adapter architecture but also the adaptation location [59], and we empirically find that stage-wise adaptation produces inferior performance and requires much more parameters. Conv-Adapter is flexible to be inserted into every residual block of the ConvNet backbone for transferability of features from different depths, as in [39, 48]. Other backbones such as ConvNext [36], and even Swin-Transformer [34] can be adapted following the same guideline (see experiments).

## 4. Experiments

This section verifies the transferability and parameter efficiency of Conv-Adapter from various aspects, including image classification, few-shot classification, object detection, and semantic segmentation. Additionally, we provide an ablation study of Conv-Adapter for its design choices and an analysis of its performance.

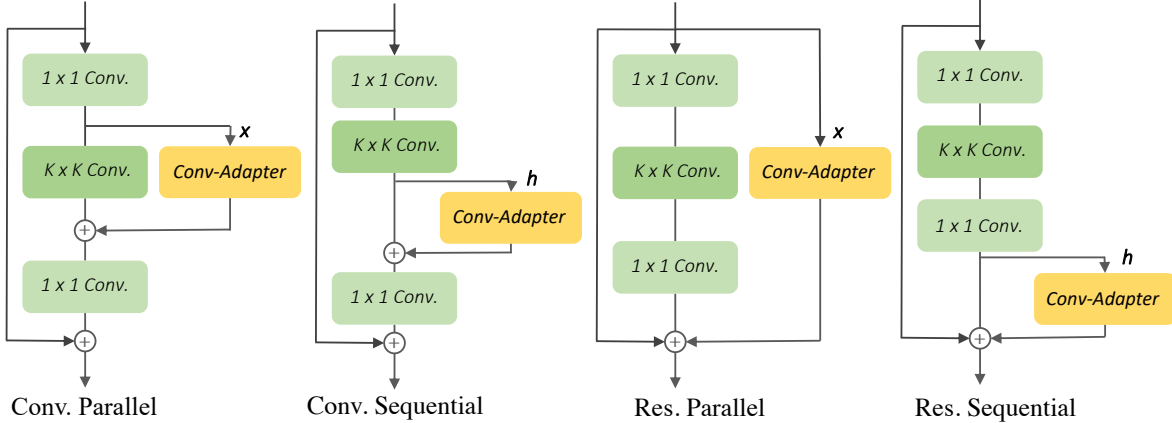


Figure 3. Four adapting schemes of Conv-Adapter to ResNet50: Convolution Parallel, Convolutional Sequential, Residual Parallel, and Residual Sequential. The schemes differ regarding the position of the modified representation and corresponding insertion form. Other networks can be adapted similarly following the illustration. Green modules are frozen during fine-tuning.

## 4.1. Transferability of Conv-Adapter

### 4.1.1 Setup

We first evaluate the transferability of Conv-Adapter on classification tasks. We experiment on two benchmarks: VTAB-1k [60] and FGVC. VTAB-1k includes 19 diverse visual classification tasks, which are grouped into three categories: *Natural*, *Specialized*, and *Structured* based on the domain of the images. Each task in VTAB-1k contains 1,000 training examples. FGVC consists of 4 *Fine-Grained Visual Classification* tasks: CUB-200-2011 [57], Stanford Dogs [26], Stanford Cars [29], and NABirds [55].

For evaluation, we compare the 4 variants of Conv-Adapter with full fine-tuning (FT) and 3 baseline methods: linear probing (LP), bias tuning (Bias) [7], and visual prompt tuning (VPT) [1, 25]. We test each method on ResNet50 [19, 27] with ImageNet21k pre-training. To find the optimal hyper-parameters of Conv-Adapter (and baseline methods), we conduct a grid search of the learning rate, weight decay, and compression factor  $\gamma$  for each dataset using the validation data split from training data for both benchmarks. For VTAB-1k, we use the recommended optimal data augmentations in [60], rather than solely Resize and Centre Crop as in [1, 63]. We find the recommended augmentations produces better results for full-tuning. For FGVC, we use RandomResized Crop with a minimum scale of 0.2 and Horizontal Flip [50] as augmentation. More details of the hyper-parameters are shown in Appendix.

### 4.1.2 Results and Discussion

Results are reported in Tab. 1. Conv-Adapter not only demonstrates significant improvements over the baseline methods, but also achieves the same level of performance

Table 1. Performance of Conv-Adapter adapting schemes on ResNet-50 BiT-M. Each setting includes three runs and averaged top-1 accuracy (%) over datasets and the averaged total trainable parameters (M) over all datasets are reported. We compare proposed variants of Conv-Adapter (in gray) with full Fine-Tuning (FT), Linear Probing (LP), Bias Tuning (Bias), and Visual Prompt Tuning (VPT). We report the number of wins of (-) for each method compared to FT. **Bold** and underline refer to the top and second result separately.

Tuning	# Param.	FGVC	VTAB-1k		
			Natural	Specialized	Structured
# Tasks	-	4	7	4	8
FT	23.89	83.46	72.19	<b>85.86</b>	<b>66.72</b>
LP	0.37	75.44 (1)	67.42 (4)	81.42 (0)	37.92 (0)
Bias	0.41	64.98 (0)	66.06 (4)	80.34 (0)	32.18 (0)
VPT	0.42	74.79 (1)	65.43 (2)	80.35 (0)	37.64 (0)
Conv. Par.	0.85	83.77 (3)	<b>72.60 (5)</b>	84.21 (1)	56.70 (1)
Conv. Seq.	0.87	79.68 (2)	72.28 (4)	83.85 (0)	58.50 (1)
Res. Par.	8.21	<b>84.24 (3)</b>	71.75 (4)	84.70 (0)	61.34 (1)
Res. Seq.	3.53	83.45 (2)	71.74 (4)	84.84 (0)	61.33 (2)

or even surpasses their fine-tuning counterparts on all domains evaluated, by introducing only around **3.5%** of full fine-tuning parameters for ResNet-50. Notably, there is a considerable performance gap, i.e., an improvement of **23.44%**, of Conv-Adapter over previous baseline methods on *Structured* datasets of VTAB-1k.

One can observe that the proposed four variants of Conv-Adapter all achieve comparable performance compared to full fine-tuning. Among the four variants, **Convolution Parallel** achieves the best trade-off between performance and parameter efficiency. On the evaluated classification tasks, inserting Conv-Adapter in parallel generally outperforms inserting sequentially. In terms of the modified repre-

sensation, one can find that, on most of the datasets, adapting only the  $K \times K$  convolutions of ResNet-50 can achieve performance close to fine-tuning. However, on *Structured* datasets, adapting whole residual blocks is far better than adjusting only the middle convolutions with more parameters, demonstrating the superior capacity of adjusting residual blocks when there is a more significant domain gap.

## 4.2. Universality of Conv-Adapter

### 4.2.1 Setup

We evaluate the universality of Conv-Adapter on classification tasks in this section, where Conv-Adapter is inserted to various ConvNets architectures with different pre-training. We adopt the simple yet effective adapting scheme – Convolution Parallel, and mainly compare it with full fine-tuning. More specifically, we adopt ImageNet-21k pre-trained ResNet50 [27], ConvNext-B and ConvNext-L [36], and even Swin-B and Swin-L [34]. Apart from ImageNet-21k, we evaluate ImageNet-1k, CLIP [20], and MoCov3 [11] pre-training. Similarly, we conduct a hyper-parameter search on the validation set, and report the accuracy on the test set of FGVC and VTAB-1k. Model details are shown in Appendix.

### 4.2.2 Results and Discussion

We present the results in Tab. 2. On various ImageNet-21k pre-trained ConvNets, Conv-Adapter demonstrates its universality with comparable performance to fine-tuning. For large models such as ConvNext-L and Swin-L, conducting traditional fine-tuning requires training nearly 196M parameters, whereas Conv-Adapter improves the parameter efficiency with only 7.8% and 4.5% of the fine-tuning parameters on ConvNext-L and Swin-L respectively. Although the transfer performance of Conv-Adapter on ImageNet-1k pre-trained models is more limited, compared to ImageNet-21k pre-training, Conv-Adapter still demonstrates its superior parameter efficiency and shows improvement over fine-tuning on several tasks. For the CLIP vision models, Conv-Adapter consistently outperforms fine-tuning on Structured tasks of VTAB-1k. We observe a performance gap of Conv-Adapter on MoCov3 pre-trained [11], and we argue this is possibly due to the difference in feature space of self-supervised and supervised models in CV [25].

## 4.3. Few-Shot Classification

### 4.3.1 Setup

PET methods usually present superior performance for tasks with low-data regimes [18, 31]. We thus evaluate Conv-Adapter on few-shot classification using ImageNet-21k pre-trained ResNet50 Bit-M [27] and ConvNext-B [36]. We evaluate 5 FGVC datasets using 1, 2, 4, 8 shots for

Table 2. Comparing Conv-Adapter (CA) with full Fine-Tuning (FT) using various backbone architectures of different pre-training. Each setting includes three runs and averaged top-1 accuracy (%) over datasets and the averaged total trainable parameters (M) over all datasets are reported. We report the number of wins of (·) for each method in compared to FT. **Bold** indicates the best results.

Pre-train	Backbone	Tuning	# Param.	FGVC	VTAB-1k		
					Natural	Specialized	Structured
# Tasks				4	7	4	8
ImageNet 21k	ResNet50	FT	23.89	83.46	72.19	<b>85.86</b>	<b>66.72</b>
		CA	0.85	<b>83.77 (3)</b>	<b>72.60 (5)</b>	84.21 (1)	56.70 (1)
	ConvNext-B	FT	87.75	<b>89.48</b>	<b>81.59</b>	<b>87.32</b>	<b>65.77</b>
		CA	6.83	89.28 (1)	80.62 (4)	86.29 (0)	64.88 (2)
	ConvNext-L	FT	196.50	90.64	<b>82.25</b>	<b>87.94</b>	<b>67.65</b>
		CA	15.52	<b>90.69 (3)</b>	81.7 (2)	86.85 (0)	64.98 (3)
	Swin-B	FT	86.92	<b>90.01</b>	78.65	<b>87.59</b>	<b>64.69</b>
		CA	4.98	88.55 (1)	<b>80.00 (4)</b>	85.84 (0)	62.57 (2)
	Swin-L	FT	195.27	<b>91.04</b>	80.64	<b>87.85</b>	<b>66</b>
		CA	8.86	90.54 (2)	<b>81.39 (3)</b>	86.29 (1)	63.19 (2)
ImageNet 1k	ResNet50	FT	23.87	<b>85.84</b>	<b>67.15</b>	<b>83.53</b>	<b>53.32</b>
		CA	0.72	83.48 (0)	64.20 (0)	81.33 (1)	52.74 (2)
	ConvNext-B	FT	87.75	<b>88.95</b>	74.51	<b>85.33</b>	61.34
		CA	10.82	87.84 (1)	<b>74.72 (4)</b>	84.29 (0)	<b>63.77 (2)</b>
CLIP	ResNet50	FT	38.50	<b>81.38</b>	<b>58.53</b>	<b>80.8</b>	57.18
		CA	2.23	76.64 (0)	56.33 (3)	79.12 (0)	<b>58.96 (4)</b>
	ResNet50x4	FT	87.17	<b>84.23</b>	<b>65.71</b>	<b>82.22</b>	58.84
		CA	6.14	82.71 (0)	62.54 (2)	80.72 (1)	<b>59.10 (4)</b>
MoCov3	ResNet50	FT	23.87	<b>83.92</b>	<b>66.25</b>	<b>83.89</b>	<b>60.26</b>
		CA	0.89	79.69 (0)	65.31 (3)	81.59 (0)	53.87 (1)

each class following following previous studies [25, 45, 63] including Food101 [4], Oxford Flowers [41], Oxford Pets [42], Stanford Cars [29], and Aircraft [38]. Averaged top-1 accuracy is reported in Tab. 3. We search from the same range as before and adopt the same augmentations as for FGVC tasks. The detailed hyper-parameters and more results for each dataset are in Appendix.

### 4.3.2 Results and Discussion

Compared with Fine-tuning, Conv-Adapter boosts few-shot classifications with an average 3.39% margin over different shots using only around 5% trainable parameters. Especially for 1/2-shot cases, Conv-Adapter shows supreme performance compared with Fine-tuning and VPT [25] (11.07% on 1-shot and 6.99% on 2-shot with larger architecture ConvNext-B). Meanwhile, Conv-Adapter provides a better accuracy-efficiency trade-off than Visual Prompt Tuning on few-shot classifications. It surpasses VPT with an average margin of 1.35% with ResNet50 Bit-M and 3.69% with ConvNext-B. In the 8-shot case, VPT drops around 8% performance compared with Fine-tuning due to limited capacity, while Conv-Adapter can achieve comparable or better performance to Fine-tuning and maintain parameter efficiency.

Table 3. Few-shot classification: the average Top-1 accuracy over 5 FGVC datasets, with 1, 2, 4, 8 shots. We compare Conv-Adapter (CA), Visual Prompt Tuning (VPT), and full Fine-Tuning (FT). **Bold** indicates the best results.

Backbone	Tuning	# Param	1	2	4	8
ResNet50 Bit-M	FT	23.72	29.30	38.96	50.09	61.27
	VPT	0.24	32.56	42.18	52.21	59.37
	CA	1.02	<b>34.31</b>	<b>43.55</b>	<b>52.43</b>	<b>61.42</b>
ConvNext-B	FT	87.68	36.34	48.83	<b>63.69</b>	<b>76.91</b>
	VPT	0.13	42.25	51.85	62.89	69.04
	CA	4.6	<b>47.41</b>	<b>55.82</b>	63.25	74.29

Table 4. Object detection & Semantic Segmentation results. We report the results of fine-tuning and Conv-Adapter with the Residual Parallel scheme.

Object Detection with Faster-RCNN					
Backbone	Tuning	# Param	AP	AP <sub>50</sub>	AP <sub>75</sub>
ResNet50	FT	41.53	38.1	59.7	41.5
	CA	35.72	<b>38.4</b>	<b>61.1</b>	<b>41.5</b>
ConvNeXt-B	FT	67.09	<b>45.2</b>	<b>67.2</b>	<b>49.9</b>
	CA	24.62	41.9	64.5	45.7

Semantic Segmentation with UPerNet			
Backbone	Tuning	# Param (M)	mIoU
ResNet50	FT	66.49	42.1
	CA	45.65	<b>43.0</b>
ConvNeXt-B	FT	81.87	<b>48.7</b>
	CA	39.40	46.9

## 4.4. Object Detection and Semantic Segmentation

### 4.4.1 Setup

Beyond image classification tasks, we also validate the generalization of Conv-Adapter on dense prediction tasks, including object detection and semantic segmentation. We use ImageNet-21k pre-trained ResNet50 and ConvNeXt-S as backbones. For object detection, we implement Conv-Adapter with Faster-RCNN using the MMDetection [8] framework compared with fine-tuning. We report the average precision (AP) results on the validation split of the MS-COCO dataset [32]. For semantic segmentation, we implement Conv-Adapter with UPerNet [58] using MM-Segmentation framework [12] and conduct experiments on the ADE20K dataset [64], with mIoU reported on the validation split.

For object detection, we compare all four schemes of Conv-Adapter with the fine-tuning baseline. Specifically, we follow a standard 1x training schedule: all models are trained with a batch size of 16 and optimized by AdamW with an initial learning rate of 0.0002 for Faster RCNN and 0.0001 for RetinaNet, which are then dropped by a factor of 10 at the 8-th and 11-th epoch. The shorter side of the input image is resized to 800 while maintaining the original aspect ratio. For segmentation, we train all models for 80k iterations with a random cropping augmentation of  $512 \times$

512 input resolution. For ConvNeXt models, we use a larger input resolution of  $640 \times 640$  and train the models for 160k iterations. We apply AdamW optimizer with a polynomial learning rate decay schedule. More detailed training setting and hyper-parameters are shown in Appendix.

### 4.4.2 Results and Discussion

The dense prediction results are summarized in Tab. 4. We observe a different effect of Conv-Adapter on two types of backbones. On ResNet50, Conv-Adapter surpasses fine-tuning with fewer trainable parameters (including the dense prediction heads) for object detection and semantic segmentation. On ConvNeXt-S, the performance is lower than their fine-tuning counterparts. We argue that the inferior performance of Conv-Adapter on ConvNeXt-S on dense prediction tasks is due to severely reduced model capacity as the number of trainable parameters is reduced by more than 50%. Nevertheless, they can still outperform the ResNet50 with fewer total parameters. This indicates there might be overfitting issues, and we encourage more future studies on this topic.

## 4.5. Ablation Study

### 4.5.1 Setup

We provide an ablation study on the design choices of Conv-Adapter, where we explore different architectures and adapting schemes. In this section, we mainly report the Top-1 accuracy on the validation set of VTAB-1k.

### 4.5.2 Architecture and Adapting Schemes

We first compare the performance of Conv-Adapter using depth-wise separable, regular, and  $1 \times 1$  convolutions (linear layers). As shown in Tab. 5, depth-wise separable convolution introduces the minimal parameter budget while achieving the best results. Apart from 4 adapting variants proposed in this work, we also explore other design choices used in previous works. We experiment on spatial down-sampling of feature maps [7]. Compared to channel down-sampling with a bottleneck in Conv-Adapter, spatial down-sampling introduces nearly 27 times of parameters with inferior accuracy. We also validate the adapting scheme of applying  $1 \times 1$  convolution to all convolutional layers [48], which introduces nearly 16 times of parameters to Conv-Adapter with -12.27% accuracy gain. Finally, we evaluate the adapting scheme that inserts Conv-Adapter stage-wisely, which is less effective in both parameter size and performance than the proposed schemes.

Table 5. Ablation study on more adapting scheme and more architectures of Conv-Adapter. The different schemes and architectures mainly come from previous works. The proposed adaptation and architecture achieve the best results.

Adapting Scheme	Down-sample	# Convs	Type of Conv.	# Param	VTAB-1k
$K \times K$ Conv. Par.	Channel	2	Depth-wise	0.67	<b>71.03</b>
$K \times K$ Conv. Par.	Channel	2	Regular	5.66	70.52
$K \times K$ Conv. Par.	Channel	2	Linear	1.22	68.32
$K \times K$ Conv. Par.	Spatial	2	Depth-wise	18.45	68.54
All Conv. Par	-	1	Linear	10.74	58.75
Stage Par.	Channel	2	Depth-wise	1.90	65.06




Figure 4. Sensitivity to hyper-parameters of initialization of learnable scaling vector  $\alpha$  and compression factor  $\gamma$ .

#### 4.5.3 Sensitivity to $\gamma$ and initialization of $\alpha$

We explicitly study the sensitivity of the transfer performance to the initialization of the learnable scaling vector  $\alpha$  and compression factor  $\gamma$  in Conv-Adapter, as shown in Fig. 4. When initializing  $\alpha$  as ones, Conv-Adapter achieves the best performance on the validation set of VTAB-1k. Compared to  $\alpha$ , Conv-Adapter is more robust to the compression factor  $\gamma$ , achieving similar performance with the compression factor of 1, 2, and 4. Setting  $\gamma$  with a larger value results in inferior performance with a more limited capacity of Conv-Adapter.

#### 4.5.4 Kernel size in Conv-Adapter

We show the performance of Conv-Adapter on VTAB-1k validation set in Fig 5, of using different kernel size for the depth-wise convolution to verify our argument of the *loss of locality*. One can observe that, for both ResNet50 and ConvNext-B, using smaller kernel size results in inferior performance. When setting the kernel size larger to that of the residual blocks, i.e., 5 and 7 for ResNet50, the performance is further boosted, with more parameters introduced.

#### 4.5.5 CKA Similarity of Conv-Adapter

We observe from Tab. 1 and Tab. 2 that, on datasets with large domain shifts, Conv-Adapter (and baseline methods) may fail to generalize well. To investigate the reason, we compute the CKA similarity [28, 47] between weights of




Figure 5. Sensitivity to kernel size of depth-wise convolution in Conv-Adapter, for both ResNet50 and ConvNext-B.

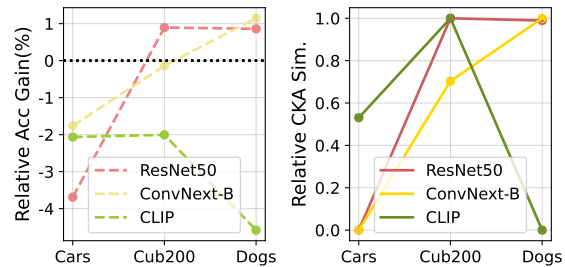


Figure 6. CKA similarity and accuracy gap between Conv-Adapter and fully fine-tuning for FGVC datasets.

convolutional filters for the pre-trained and fine-tuned backbone. The lower the CKA similarity, the larger capacity is required for good transfer performance. We plot the CKA similarity and the relative accuracy gain of Conv-Adapter to fine-tuning in Fig. 7, where the same trends over datasets exhibit for different architectures. When fully fine-tuning only leads to small changes in filter weights (larger CKA similarities), Conv-Adapter is more likely to surpass the performance of fully fine-tuning. More detail on CKA similarity comparison is in Appendix.

## 5. Conclusions

In this work, we propose Conv-Adapter, a parameter efficient tuning module for ConvNets. Conv-Adapter is light-weight, domain-transferable, and model-agnostic. Extensive experiments on classification and dense prediction tasks show it can achieve performance comparable to full fine-tuning with much fewer parameters. We find Conv-Adapter might fail on tasks with large domain shifts and subject to feature quality determined by pre-training. Future work includes more exploration of Conv-Adapter on domain robustness and dense predictions and NAS for Conv-Adapter.

## References

[1] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Exploring visual prompts for adapting large-



- scale models. *arXiv preprint arXiv:2203.17274*, 2022. 2, 5
- [2] Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2022. 1, 2
- [3] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, et al. On the opportunities and risks of foundation models, 2021. 1
- [4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014. 6
- [5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. 1, 2
- [6] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019. 2, 3
- [7] Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. Tinytl: Reduce memory, not parameters for efficient on-device learning. In *Advances in Neural Information Processing Systems*, pages 11285–11297. Curran Associates, Inc., 2020. 1, 4, 5, 7
- [8] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 7
- [9] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 3
- [10] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition, 2022. 2
- [11] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. *arXiv preprint arXiv:2104.02057*, 2021. 6
- [12] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 7
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2018. 1, 2
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2020. 1
- [15] Gamaleldin F. Elsayed, Ian Goodfellow, and Jascha Sohl-Dickstein. Adversarial reprogramming of neural networks, 2018. 2
- [16] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity, 2021. 1, 2
- [17] Yunhui Guo, Yandong Li, Liqiang Wang, and Tajana Rosing. Depthwise convolution is all you need for learning multiple visual domains. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:8368–8375, 2019. 3, 4
- [18] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *International Conference on Learning Representations*, 2022. 1, 2, 3, 4, 6
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 3, 4, 5
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6, 2
- [21] 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 *Proceedings of the 36th International Conference on Machine Learning*, 2019. 2, 3
- [22] 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 *Proceedings of the 36th International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019. 1
- [23] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017. 1, 2, 3, 4
- [24] Edward Hu, Yelong Shen, Phil Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. 1, 2, 3
- [25] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 3, 5, 6
- [26] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Li Fei-Fei. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, CO, 2011. 5

- [27] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning, 2019. [2](#), [5](#), [6](#), [3](#), [4](#)
- [28] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International Conference on Machine Learning*, pages 3519–3529. PMLR, 2019. [8](#), [3](#)
- [29] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. [5](#), [6](#)
- [30] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic, 2021. Association for Computational Linguistics. [1](#), [2](#), [3](#)
- [31] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation, 2021. [1](#), [2](#), [3](#), [6](#)
- [32] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context, 2014. cite arxiv:1405.0312Comment: 1) updated annotation pipeline description and figures; 2) added new section describing datasets splits; 3) updated author list. [7](#)
- [33] 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, 2021. [2](#)
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [1](#), [4](#), [6](#), [2](#)
- [35] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, Furu Wei, and Baining Guo. Swin transformer v2: Scaling up capacity and resolution. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [36] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [2](#), [4](#), [6](#), [3](#)
- [37] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2017. [2](#)
- [38] S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. Technical report, 2013. [6](#)
- [39] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights, 2018. [2](#), [4](#)
- [40] Pramod Kaushik Mudrakarta, Mark Sandler, Andrey Zhmoginov, and Andrew Howard. K for the price of 1: Parameter-efficient multi-task and transfer learning, 2018. [2](#)
- [41] M-E Nilsback and Andrew Zisserman. A visual vocabulary for flower classification. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, pages 1447–1454. IEEE, 2006. [6](#)
- [42] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012. [6](#)
- [43] Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. AdapterHub: A framework for adapting transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54, Online, 2020. Association for Computational Linguistics. [2](#)
- [44] Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapterfusion: Non-destructive task composition for transfer learning. *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 2021. [2](#)
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. [1](#), [2](#), [6](#)
- [46] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollar. Designing network design spaces. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [1](#)
- [47] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems*, 34:12116–12128, 2021. [8](#), [3](#)
- [48] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Efficient parametrization of multi-domain deep neural networks. In *CVPR*, 2018. [1](#), [2](#), [3](#), [4](#), [7](#)
- [49] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):640–651, 2017. [3](#)
- [50] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. [5](#), [1](#)
- [51] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. [4](#)
- [52] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks, 2019. [1](#), [4](#)
- [53] Mingxing Tan and Quoc V. Le. Efficientnetv2: Smaller models and faster training, 2021. [4](#)

- [54] Sebastian Thrun. Lifelong learning algorithms. In *Learning to learn*, pages 181–209. Springer, 1998. [1](#)
- [55] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 595–604, 2015. [5](#)
- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017. [2](#)
- [57] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. [5](#)
- [58] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 418–434, 2018. [7](#)
- [59] Li Yang, Adnan Siraj Rakin, and Deliang Fan. Rep-net: Efficient on-device learning via feature reprogramming. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12277–12286, 2022. [3](#), [4](#)
- [60] Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruysen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, Lucas Beyer, Olivier Bachem, Michael Tschannen, Marcin Michalski, Olivier Bousquet, Sylvain Gelly, and Neil Houlsby. A large-scale study of representation learning with the visual task adaptation benchmark, 2019. [2](#), [5](#), [1](#)
- [61] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Haibin Lin, Zhi Zhang, Yue Sun, Tong He, Jonas Mueller, R. Manmatha, Mu Li, and Alexander Smola. Resnet: Split-attention networks, 2020. [4](#)
- [62] Ningyu Zhang, Luoqiu Li, Xiang Chen, Shumin Deng, Zhen Bi, Chuanqi Tan, Fei Huang, and Huajun Chen. Differentiable prompt makes pre-trained language models better few-shot learners, 2021. [1](#)
- [63] Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. Neural prompt search, 2022. [1](#), [5](#), [6](#)
- [64] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [7](#)