

LESSONS LEARNED FROM A UNIFYING EMPIRICAL STUDY OF PARAMETER-EFFICIENT TRANSFER LEARNING (PETL) IN VISUAL RECOGNITION

Anonymous authors

Paper under double-blind review

ABSTRACT

Parameter-efficient transfer learning (PETL) has attracted significant attention lately, due to the increasing size of pre-trained models and the need to fine-tune them for superior downstream performance. This community-wide enthusiasm has sparked a plethora of approaches. Nevertheless, a systematic study to understand their performance and suitable application scenarios is lacking, leaving questions like “*when to apply PETL*” and “*which approach to use*” largely unanswered, especially in visual recognition. In this paper, we conduct a unifying empirical study of representative PETL approaches in the context of Vision Transformers (ViT). We systematically tune their hyper-parameters to fairly compare their accuracy on downstream tasks. Our study not only offers a valuable user guide but also unveils several new insights. First, if tuned carefully, different PETL approaches can obtain quite similar accuracy in the low-shot benchmark VTAB-1K. This includes simple approaches like fine-tuning the bias terms that were reported inferior. Second, though with similar accuracy, we find that PETL approaches make different mistakes and high-confidence predictions, likely due to their different inductive biases. Such an inconsistency (or complementariness) opens up the opportunity for ensemble methods, and we make preliminary attempts at this. Third, going beyond the commonly used low-shot tasks, we find that PETL is also useful in many-shot regimes — it achieves comparable and sometimes better accuracy than full fine-tuning, using much fewer learnable parameters. Last but not least, we investigate PETL’s ability to preserve a pre-trained model’s robustness to distribution shifts (*e.g.*, a CLIP backbone). Perhaps not surprisingly, PETL approaches outperform full fine-tuning alone. However, with weight-space ensembles, the fully fine-tuned model can better balance target (*i.e.*, downstream) distribution and distribution shift performance, suggesting a future research direction for PETL.

1 INTRODUCTION

Pre-training and then fine-tuning has become the standard practice to tackle visual recognition problems (Bommasani et al., 2021). The community-wide enthusiasm for open-sourcing has made it possible to access large, powerful pre-trained models learned from a gigantic amount of data, *e.g.*, ImageNet-21K (Ridnik et al., 2021) or LAION-5B (Schuhmann et al., 2022). More research focus has thus been on how to fine-tune such large models (Yu et al., 2023a). Among existing efforts, parameter-efficient transfer learning (PETL), *a.k.a* parameter-efficient fine-tuning (PEFT), has attracted increasing attention lately (Han et al., 2024; Ding et al., 2023). Instead of fine-tuning the whole model (*i.e.*, full fine-tuning) or the last fully connected layer (*i.e.*, linear probing), PETL approaches seek to update or insert a relatively small number of parameters to the pre-trained model (Xin et al., 2024). Doing so has several noticeable advantages. First, as named, PETL is parameter-efficient. For one downstream task (*e.g.*, recognizing bird species or car brands), it only needs to learn and store a tiny fraction of parameters on top of the pre-trained model. Second, accuracy-wise, PETL has been shown to consistently outperform linear probing and often beat full fine-tuning, as reported on the commonly used low-shot image classification benchmark VTAB-1K (Zhai et al., 2019).

To date, a plethora of PETL approaches have been proposed, bringing in inspiring ideas and promising results. Along with this come several excellent surveys that summarize existing PETL approaches




Figure 1: Highlights of our insights. **(a) Downstream accuracy:** if tuned carefully, different PETL methods achieve similar accuracy (● for the range from the most to the least accurate methods) and consistently outperform linear probing (×) and full fine-tuning (■) on VTAB-1K. **(b) Distribution shift accuracy:** fine-tuning a CLIP ViT-B/16, known for its generalizability across domains, with PETL on ImageNet-1K (100 samples/class) better preserves the distribution shift accuracy (Y-axis, averaged across ImageNet-V2, ImageNet-S, ImageNet-R, and ImageNet-A) than full fine-tuning, evidenced by the ★ points. Interestingly, *weight-space ensembles* (WiSE) (Wortsman et al., 2022) is applicable between PETL’s fine-tuned model and the pre-trained model (■), but not as effective as applying it to the fully fine-tuned model. Details are in section 3 and section 7.

(Yu et al., 2023a; Xin et al., 2024; Ding et al., 2023). Yet, a systematic understanding of the PETL paradigm seems still missing.

For example, with so many PETL approaches, there is a lack of unifying references for when and how to apply them. Though superior accuracy was reported on the low-shot benchmark VTAB-1K, there is not much discussion on how PETL approaches achieve it. Does it result from PETL’s ability to promote transferability or prevent over-fitting? The current evaluation also raises the question of whether PETL is useful beyond a low-shot scenario. Last but not least, besides superior accuracy, do existing PETL approaches offer different, ideally, complementary information?

Attempting to answer these questions, we conduct a unifying empirical study of representative PETL approaches in the context of Vision Transformers (ViT) (Dosovitskiy et al., 2020). This includes Low-Rank Adaptation (LoRA) (Hu et al., 2021), Visual Prompt Tuning (VPT) (Jia et al., 2022), Adapter (Houlsby et al., 2019), and ten other approaches. We systematically tune their hyper-parameters to fairly compare their accuracy on the low-shot benchmark VTAB-1K. This includes learning rate, weight decay, and approach-specific parameters like the size of PETL parameters. Besides VTAB-1K, we examine PETL approaches on full-size downstream datasets such as CIFAR-100 (Krizhevsky et al., 2009), RESISC for remote sensing image scene classification (Cheng et al., 2017), and Clevr-Distance for depth classification with synthetic data (Zhai et al., 2019; Johnson et al., 2017). We also conduct a study on ImageNet (Deng et al., 2009) and its variants with domain shifts (Hendrycks et al., 2021a; Gao et al., 2023; Hendrycks et al., 2021b; Recht et al., 2019).

We summarize our key findings and extended analyses as follows.

Representative PETL approaches perform similarly on VTAB-1K, if properly implemented.

This includes methods previously considered less effective, such as fine-tuning the bias terms (Zaken et al., 2022) in the pre-trained backbone and methods originally proposed for NLP, like Adapter (Houlsby et al., 2019). Among all the hyper-parameters, we find the drop path rate (Huang et al., 2016) quite important. Ignoring it (*i.e.*, setting it to 0) significantly degrades the performance. Overall, PETL approaches consistently outperform linear probing and full fine-tuning on all 19 image classification tasks (each with 1,000 training examples) in VTAB-1K.

While similarly accurate on average, PETL approaches make different predictions. The above finding seems daunting: *if existing PETL approaches all perform similarly in terms of accuracy, do we learn anything useful beyond a single approach?* This is particularly worrisome given that they fine-tune the same backbone using the same downstream data. Fortunately, our analysis shows that different PETL methods learn differently from the same data, resulting in diverse prediction errors and confidence. We attribute this to their difference in inductive biases (Neyshabur et al., 2014) — they explicitly specify different parameters to be updated or inserted. This opens up the door to leverage their discrepancy for improvement, *e.g.*, through ensemble methods (Dietterich, 2000; Zhou, 2012) or co-training (Blum & Mitchell, 1998; Balcan et al., 2004) and we provide preliminary studies.

PETL is also effective in many-shot regimes. We apply PETL beyond the low-shot regime and find it effective even with ample downstream training data — PETL can be on par or outperform full fine-tuning. This suggests that varying a fraction of parameters of a properly chosen pre-trained backbone (e.g., pre-trained on ImageNet-21K (Dosovitskiy et al., 2020)) could already offer a sufficient effective capacity (Zhang et al., 2021) to reach a performant hypothesis for downstream tasks.

PETL is more robust than full fine-tuning to distribution shifts, but with weight-space ensembles, the observation is overturned. We also evaluate PETL’s robustness to distribution shifts, inspired by (Wortsman et al., 2022). We consider a CLIP backbone (Radford et al., 2021), known for its superior generalizability to distribution shifts, and apply PETL to fine-tune it with ImageNet-1K. We find that PETL preserves CLIP’s generalizability (e.g., to samples from ImageNet-Sketch or ImageNet-Rendition) better than full fine-tuning. This may not be surprising. What is interesting is that the weight-space ensembles (WiSE) between the fine-tuned and pre-trained models (Wortsman et al., 2022) apply to PETL as well to further improve the robustness without sacrificing the downstream accuracy. Nevertheless, full fine-tuning with WiSE can achieve even higher accuracy in both downstream and distribution shift data than PETL, suggesting a further research direction in PETL.

What lead to PETL’s success? We attempt to answer this fundamental question by analyzing the results of our study. On VTAB-1K with 19 tasks, we find two cases: 1) in some tasks, full fine-tuning outperforms linear probing, suggesting the need to update the backbone; 2) in some tasks, linear probing outperforms full fine-tuning, suggesting either the backbone is good enough or updating it risks over-fitting. The superior accuracy of PETL in both cases suggests that PETL acts as an *effective regularizer* during low-shot training. Still using VTAB but with ample training data, we find that for tasks in case 1), PETL is on par with full fine-tuning, suggesting that its regularization role does not prevent the fine-tuned model from learning sufficiently from the data. For tasks in case 2), PETL can surprisingly still outperform full fine-tuning, suggesting that it effectively transfers (or preserves) some useful pre-trained knowledge that full fine-tuning may wash away. In sum, PETL succeeds as a **high-capacity learner** equipped with an **effective regularizer**.

Contributions. Instead of chasing the leaderboard, we systematically understand existing approaches via a unifying study. Our contribution is thus not a technical novelty, but: (1) a **systematic framework** enabling consistent and reproducible evaluations of PETL methods; (2) a set of **empirical recommendations** on when and how to use PETL methods for practitioners; (3) **new insights for future research** including leveraging PETL’s prediction differences and exploring robust fine-tuning.

What do we not investigate? There are many aspects that one can ask about PETL. Our study does not consider computation-specific properties like memory usage and FLOPS.

2 BACKGROUND

2.1 LARGE PRE-TRAINED MODELS

Pre-trained models have become an indispensable part of modern AI development (Bommasani et al., 2021). Building upon neural networks with millions if not billions of parameters and gigantic amounts of training data, these large pre-trained models have led to groundbreaking results in various downstream tasks (Liang et al., 2024; Moor et al., 2023) and shown several emerging capabilities not observed previously (Khan et al., 2022; Li et al., 2024; Bommasani et al., 2021). For example, in computer vision, a Vision Transformer (ViT) (Dosovitskiy et al., 2020) trained with ImageNet-21K (around 14M images) leads to consistent gains v.s. a ViT trained with ImageNet-1K (around 1.3M images) (Dosovitskiy et al., 2020). ViTs pre-trained with millions of image-text pairs via a contrastive objective function (e.g., a CLIP-ViT model) (Radford et al., 2021; Cherti et al., 2023) show an unprecedented zero-shot capability and robustness to distribution shifts (Radford et al., 2021). In this paper, we focus on the ImageNet-21K-ViT and use the CLIP-ViT in a robustness study.

Vision Transformer (ViT).

We briefly review ViTs (Dosovitskiy et al., 2020), which are adapted from the Transformer-based models (Vaswani et al., 2017) in NLP. ViTs divide an image into a sequence of N fixed-sized patches and treat them like NLP tokens. Each patch is first embedded into a D -dimensional vector $\mathbf{x}_0^{(n)}$ with positional encoding. The sequence of vectors is then prepended with a “CLS” vector

$\mathbf{x}_0^{(\text{Class})}$ to generate the input $\mathbf{Z}_0 = [\mathbf{x}_0^{(\text{Class})}, \mathbf{x}_0^{(1)}, \dots, \mathbf{x}_0^{(N)}] \in \mathbb{R}^{D \times (1+N)}$ to a ViT, composed of M Transformer layers. We use super-/sub-script to index token/layer. The output of the “CLS” token $\mathbf{x}_M^{(\text{Class})}$ is used as the image representation.

Each of the ViT’s M Transformer layers consists of a multi-head self-attention (MSA) block, a multi-level perceptron (MLP) block, two Layer Normalization (LN) blocks (Ba et al., 2016), and two residual links. The m -th Transformer layer can be formulated as

$$\mathbf{Z}'_m = \text{MSA}(\text{LN}(\mathbf{Z}_{m-1})) + \mathbf{Z}_{m-1}, \quad (1)$$

$$\mathbf{Z}_m = \text{MLP}(\text{LN}(\mathbf{Z}'_m)) + \mathbf{Z}'_m, \quad (2)$$

where $\mathbf{Z}_{m-1} = [\mathbf{x}_{m-1}^{(\text{Class})}, \mathbf{x}_{m-1}^{(1)}, \dots, \mathbf{x}_{m-1}^{(N)}] \in \mathbb{R}^{D \times (1+N)}$ is the output of the preceding $(m-1)$ -th Transformer layer. The MLP is applied to each column vector of \mathbf{Z}'_m independently.

Without loss of generality, let us consider an MSA block with a single head. Given a generic input $\mathbf{Z} \in \mathbb{R}^{D \times (1+N)}$, this block first projects it into three matrices, Query \mathbf{Q} , Key \mathbf{K} , and Value \mathbf{V}

$$\mathbf{Q} = \mathbf{W}_Q \mathbf{Z}, \quad \mathbf{K} = \mathbf{W}_K \mathbf{Z}, \quad \mathbf{V} = \mathbf{W}_V \mathbf{Z}, \quad (3)$$

where $\mathbf{W}_{Q/K/V} \in \mathbb{R}^{D \times D}$ are projection matrices. The output of this block is then formulated as

$$\mathbf{V} \times \text{Softmax}\left(\frac{\mathbf{K}^\top \mathbf{Q}}{\sqrt{D}}\right) \in \mathbb{R}^{D \times (1+N)}. \quad (4)$$

2.2 PARAMETER EFFICIENT TRANSFER LEARNING (PETL)

Fine-tuning is arguably the most common way to tailor a pre-trained model for downstream tasks. As the size of pre-trained models gets larger, copying and updating all the parameters for one downstream task becomes inefficient. PETL has thus emerged as a promising paradigm.

PETL was originally developed in NLP (He et al., 2021a; Lester et al., 2021; He et al., 2022b; Mao et al., 2022; Sung et al., 2021; Zaken et al., 2022; Asai et al., 2022; Vu et al., 2022; Liu et al., 2022a; Su et al., 2022; Zhong et al., 2022) and has attracted increasing attention in vision (Jia et al., 2022; Chen et al., 2022b; Jie & Deng, 2022; Zhang et al., 2022b; Liu et al., 2022b; Lian et al., 2022). Existing approaches can generally be categorized into four groups: prompt-based, adapter-based, direct selective parameter tuning, and efficient selective parameter tuning. *We focus on visual recognition and compare representative PETL approaches applicable to ViTs.* During fine-tuning, all approaches learn a new FC layer for prediction.

Prompt-based approaches. Prompt-based learning emerged in NLP (Liu et al., 2023; Lialin et al., 2023). The core concept is to augment the input data with task-specific hints (prompts). **Visual Prompt Tuning (VPT)** (Jia et al., 2022) adapts such an idea to ViTs. Specifically, its deep version (VPT-Deep) prepends a set of soft prompts to the input tokens of each Transformer layer (*i.e.*, $\{\mathbf{Z}_m\}_{m=0}^{M-1}$) and only optimizes the prompts during fine-tuning. Other representative works in this category include (Yu et al., 2023b; Tu et al., 2023; Gu et al., 2023).

Adapter-based approaches. This category typically introduces additional trainable parameters (*e.g.*, an MLP block) to the frozen pre-trained model (Lialin et al., 2023). It was initially developed for multi-domain adaptation (Rebuffi et al., 2017; 2018) and continual learning (Rosenfeld & Tsotsos, 2018; Mai et al., 2022), and was subsequently extended to the NLP and vision domains to adapt Transformer-based models (Houlsby et al., 2019; Yu et al., 2023b).

We consider five popular adapter-based methods. **Houl. Adapter** (Houlsby et al., 2019) is the first adapter-based PETL approach. It inserts two Adapters — a two-layer bottleneck-structured MLP with a residual link — into each Transformer layer, one after the MSA block and the other after the MLP block. **Pfeif. Adapter** (Pfeiffer et al., 2021) inserts the Adapter solely after the MLP block, a strategy shown effective in recent studies (Hu et al., 2021). **AdaptFormer** (Chen et al., 2022b) inserts the Adapter in parallel with the original MLP block in a Transformer layer, different from the sequential design of Houl. and Pfeif. Adapter. One can view it as an ensemble, summing the task-specific features (by the Adapter) and the task-agnostic features (by the original MLP) to form \mathbf{Z}_m in Equation 2. **ConvPass** (Jie & Deng, 2022) introduces a convolutional-based bottleneck module (without a skip link) that explicitly encodes visual inductive biases: the 2D convolution is performed

over tokens of nearby patches. The module is inserted in parallel with the MSA and/or MLP block. **RepAdapter** (Luo et al., 2023) introduces a linear Adapter with group-wise transformations (Luo et al., 2022) and sequentially inserts two such modules after both MSA and MLP blocks.

Direct selective parameter tuning. This category selectively updates a subset of parameters of the pre-trained model, seen as a trade-off between full fine-tuning and linear probing. We consider three approaches. **BitFit** (Zaken et al., 2022) updates the bias terms, including those in the Q/K/V projections, the MLP blocks, the LN blocks, and the projection for patch embeddings. **LayerNorm** (Basu et al., 2023) updates the trainable parameters of the LN blocks in each Transformer layer. **DiffFit** (Xie et al., 2023) updates both the bias terms and the LN blocks and inserts learnable factors to scale the features after the MSA and the MLP blocks. Instead of updating parameters, **SSF** (Lian et al., 2022) linearly adapts intermediate features, motivated by feature modulation (Huang & Belongie, 2017; Perez et al., 2018). For an intermediate feature $\mathbf{Z} \in \mathbb{R}^{D \times (N+1)}$, SSF learns a D -dimensional scaling vector and a D -dimensional additive vector broadcasting to the tokens.

Efficient selective parameter tuning. Unlike the above category which directly updates parameters, this category learns *additive residuals* (e.g., $\Delta\mathbf{W}$) to the original parameters (e.g., \mathbf{W}). By injecting a low-rank constraint to the residuals, this category effectively reduces the learnable parameters. **LoRA** (Hu et al., 2021), arguably the most well-known approach, parameterizes the residuals by low-rank decomposition to update the Query/Value projection matrices $\mathbf{W}_{Q/V} \in \mathbb{R}^{D \times D}$. Concretely, to update a $\mathbf{W} \in \mathbb{R}^{D \times D}$ matrix, LoRA learns $\mathbf{W}_{\text{down}} \in \mathbb{R}^{r \times D}$ and $\mathbf{W}_{\text{up}} \in \mathbb{R}^{D \times r}$ with $r \ll D$, and forms the additive residual by $\Delta\mathbf{W} = \mathbf{W}_{\text{up}}\mathbf{W}_{\text{down}} \in \mathbb{R}^{D \times D}$. **Factor Tuning (FacT)** (Jie & Deng, 2023) extends the idea of matrix decomposition into tensor decomposition. It stacks the $D \times D$ learnable matrices in all the Transformer layers into a 3D tensor and learns an additive residual parameterized by the well-established Tensor-Train (TT) (Oseledets, 2011) and Tucker (TK) (De Lathauwer et al., 2000) formulations.

More detailed descriptions of ViT and PETL methods can be found in Appendix B.

2.3 RELATED WORK AND COMPARISON

The community-wide enthusiasm for PETL has led to multiple survey articles (Yu et al., 2023a; Xin et al., 2024; Han et al., 2024). Meanwhile, several empirical, integrative, and theoretical studies were presented, mostly based on NLP tasks, attempting to provide a holistic understanding. (He et al., 2021a; Mao et al., 2021) provided unified views to methodologically connect PETL approaches. (Chen et al., 2022a; Ding et al., 2023; He et al., 2021b) and (He et al., 2022a) empirically compared PETL approaches on NLP and vision tasks, respectively, while (Fu et al., 2023) offered a theoretical stability and generalization analysis. Accuracy-wise, (Chen et al., 2022a; Ding et al., 2023; He et al., 2021b) found that PETL is robust to over-fitting and quite effective in NLP tasks under low-data regimes. *This is, however, not the case for vision tasks: (He et al., 2022a) showed that representative PETL approaches like LoRA and Adapter cannot consistently outperform either full fine-tuning or linear probing.* In terms of why PETL works, (Fu et al., 2023) framed PETL as sparse fine-tuning and showed that it imposes a regularization by controlling stability; (Ding et al., 2023; He et al., 2022a) framed PETL as (subspace) optimization; (Ding et al., 2023) further discussed the theoretical principle inspired by optimal control.

Our study strengthens and complements the above studies and offers new insights. First, we compared over ten PETL approaches, more than any of the above. We carefully tune the hyper-parameters, aiming to reveal the faithful accuracy of each approach. This is particularly important for the vision community because there have been no unifying references for PETL accuracy; simple approaches like BitFit have often been reported as quite inferior; the effectiveness of other approaches was reported quite discrepant from the study in NLP. Second, we go beyond a *competition* perspective to investigate a *complementary* perspective of PETL approaches. We show that different PETL approaches offer effective base learners for model ensembles. Third, we go beyond downstream accuracy to investigate PETL’s effectiveness in maintaining out-of-distribution robustness. Fourth, we systematically analyze the results from low-shot and many-shot regimes and identify two distinct patterns among PETL, full fine-tuning, and linear probing, extending the understanding of PETL.

Method	Natural								Specialized								Structured										Overall Mean	Tunable Params
	CIFAR-100	Celeb101	DTD	Flowers102	Pets	SVHN	Sun397	Mean	Camelyon	EuroSAT	Resis45	Retinopathy	Mean	Clevr-Count	Clevr-Dist	DMLab	KITTI-Dist	dSpr-Loc	dSpr-On	sNOB-Azim	sNOB-Elev	Mean						
Linear Full	78.1	86.6	65.7	98.9	89.3	41.5	53.2	72.5	83.1	90.0	74.9	74.6	80.6	37.5	35.1	36.5	64.6	16.2	29.4	17.3	23.7	32.5	61.9	0				
VPT-Shallow	80.2	88.7	67.9	99.1	89.6	77.0	54.2	79.4	81.8	90.3	77.2	74.4	80.9	42.2	52.4	38	66.5	52.4	43.1	15.2	23.2	41.6	67.3	0.07				
VPT-Deep	84.8	91.5	69.4	99.1	91.0	85.6	54.7	81.8	86.4	94.9	84.2	73.9	84.9	79.3	62.4	48.5	77.9	80.3	56.4	33.2	43.8	60.2	75.6	0.43				
BitFit	86.5	90.5	70.3	98.9	91.0	91.2	54.2	82.6	86.7	95.0	85.3	75.5	85.6	77.2	63.2	51.2	79.2	78.6	53.9	30.1	34.7	58.5	75.6	0.1				
DiffFit	86.3	90.2	71.2	99.2	91.7	91.2	56.1	83.2	85.8	94.1	80.9	75.2	84.0	80.1	63.4	50.9	81.0	77.8	52.8	30.7	35.5	59.0	75.4	0.14				
LayerNorm	86.0	89.7	72.2	99.1	91.4	90.0	56.1	83.0	84.7	93.8	83.0	75.2	84.2	77.5	62.2	49.9	78.1	78.0	52.1	24.3	34.4	57.1	74.7	0.04				
SSF	86.6	89.8	68.8	99.1	91.4	91.2	56.5	82.8	86.1	94.5	83.2	74.8	84.7	80.1	63.6	53.0	81.4	85.6	52.1	31.9	37.2	60.6	76.0	0.21				
Pfeif. Adapter	86.3	91.5	72.1	99.2	91.4	88.5	55.7	83.0	86.2	95.5	85.3	76.2	85.8	83.1	65.2	51.4	80.2	83.3	56.6	33.8	41.1	61.8	76.9	0.67				
Houli. Adapter	84.3	92.1	72.3	98	91.7	90.0	55.4	83.2	88.7	95.3	86.5	75.2	86.4	82.9	63.6	53.8	79.6	84.4	54.3	34.2	44.3	62.1	77.2	0.77				
AdaptFormer	85.8	91.8	70.5	99.2	91.8	89.4	56.7	83.2	86.8	95.0	86.5	76.3	86.2	82.9	64.1	52.8	80.0	84.7	53.0	33.0	41.4	61.5	76.9	0.46				
RepAdapter	86.0	92.5	69.1	99.1	90.9	90.9	55.4	82.9	86.9	95.3	86.0	75.4	85.9	82.5	63.5	51.4	80.2	85.4	52.1	35.7	41.7	61.6	76.8	0.53				
Compass	85.0	92.1	72.0	99.3	91.3	90.8	55.9	83.5	87.7	95.8	85.9	75.9	86.3	82.3	65.2	53.8	78.1	86.5	55.3	38.6	45.1	63.1	77.6	0.49				
LoRA	85.7	92.6	69.8	99.1	90.5	88.5	55.5	82.6	87.5	94.9	85.9	75.7	86.0	82.9	63.9	51.8	79.9	86.6	47.2	33.4	42.5	61.0	76.5	0.55				
FacT_TT	85.8	91.8	71.5	99.3	91.1	90.8	55.9	83.4	87.7	94.9	85.0	75.6	85.8	83.0	64.0	49.0	79.3	85.8	53.1	32.8	43.7	61.3	76.8	0.13				
FacT_TK	86.2	92.5	71.8	99.1	90.1	91.2	56.2	83.4	85.8	95.5	86.0	75.7	85.8	82.7	65.1	51.5	78.9	86.7	53.1	27.8	40.8	60.8	76.6	0.23				
Relative Std Dev	0.81	1.13	1.78	0.34	0.54	1.82	1.24	0.54	1.20	0.59	1.95	0.83	0.94	2.67	1.50	3.22	1.37	4.11	4.46	11.02	9.30	2.70	1.09	-				

Table 1: Results on VTAB-1K (19 tasks from 3 groups). Based on the accuracy among PETL, linear probing, and full fine-tuning, we find two task groups (purple and orange), as discussed in section 6.




Figure 2: Ranking frequency of 15 methods (14 PETL + linear probing) for three groups in VTAB-1K. Element (i, j) is the number of times method i ranks j -th in each group. Methods are ordered by mean ranks (in brackets). The parameters column shows the # of trainable parameters in millions. More details are in Appendix C.

3 HOW DO PETL METHODS PERFORM IN LOW-SHOTS REGIME?

Pre-trained models are meant to ease downstream applications. One representative scenario is low-shot learning: supervised fine-tuning with a small number of examples per class. Indeed, low-shot learning has been widely used to evaluate PETL performance.

Dataset. VTAB-1K (Zhai et al., 2019) consists of 19 image classification tasks from three groups. The **Natural** group comprises natural images captured with standard cameras. The **Specialized** group contains images captured by specialist equipment for remote sensing and medical purposes. The **Structured** group evaluates the scene structure comprehension, such as object counting and 3D depth estimation. Following Zhai et al. (2019), we perform an 80/20 split on the **1000** training images in each task for hyperparameter tuning. The reported top1 accuracy is obtained by training on the 1000 training images and evaluating on the original test set.

Methods. We consider linear probing, full fine-tuning, and **14** PETL methods including **2** prompt-based (Jia et al., 2022), **5** adapter-based (Houlsby et al., 2019; Pfeiffer et al., 2021; Chen et al., 2022b; Jie & Deng, 2022; Luo et al., 2023), **4** Direct selective (Zaken et al., 2022; Basu et al., 2023; Xie et al., 2023; Lian et al., 2022), and **3** Efficient selective (Hu et al., 2021; Jie & Deng, 2023). Please refer to subsection 2.2 for details.

Setup. We employ the ViT-B/16 model (Dosovitskiy et al., 2020) pre-trained on ImageNet-21K (Deng et al., 2009) as the backbone. The prediction head is randomly initialized for each dataset. Images are resized to 224×224 . We systematically tune 1) learning rate, 2) weight decay, and 3) approach-specifics like the size of PETL parameters which are often left intact. **We set a cap for 3), $\leq 1.5\%$ of ViT-B/16.** We also turn the drop path rate (Huang et al., 2016) on (e.g., 0.1) or off (i.e., 0). A detailed hyperparameter search grid and additional training details are provided in Appendix A.1.

Results. As shown in Figure 1a and Table 1, PETL methods generally outperform both linear probing and full fine-tuning across datasets. Additionally, under fair hyper-parameter tuning, we surprisingly found that most PETL methods perform similarly as the relative standard deviations (divided by




Figure 3: Prediction similarity analysis: element (i, j) shows the percentage of samples that method i and method j have the same predictions. Although different methods achieve similar accuracy, they have diverse predictions. Details in Appendix C.




Figure 4: LoRA, Adapter and SSF finetuned on CIFAR100(VTAB-1K). (a): correct prediction overlap for the 5K most confident data. (b): wrong prediction overlap for the 5K least confident data. Details in Appendix C.

the means) in all three groups are quite low. Simple methods (*e.g.*, Bitfit) and PETL methods originally proposed for NLP (*e.g.*, LoRA and Adapter), which were previously reported as inferior due to un-optimized implementations and hyperparameter tuning, now demonstrate competitive performance with SOTA visual PETL methods. To understand the relative advantages of different approaches, we provide the ranking frequency of PETL methods across different groups in Figure 2, where the element (i, j) in each ranking matrix represents the frequency that method i ranks j^{th} in each group. Methods are ordered by their mean ranks (in brackets), and the parameters column indicates the number of trainable parameters in millions. In the natural group, simpler methods with fewer trainable parameters—such as DiffFit and Fact-TT—offer a cost-effective solution without compromising performance. Conversely, in the specialized and structured groups, methods with more parameters generally yield better performance. We hypothesize that this performance discrepancy arises from the **domain affinity** between the pre-trained domain (ImageNet) and the downstream domains. The natural group, sharing a stronger affinity with ImageNet, allows simpler methods like BitFit to adjust the features effectively. In contrast, the specialized and structured groups necessitate more complex methods with more trainable parameters to bridge the domain gap.

Recipes. In low-shot regimes, when the downstream data are similar to the pre-trained data, simple methods (*e.g.*, DiffFit) with decent accuracy and much fewer parameters are preferred—aligning with Occam’s razor. When there is a substantial domain gap, complex methods with higher accuracy (but more parameters) become competitive. Learning with low-shot data is prone to over-fitting. We find that if the drop path rate — which stochastically drops a transformer block per sample (Huang et al., 2016) — is set not as default (*i.e.*, nonzero), all the methods can benefit from such a regularization, as shown in Figure 10 in the Appendix.

4 PETL APPROACHES OFFER COMPLEMENTARY INFORMATION

The previous section demonstrates that all PETL methods perform similarly across various domains. Given that different PETL methods are trained on the same downstream data using the same backbone and achieve comparable accuracy, one might expect them to learn similar knowledge from the data, resulting in similar predictions. Contrary to this expectation, our findings below reveal that different PETL methods acquire **distinct** and **complementary** knowledge from the same downstream data, even when built upon the same backbone, leading to diverse predictions.

We start by analyzing their prediction similarity on the same dataset in VTAB-1K. It is expected that their predictions are similar for datasets with very high accuracy, such as Flowers102 (avg 99.1%) and Caltech101 (avg 91.4%). Beyond them, we find that most PETL methods show diverse predictions in other datasets in VTAB-1K. Figure 3 shows the prediction similarities between 14 PETL methods in DTD, Retinopathy, and DMLab, which belong to natural, specialized, and structured groups, respectively. In DTD and Retinopathy, most methods differ in about 20% of their predictions, while in DMLab, this difference increases to approximately 35%, even though they achieve similar accuracies. This prediction diversity may be attributed to the different *inductive biases* (Neyshabur et al., 2014) of PETL methods — they explicitly select specific parameters to update or insert different modules at various locations within the model. More analyses and details about Figure 3 are offered in Appendix C.

Such diverse predictions across methods open up the possibility of leveraging their heterogeneity for further improvement. The most straightforward approach is ensemble (Gontijo-Lopes et al., 2021), *e.g.*, average logits over methods. Figure 5 demonstrates the ensemble performance gain over all the PETL methods in each dataset, where we use the worst PETL method as the baseline. Due to the diverse predictions across methods, the ensemble can provide consistent gain.

Also, we analyze if PETL methods make similar correct predictions for high-confidence samples and similar mistakes for low-confidence samples. Figure 4 shows the correct prediction overlap for the 5K most confident samples (per method) and the wrong prediction overlap for the 5K least confident samples (per method). For demonstration purposes, we select one method from each PETL category (LoRA, Adapter, SSF) and they are fine-tuned on CIFAR-100 in VTAB-1K. Methods within the same category also show diverse predictions (Appendix C). Since they make different predictions in both high and low-confidence regimes, this paves the way for new possibilities of using different PETL methods to generate **diverse pseudo-labels** for semi-supervised learning (Yang et al., 2022), domain adaptation (Farahani et al., 2021), and transfer learning (Zhuang et al., 2020).




Figure 5: Ensemble (avg logits) provides consistent gain on most datasets thanks to the diverse prediction. Details in Appendix C.

5 HOW DO PETL METHODS PERFORM IN MANY-SHOT REGIME?

Recent works in NLP Chen et al. (2022a) have indicated that PETL methods may not perform as competitively as full fine-tuning when data is abundant. We thus aim to investigate PETL’s performance in many-shot regimes by addressing the following questions: (1) Should we use PETL or full fine-tuning when data is sufficient? (2) How should we adjust the number of trainable parameters for PETL methods in many-shot regimes?

Dataset. We select one representative dataset from each of the natural, specialized, and structured groups in VTAB: (1) CIFAR-100 Krizhevsky et al. (2009), a natural image dataset comprising 50K training images across 100 classes; (2) RESISC Cheng et al. (2017), a remote sensing dataset for scene classification with 25.2K training samples across 45 classes; and (3) Clevr-Distance Zhai et al. (2019); Johnson et al. (2017), a synthetic image dataset for predicting the depth of the closest object from the camera with 6 depth classes and 70K samples. The reported results are obtained by training on the **full** training set and evaluating on the original test set.

Setup. The model setup follows the VTAB-1K experiment. More details about setup and hyperparameter search are provided in Appendix A.

Results. In many-shot regimes, with sufficient downstream data, full fine-tuning may catch up and eventually outperform PETL methods. However, from Figure 6, we found that even in many-shot regimes, PETL can achieve **comparable results with full fine-tuning**, even just using 2% of fine-tuning parameters. (The performance gain, however, quickly **diminishes and plateaus after 5%** of tunable parameters.) By comparing the results on the domain-close CIFAR-100 and domain-different RESISC and Clevr, we have some further observations. On the one hand, downstream tasks with larger domain gaps suggest the need to update, perhaps many, parameters to obtain high accuracy. With sufficient downstream data, full fine-tuning is less prone to over-fitting and indeed attains a high accuracy. But interestingly, PETL methods, with only 2 ~ 5% of tunable parameters, achieve similar accuracy, suggesting that its design principle does offer sufficient effective capacity for the model to learn Zhang et al. (2021). On the other hand, downstream tasks with smaller domain gaps suggest that the pre-trained model had learned sufficient knowledge about them; fully fine-tuning it thus risks washing such knowledge away. In fact, we found that PETL notably outperforms full fine-tuning on CIFAR-100, suggesting it as a more robust *transfer learning* algorithm for downstream tasks.

Recipes. In many-shot regimes, PETL methods with sufficient parameters (2 ~ 5%) appear more favorable than full fine-tuning and linear probing. On the one hand, they achieve comparable and even




Figure 6: PETL accuracy in many-shot regimes, with different parameter sizes (X-axis) on three datasets from different domains. Even 2%-5% trainable parameters allow the models to have sufficient capacity to learn from full data. (Details are in Appendix C)

better accuracy than full fine-tuning. On the other hand, the tunable parameters remain manageable. The parameter efficiency of PETL also often implies less training GPU memory usage and training time, making PETL methods a favorable alternative in many-shot regimes. For a downstream domain that is close to the pre-training domain, PETL shows much pronounced *transferability*. For a downstream domain that is quite different, the limited tunable parameters (controversially, 2 ~ 5% already amounts to a few million) already allow the model to learn sufficiently.

6 WHY DO PETL METHODS WORK? ¹

Putting together section 3 and section 5, we identify two distinct patterns regarding the performance among linear probing, full fine-tuning, and PETL. Within 19 VTAB-1K tasks, we see: (1) Full fine-tuning outperforms linear probing. As linear probing reflects the pre-trained feature quality for downstream tasks, case (1) suggests the necessity to update the backbone to close the gap between pre-trained and downstream domains. (2) Linear probing surpasses full fine-tuning, suggesting the pre-trained features are good enough (at least in a low-shot scenario). Recklessly updating them may risk over-fitting. Figure 7 (a-b) summarizes the low-shot accuracy comparison based on the categorization above; each line corresponds to one task. Linear probing, PETL, and fine-tuning are located in order, from left to right, to reflect their tunable parameter sizes. PETL’s superiority in both cases showcases its **capacity** to learn and its **regularization role** to prevent over-fitting.

We also draw the many-shot accuracy in Figure 7 (c-d) based on the same categorization: RESISC and Clevr in case (1), and CIFAR-100 in case (2). In the many-shot setting, full fine-tuning consistently outperforms linear probing, which seems to suggest *no more risk of over-fitting*. However, on CIFAR-100 (Figure 7 (d)), we again see a noticeable gap between PETL and full fine-tuning, just like in Figure 7 (b). Such a concave shape reminds us of the long-standing under-fitting-over-fitting curve, suggesting that even with sufficient downstream data, full fine-tuning still risks over-fitting.

Taking into account PETL’s comparable performance to full fine-tuning on RESISC and Clevr with large domain gaps, we conclude — PETL succeeds as a **high-capacity** learner equipped with an **effective regularizer**; the two roles trade-offs well such that PETL can excel in both low-affinity and high-affinity domains under both low-shot and many-shot settings.

7 ARE PETL METHODS MORE ROBUST TO DISTRIBUTION SHIFTS?

Large pre-trained models such as CLIP Radford et al. (2021) and ALIGN Jia et al. (2021) have demonstrated unprecedented accuracy across a range of data distributions when performing zero-shot inference. However, recent studies Wortsman et al. (2022); Radford et al. (2021) have shown that fine-tuning on downstream data, while significantly boosting performance on the target distribution, often compromises the model’s robustness to distribution shifts. Given that PETL only updates a

¹Our intention is not to offer a definitive conclusion about why PETL works. As discussed in subsection 2.3, there is currently no universally agreed-upon explanation for the effectiveness of PETL. We hope our empirical findings will contribute to the ongoing efforts to understand the underlying principles of PETL methods.




Figure 7: (a): VTAB-1K tasks in case (1), PETL > full > linear. (b) VTAB-1K tasks in case (2), PETL > linear > full. (c) RESISC & Clevr in case (1) with enough data, PETL \approx full > linear. (d) CIFAR in case (2) with enough data, PETL > full > linear. Within each figure, left for linear, middle for PETL, and right for full. More details are in Appendix C.

	Full	BitFit	Layer-Norm	Houl. Adapter	Adapt-Former	Rep-Adapter	Convpass	LoRA	FacT_TK
100-shot ImageNet	75.0	75.27	74.8	75.0	75.6	76.5	76.3	76.6	74.7
Avg. distribution shift Acc	42.5	55.4 (12.9) \uparrow	55.9 (13.4) \uparrow	56.9 (14.4) \uparrow	56.1 (13.6) \uparrow	56.2 (13.7) \uparrow	54.7 (12.2) \uparrow	55.9 (13.4) \uparrow	56.1 (13.6) \uparrow

Table 2: The “Avg. distribution shift Acc” denotes the average performance of ImageNet-(V2, S, R, A) evaluated on the CLIP model fine-tuned on ImageNet. (\uparrow) indicates the gain over full fine-tuning.

limited number of parameters in the model, we investigate whether PETL can offer a more robust alternative to full fine-tuning for pre-trained models.

Dataset. We use 100-shot ImageNet-1K as our target distribution, with each class containing 100 images. Following Wortsman et al. (2022), we consider 4 natural distribution shifts from ImageNet: **ImageNet-V2** Recht et al. (2019), a new ImageNet test set collected with the original labeling protocol; **ImageNet-R** Hendrycks et al. (2021a), renditions for 200 ImageNet classes; **ImageNet-S** Gao et al. (2023), sketch images for 1K ImageNet classes; **ImageNet-A** Hendrycks et al. (2021b), a test set of natural images misclassified by a ResNet-50 He et al. (2015) for 200 ImageNet classes.

Setup. We focus on the CLIP ViT-B/16 model, which comprises a visual encoder and a text encoder, pre-trained via contrastive learning on image-text pairs. Following Wortsman et al. (2022), we add an FC layer as the prediction head with zero-initialized bias and initialize weights using the class label text embedded by the text encoder. Subsequently, we discard the text encoder and apply PETL methods to the visual encoder, fine-tuning only the PETL modules and the head. More details about the CLIP model and experiment setup can be found in Appendix A.1.

Results. As shown in Table 2, while some PETL methods may not surpass full fine-tuning on the target distribution, they consistently demonstrate more robust performance on distribution shift data. This is likely because PETL updates only a small fraction of the parameters, thus preserving the robust features of the foundation models. Given the similar target distribution performance, should we blindly use PETL methods for more robustness?

Weight-space ensembles (WiSE) for PETL. WiSE (Wortsman et al., 2022), which linearly interpolates the full fine-tuned and original models, is a popular fine-tuning approach to enhance robustness. We explore whether WiSE can enhance the robustness of PETL. To apply WiSE to PETL, we first linearly interpolate the prediction head with a mixing coefficient α . For direct selective tuning methods (e.g. BitFit), we directly interpolate with the original model. Since most Adapter-based methods have residual connections, we can multiply the adapter modules with α to control their strengths. A similar approach can be applied to efficient selective methods (e.g. LoRA) as they learn additive residuals to the original parameters. As shown in Figure 1b (more results in Appendix C), WiSE improves both fine-tuning and distribution shift performance of PETL methods. Interestingly, even though full fine-tuning is generally less robust than PETL methods, applying WiSE allows it to achieve better performance in both target distribution and distribution shift data, which suggests a promising research direction for robust PETL.

8 CONCLUSION

We conduct a unifying empirical study of parameter-efficient fine-tuning (PETL), an emerging topic in the large model era. We have several new insights and implications, including PETL methods’ complementary expertise, suitable application regimes, and robustness to domain shifts. We expect our study to open new research directions and serve as a valuable user guide in practice.

9 REPRODUCIBILITY STATEMENT

We have made extensive efforts to ensure the reproducibility of our results. All 14 PETL algorithms and two baseline models, along with data preprocessing routines and data loaders for all datasets—including 19 low-shot datasets, 3 many-shot (full) datasets, and 5 robustness-related datasets—are implemented within a systematic and extensible framework. This framework is designed to facilitate the easy addition of new PETL methods and datasets, modification of backbones, and incorporation of additional scenarios, serving as a convenient tool for future research. Detailed explanations of our implementations, raw results for all experiments, and commands to reproduce the results are thoroughly documented in the README file. We provide the anonymous source code in the supplementary material.

10 ETHICS STATEMENT

Our study provides a unifying study of PETL in visual recognition. We expect it to serve as a valuable practical user guide to benefit society. Specifically, fine-tuning large models needs significant computation. A unifying study of PETL will ease end-users to apply more parameter-efficient and computation-efficient ways for fine-tuning. To our knowledge, our paper does not introduce any additional negative societal impacts compared to existing papers on PETL.

REFERENCES

- Armen Aghajanyan, Sonal Gupta, and Luke Zettlemoyer. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 7319–7328, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.568. URL <https://aclanthology.org/2021.acl-long.568>.
- Akari Asai, Mohammadreza Salehi, Matthew E Peters, and Hannaneh Hajishirzi. Attentional mixtures of soft prompt tuning for parameter-efficient multi-task knowledge sharing. *arXiv preprint arXiv:2205.11961*, 2022.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Maria-Florina Balcan, Avrim Blum, and Ke Yang. Co-training and expansion: Towards bridging theory and practice. *Advances in neural information processing systems*, 17, 2004.
- Samyadeep Basu, Daniela Massiceti, Shell Xu Hu, and Soheil Feizi. Strong baselines for parameter efficient few-shot fine-tuning. *arXiv preprint arXiv:2304.01917*, 2023.
- Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pp. 92–100, 1998.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Guanzheng Chen, Fangyu Liu, Zaiqiao Meng, and Shangsong Liang. Revisiting parameter-efficient tuning: Are we really there yet? *arXiv preprint arXiv:2202.07962*, 2022a.
- Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptorformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems*, 35:16664–16678, 2022b.
- Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017.

- Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2818–2829, 2023.
- Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. A multilinear singular value decomposition. *SIAM journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pp. 1–15. Springer, 2000.
- Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5(3):220–235, 2023.
- Xiaohan Ding, Xiangyu Zhang, Ningning Ma, Jungong Han, Guiguang Ding, and Jian Sun. Repvgg: Making vgg-style convnets great again. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 13733–13742, 2021.
- 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 *International Conference on Learning Representations*, 2020.
- Abolfazl Farahani, Sahar Voghoei, Khaled Rasheed, and Hamid R Arabnia. A brief review of domain adaptation. *Advances in data science and information engineering: proceedings from ICDATA 2020 and IKE 2020*, pp. 877–894, 2021.
- Zihao Fu, Haoran Yang, Anthony Man-Cho So, Wai Lam, Lidong Bing, and Nigel Collier. On the effectiveness of parameter-efficient fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 12799–12807, 2023.
- Shanghua Gao, Zhong-Yu Li, Ming-Hsuan Yang, Ming-Ming Cheng, Junwei Han, and Philip Torr. Large-scale unsupervised semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):7457–7476, June 2023. ISSN 1939-3539. doi: 10.1109/tpami.2022.3218275. URL <http://dx.doi.org/10.1109/TPAMI.2022.3218275>.
- Raphael Gontijo-Lopes, Yann Dauphin, and Ekin Dogus Cubuk. No one representation to rule them all: Overlapping features of training methods. In *International Conference on Learning Representations*, 2021.
- Jindong Gu, Zhen Han, Shuo Chen, Ahmad Beirami, Bailan He, Gengyuan Zhang, Ruotong Liao, Yao Qin, Volker Tresp, and Philip Torr. A systematic survey of prompt engineering on vision-language foundation models. *arXiv preprint arXiv:2307.12980*, 2023.
- Zeyu Han, Chao Gao, Jinyang Liu, Sai Qian Zhang, et al. Parameter-efficient fine-tuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*, 2024.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*, 2021a.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.
- Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jia-Wei Low, Lidong Bing, and Luo Si. On the effectiveness of adapter-based tuning for pretrained language model adaptation. *arXiv preprint arXiv:2106.03164*, 2021b.
- Xuehai He, Chunyuan Li, Pengchuan Zhang, Jianwei Yang, and Xin Eric Wang. Parameter-efficient fine-tuning for vision transformers. *arXiv preprint arXiv:2203.16329*, 3, 2022a.

- Yun He, Steven Zheng, Yi Tay, Jai Gupta, Yu Du, Vamsi Aribandi, Zhe Zhao, YaGuang Li, Zhao Chen, Donald Metzler, et al. Hyperprompt: Prompt-based task-conditioning of transformers. In *International Conference on Machine Learning*, pp. 8678–8690. PMLR, 2022b.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization, 2021a.
- Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples, 2021b.
- 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 *International Conference on Machine Learning*, pp. 2790–2799. PMLR, 2019.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.
- Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pp. 646–661. Springer, 2016.
- Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pp. 1501–1510, 2017.
- Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2704–2713, 2018.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pp. 4904–4916. PMLR, 2021.
- 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*, pp. 709–727. Springer, 2022.
- Shibo Jie and Zhi-Hong Deng. Convolutional bypasses are better vision transformer adapters. *arXiv preprint arXiv:2207.07039*, 2022.
- Shibo Jie and Zhi-Hong Deng. Fact: Factor-tuning for lightweight adaptation on vision transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 1060–1068, 2023.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2901–2910, 2017.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pp. 4171–4186, 2019.
- Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM computing surveys (CSUR)*, 54(10s):1–41, 2022.
- Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2661–2671, 2019.

- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Ananya Kumar, Aditi Raghunathan, Robbie Matthew Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations*, 2021.
- 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*, pp. 3045–3059. Association for Computational Linguistics, 2021.
- Chunyu Li, Heerad Farkhor, Rosanne Liu, and Jason Yosinski. Measuring the intrinsic dimension of objective landscapes. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=ryup8-WCW>.
- Chunyu Li, Zhe Gan, Zhengyuan Yang, Jianwei Yang, Linjie Li, Lijuan Wang, Jianfeng Gao, et al. Multimodal foundation models: From specialists to general-purpose assistants. *Foundations and Trends® in Computer Graphics and Vision*, 16(1-2):1–214, 2024.
- Vladislav Lialin, Vijeta Deshpande, and Anna Rumshisky. Scaling down to scale up: A guide to parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.15647*, 2023.
- Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. *Advances in Neural Information Processing Systems*, 35: 109–123, 2022.
- Yuxuan Liang, Haomin Wen, Yuqi Nie, Yushan Jiang, Ming Jin, Dongjin Song, Shirui Pan, and Qingsong Wen. Foundation models for time series analysis: A tutorial and survey. *arXiv preprint arXiv:2403.14735*, 2024.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35, 2023.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 61–68, 2022a.
- Yen-Cheng Liu, Chih-Yao Ma, Junjiao Tian, Zijian He, and Zsolt Kira. Polyhistor: Parameter-efficient multi-task adaptation for dense vision tasks. *arXiv preprint arXiv:2210.03265*, 2022b.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Yan Wang, Liujuan Cao, Yongjian Wu, Feiyue Huang, and Rongrong Ji. Towards lightweight transformer via group-wise transformation for vision-and-language tasks. *IEEE Transactions on Image Processing*, 31:3386–3398, 2022.
- Gen Luo, Minglang Huang, Yiyi Zhou, Xiaoshuai Sun, Guannan Jiang, Zhiyu Wang, and Rongrong Ji. Towards efficient visual adaption via structural re-parameterization. *arXiv preprint arXiv:2302.08106*, 2023.
- Zheda Mai, Ruiwen Li, Jihwan Jeong, David Quispe, Hyunwoo Kim, and Scott Sanner. Online continual learning in image classification: An empirical survey. *Neurocomputing*, 469:28–51, 2022.
- Yuning Mao, Lambert Mathias, Rui Hou, Amjad Almahairi, Hao Ma, Jiawei Han, Wen-tau Yih, and Madian Khabza. Unipelt: A unified framework for parameter-efficient language model tuning. *arXiv preprint arXiv:2110.07577*, 2021.
- Yuning Mao, Lambert Mathias, Rui Hou, Amjad Almahairi, Hao Ma, Jiawei Han, Scott Yih, and Madian Khabza. UniPELT: A unified framework for parameter-efficient language model tuning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 6253–6264. Association for Computational Linguistics, 2022.

- Michael Moor, Oishi Banerjee, Zahra Shakeri Hossein Abad, Harlan M Krumholz, Jure Leskovec, Eric J Topol, and Pranav Rajpurkar. Foundation models for generalist medical artificial intelligence. *Nature*, 616(7956):259–265, 2023.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.
- Ivan V Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5): 2295–2317, 2011.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- JonLayer Noras Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapterfusion: Non-destructive task composition for transfer learning. In *16th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2021*, pp. 487–503. Association for Computational Linguistics (ACL), 2021.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *Advances in neural information processing systems*, 30, 2017.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Efficient parametrization of multi-domain deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8119–8127, 2018.
- Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet?, 2019.
- Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021.
- Amir Rosenfeld and John K Tsotsos. Incremental learning through deep adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 42(3):651–663, 2018.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- Yusheng Su, Xiaozhi Wang, Yujia Qin, Chi-Min Chan, Yankai Lin, Huadong Wang, Kaiyue Wen, Zhiyuan Liu, Peng Li, Juanzi Li, et al. On transferability of prompt tuning for natural language processing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 3949–3969, 2022.
- Yi-Lin Sung, Varun Nair, and Colin A Raffel. Training neural networks with fixed sparse masks. *Advances in Neural Information Processing Systems*, 34:24193–24205, 2021.
- 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*, pp. 1–9, 2015.
- Cheng-Hao Tu, Zheda Mai, and Wei-Lun Chao. Visual query tuning: Towards effective usage of intermediate representations for parameter and memory efficient transfer learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7725–7735, 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

- Tu Vu, Brian Lester, Noah Constant, Rami Al-Rfou', and Daniel Cer. SPoT: Better frozen model adaptation through soft prompt transfer. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5039–5059. Association for Computational Linguistics, 2022.
- Benyou Wang, Yuxin Ren, Lifeng Shang, Xin Jiang, and Qun Liu. Exploring extreme parameter compression for pre-trained language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=RftryyYyjiG>.
- Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems*, pp. 10506–10518, 2019.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pp. 38–45, 2020.
- Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7959–7971, 2022.
- Enze Xie, Lewei Yao, Han Shi, Zhili Liu, Daquan Zhou, Zhaoqiang Liu, Jiawei Li, and Zhenguo Li. DiffFit: Unlocking transferability of large diffusion models via simple parameter-efficient fine-tuning. *arXiv preprint arXiv:2304.06648*, 2023.
- Yi Xin, Siqi Luo, Haodi Zhou, Junlong Du, Xiaohong Liu, Yue Fan, Qing Li, and Yuntao Du. Parameter-efficient fine-tuning for pre-trained vision models: A survey. *arXiv preprint arXiv:2402.02242*, 2024.
- Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu. A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(9):8934–8954, 2022.
- Bruce XB Yu, Jianlong Chang, Haixin Wang, Lingbo Liu, Shijie Wang, Zhiyu Wang, Junfan Lin, Lingxi Xie, Haojie Li, Zhouchen Lin, et al. Visual tuning. *ACM Computing Surveys*, 2023a.
- Bruce XB Yu, Jianlong Chang, Haixin Wang, Lingbo Liu, Shijie Wang, Zhiyu Wang, Junfan Lin, Lingxi Xie, Haojie Li, Zhouchen Lin, et al. Visual tuning. *arXiv preprint arXiv:2305.06061*, 2023b.
- Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 1–9, 2022.
- Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruysen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, et al. A large-scale study of representation learning with the visual task adaptation benchmark. *arXiv preprint arXiv:1910.04867*, 2019.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.
- Jinnian Zhang, Houwen Peng, Kan Wu, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Minivit: Compressing vision transformers with weight multiplexing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12145–12154, 2022a.
- Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. Neural prompt search. *arXiv preprint arXiv:2206.04673*, 2022b.
- Qihuang Zhong, Liang Ding, Juhua Liu, Bo Du, and Dacheng Tao. Panda: Prompt transfer meets knowledge distillation for efficient model adaptation. *arXiv preprint arXiv:2208.10160*, 2022.

Zhi-Hua Zhou. *Ensemble methods: foundations and algorithms*. CRC press, 2012.

Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020.

APPENDIX

We provide details that are omitted from the main paper.

- Appendix A: Experiment and dataset details
- Appendix B: Detailed descriptions of ViT and compared methods.
- Appendix C: Additional results not presented in main paper
- Appendix D: Discussion about further impacts of this work

A EXPERIMENT AND DATASET DETAILS

A.1 EXPERIMENT DETAILS

VTAB-1K We employ AdamW optimizer Loshchilov & Hutter (2018) with a batch size of 64 and utilize the cosine decay learning rate scheduler. We train all methods with 100 epochs. The learning rate is tuned from [1e-3, 1e-2] and weight decay from [1e-4, 1e-3]. The method-specific hyperparameter searching grip is shown in Table 3, along with the tunable parameter ranges (in millions). Since most method-specific hyperparameters affect the number of tunable parameters in the PETL methods, we set a cap on the tunable parameters for each PETL method to be less than or equal to **1.5%** of the total parameters in ViT-B/16, which is approximately equal to the number of parameters in the Query, Key, and Value matrices of a single MSA block. Consistent with the original VTAB-1k paper (Zhai et al., 2019), most PETL studies Jie & Deng (2023; 2022); Luo et al. (2023); Jia et al. (2022); Zhang et al. (2022b); Luo et al. (2023); Lian et al. (2022) don’t apply data augmentation as it’s challenging to identify a set of augmentations that uniformly benefits all 19 datasets². To ensure that our results are directly comparable and that any performance differences are attributable to the methods themselves rather than data augmentation, we don’t apply data augmentation.

Many-shot We also employ AdamW optimizer with a batch size of 64 and a cosine decay learning rate scheduler. The learning rate is tuned from [5e-4, 1e-3] and weight decay keeps the same range of [1e-4, 1e-3]. We apply horizontal flipping for CIFAR100, horizontal and vertical flipping for Resisc, and no augmentation for Clevr. We train all methods with 40 epochs.

Robustness Model CLIP models are trained on image-caption pairs collected from the web. Given a dataset of such pairs $\{(x_1, s_1), \dots, (x_B, s_B)\}$, these models learn an image encoder g and a text encoder h that aim to maximize the similarity $\langle g(x_i), h(s_i) \rangle$ between matching image and caption embeddings while minimizing it for non-matching pairs. For zero-shot classification, the models predict the class of an input image x from a set of k class names $C = \{c_1, \dots, c_k\}$ by matching x with captions derived from these class names. Specifically, for each class c_j , a caption is formulated as $s_j = \text{“a photo of a } c_j\text{”}$. The predicted class is then determined by selecting the one whose caption embedding has the highest similarity with the image embedding: $\hat{y} = \arg \max_j \langle g(x), h(s_j) \rangle$. Alternatively, a weight matrix $W_{\text{zero-shot}} \in \mathbb{R}^{d \times k}$ can be constructed, where each column is the embedding $h(s_j)$ corresponding to class c_j . The model’s output scores for each class are then computed as $f(x) = g(x)^\top W_{\text{zero-shot}}$. To generate $W_{\text{zero-shot}}$, we ensemble the 80 prompts provided by CLIP at <https://github.com/openai/CLIP>.

Robustness Setup In section 3 and section 5, the fine-tuning train set and test set are from the same data distribution. In section 7, we fine-tune the CLIP model using ImageNet-1K training data (100 shots) and subsequently evaluate the fine-tuned model not only on the test set of ImageNet-1K but also on four additional datasets with distribution shifts: ImageNet-V2, ImageNet-R, ImageNet-S, and ImageNet-A, as shown in Figure 8. Following Wortsman et al. (2022), we set a small learning rate as 3e-5 and weight decay as 5e-3. We use a strong data augmentation following Zhang et al. (2022b).

Computation We used a workstation with eight NVIDIA RTX 6000 Ada GPUs, two AMD EPYC 9554 64-Core Processors, and 800GB of RAM.

²To demonstrate it, we apply simple data augmentations (RandomResizedCrop and RandomHorizontalFlip) on three datasets in each group, as shown in Table 4.

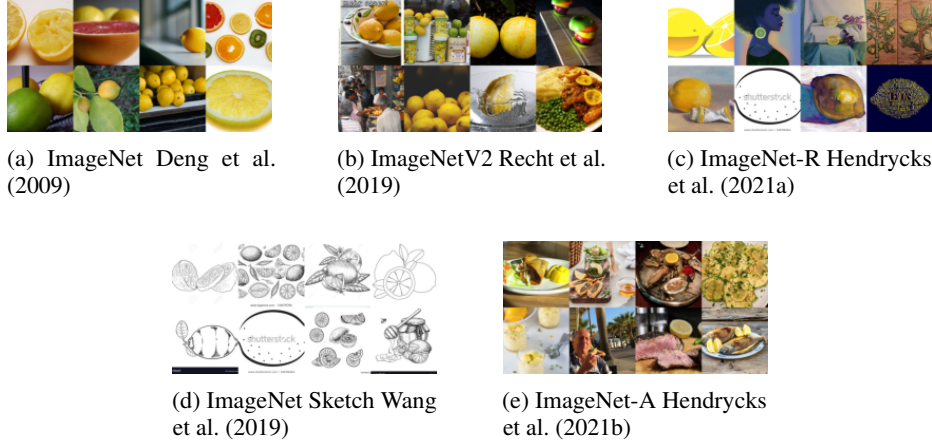


Figure 8: Samples of the class lemon, from the fine-tuned dataset ImageNet and distribution shifts datasets (ImageNet-V2, ImageNet-R, ImageNet-S, and ImageNet-A). The CLIP model is fine-tuned with PETL on ImageNet and evaluated on distribution shifts datasets to measure the robustness of fine-tuned models. The figures are modified based on Wortsman et al. (2022).

Method	Hyperparameters	#Params (M)
VPT-Shallow	Prompt Number: [5, 10, 50, 100, 200]	0.0003 \sim 0.153
VPT-Deep	Prompt Number: [5, 10, 50, 100]	0.046 \sim 0.921
BitFit	N/A	0.102
DiffFit	N/A	0.140
LayerNorm	N/A	0.038
SSF	N/A	0.205
Pfeif. Adapter	Adapter Scale Factor: [0.01, 0.1, 1, 10] Adapter Bottleneck: [4, 8, 16, 32]	0.082 \sim 0.599
Houl. Adapter	Adapter Scale Factor: [0.01, 0.1, 1, 10] Adapter Bottleneck: [4, 8, 16, 32]	0.165 \sim 1.198
AdaptFormer	Adapter Scale Factor: [0.05, 0.1, 0.2] Adapter Bottleneck: [4, 16, 32]	0.082 \sim 0.599
RepAdapter	RepAdapter Scale Factor: [0.1, 0.5, 1, 5, 10] RepAdapter Bottleneck: [8, 16, 32]	0.239 \sim 0.903
Convpass	Convpass Scale Factor: [0.01, 0.1, 1, 10, 100] Convpass Bottleneck: [8, 16] Convpass Xavier Init: [True, False]	0.327 \sim 0.664
LoRA	LoRA Bottleneck: [1, 8, 16, 32]	0.036 \sim 1.179
FacT_TT	FacT Scale Factor: [0.01, 0.1, 1, 10, 100] FacT Bottleneck: [8, 16, 32]	0.021 \sim 0.196
FacT_TK	FacT Bottleneck: [16, 32, 64] FacT Scale Factor: [0.01, 0.1, 1, 10, 100]	0.030 \sim 0.369

Table 3: Methods-specific hyperparameter searching grip for VTAB-1K experiment.

A.2 DATASET DETAILS

VTAB-1K The processed VTAB-1K can be downloaded from our official code base to ensure reproducibility.

Many-shot Datasets We perform 90/10 train-val split for CIFAR-100, RESISC and Clevr-Distance. The split details are provided in our code base for reproducibility. We apply horizontal flipping for CIFAR100, horizontal and vertical flipping for Resisc, and no augmentation for Clevr. All data are normalized by ImageNet mean and standard deviation.

			Linear	Full	VPT-Shallow	VPT-Deep	BitFit	DiffFit	LayerNorm	SSF	Picfit Adapter	Houl Adapter	Adapt-Former	Rep-Adapter	Convpass	LoRA	FacT_TT	FacT_TK
Catech101		Simple DA	84.4	76.8	84.9	84.8	83.8	85.7	85.8	86.1	87.4	86.0	84.9	86.4	85.2	86.8	85.5	86.0
		Default	86.6	89.9	88.7	91.5	90.5	90.2	89.7	89.8	91.5	92.1	91.8	92.5	92.1	92.6	91.8	92.5
		Δ	-2.2	-13.1	-3.8	-6.7	-4.5	-3.9	-3.7	-4.1	-6.1	-6.9	-6.1	-6.9	-6.9	-5.8	-6.3	-6.5
Natural	DTD	Simple DA	67.5	57.8	69.0	71.1	70.7	73.7	73.5	68.4	72.7	72.6	70.6	71.5	71.8	73.0	72.2	71.9
		Default	65.7	61.9	67.9	69.4	70.3	71.2	72.2	68.8	72.1	72.3	70.5	69.1	72.0	69.8	71.5	71.8
		Δ	1.8	-4.1	1.1	1.7	0.4	2.5	1.3	-0.4	0.6	0.3	0.1	2.4	-0.2	3.2	0.7	0.1
Flower102		Simple DA	98.1	92.2	98.2	98.6	98.0	98.8	98.8	98.8	98.4	97.5	98.7	98.0	98.9	98.7	98.7	98.7
		Default	98.9	97.4	99.1	99.1	98.9	99.2	99.1	99.1	99.2	98.0	99.2	99.1	99.3	99.1	99.3	99.1
		Δ	-0.8	-5.2	-0.9	-0.5	-0.9	-0.4	-0.3	-0.3	-0.8	-0.5	-0.5	-1.1	-0.4	-0.4	-0.6	-0.4
EuroSAT		Simple DA	87.3	91.0	88.4	92.0	92.0	91.7	91.9	92.5	92.1	92.5	92.3	93.1	92.7	92.8	93.3	92.8
		Default	90.0	88.1	90.3	94.9	95.0	94.1	93.8	94.5	95.5	95.3	95.0	95.3	95.8	94.9	94.9	95.5
		Δ	-2.7	2.9	-1.9	-2.9	-3.0	-2.4	-1.9	-2.0	-3.4	-2.8	-2.7	-2.2	-3.1	-2.1	-1.6	-2.7
Specialized	Resisc45	Simple DA	74.3	75.0	74.4	80.1	81.0	78.5	80.7	80.6	80.6	81.6	82.2	81.5	81.5	82.2	80.7	82.9
		Default	74.9	81.6	77.2	84.2	85.3	80.9	83.0	83.2	85.3	86.5	86.5	86.0	85.9	85.9	85.0	86.0
		Δ	-0.6	-6.6	-2.8	-4.1	-4.3	-2.4	-2.3	-2.6	-4.7	-4.9	-4.3	-4.5	-4.4	-3.7	-4.3	-3.1
Retinopathy		Simple DA	74.5	73.6	74.7	76.3	76.3	76.7	76.4	76.4	77.3	75.6	77.0	77.1	76.8	76.2	75.3	77.0
		Default	74.6	73.6	74.4	73.9	75.5	75.2	75.2	74.8	76.2	75.2	76.3	75.4	75.9	75.7	75.6	75.7
		Δ	-0.1	0.0	0.3	2.4	0.8	1.5	1.2	1.6	1.1	0.4	0.7	1.7	0.9	0.5	-0.3	1.3
dSpr-Ori		Simple DA	22.6	29.9	24.3	29.6	28.7	29.0	29.0	27.9	28.9	22.9	30.9	31.4	30.5	32.5	32.5	28.4
		Default	29.4	46.6	43.1	56.4	53.9	52.8	52.1	52.1	56.6	54.3	53.0	52.1	55.3	47.2	53.1	53.1
		Δ	-6.8	-16.7	-18.8	-26.8	-25.2	-23.8	-23.1	-24.2	-27.7	-31.4	-22.1	-20.7	-24.8	-16.9	-20.6	-24.7
Structured	KITTI	Simple DA	49.8	48.7	49.4	52.7	52.6	54.7	52.7	50.6	53.9	53.6	53.2	52.2	51.3	52.9	52.0	53.3
		Default	64.6	77.9	66.5	77.9	79.2	81.0	78.1	81.4	80.2	79.6	80.0	80.2	78.1	79.9	79.3	78.9
		Δ	-14.8	-29.2	-17.1	-25.2	-26.6	-26.3	-25.4	-30.8	-26.3	-26.0	-26.8	-28.0	-26.8	-27.0	-27.3	-25.6
sNORB-Azim		Simple DA	15.0	24.2	12.8	21.2	21.9	23.4	18.8	24.5	25.1	26.4	24.8	26.9	25.7	24.5	23.4	18.2
		Default	17.3	31.0	15.2	33.2	30.1	30.7	24.3	31.9	33.8	34.2	33.0	35.7	38.6	33.4	32.8	27.8
		Δ	-2.3	-6.8	-2.4	-12.0	-8.2	-7.3	-5.5	-7.4	-8.7	-7.8	-8.2	-8.8	-12.9	-8.9	-9.4	-9.6

Table 4: We apply simple data augmentations (DA) (RandomResizedCrop and RandomHorizontalFlip) on three datasets in each group. Data augmentation generally does not benefit VTAB-1K and thus, most recent PETL papers (Jie & Deng, 2023; 2022; Luo et al., 2023; Jia et al., 2022; Zhang et al., 2022b; Luo et al., 2023; Lian et al., 2022) skip it. To ensure that our results are directly comparable to existing papers, we don’t apply data augmentation.

Symbol (Abbreviation)	Definition
(H, W)	Resolution of input images
C	Number of channels (input images)
P	Resolution of patches
N	Number of patches (tokens)
N_h	Number of head in each Transformer layer
D	Embedding dimension
D_h	Embedding dimension for single-head attention
L_m	m -th Transformer layer
M	Number of Transformer layer
Z_{m-1}	Input of m -th Transformer layer
ViT	Vision Transformer
LN	Layer Normalization
MSA	Multi-head Self-Attention
MLP	Multi-Layer Perceptron
FC	Fully-connected layer

Table 5: Definitions of symbols and abbreviation used in Appendix B

B BACKGROUND

B.1 VISION TRANSFORMER

Overview of ViT. Inspired by the recent success of Transformer-based models Vaswani et al. (2017) in NLP Wolf et al. (2020), Vision Transformer (ViT) Dosovitskiy et al. (2020) has become widely used in computer vision. To handle 2D images, ViT divides an image $\mathbf{I} \in \mathbb{R}^{H \times W \times C}$ into N non-overlapping patches $\{\mathbf{I}^{(n)} \in \mathbb{R}^{P^2 \times C}\}_{n=1}^N$, where (H, W) is the resolution of the input image, C is the number of channels, $N = HW/P^2$ and (P, P) is the resolution of each patch. Each patch $\mathbf{I}^{(n)}$ is flattened and embedded into a D -dimensional vector $\mathbf{x}_0^{(n)}$ with a trainable linear projection. Incorporating the BERT design approach Kenton & Toutanova (2019), a “Class” token $\mathbf{x}_0^{(\text{Class})}$ is prepended to the sequence of embedded patches, whose output state at the last Transformer layer is utilized as the image representation. Finally, position embeddings $\mathbf{E}_{\text{pos}} \in \mathbb{R}^{D \times (1+N)}$ are added to preserve positional information and form the input $\mathbf{Z}_0 \in \mathbb{R}^{D \times (1+N)}$ to the ViT, which can be formulated by:

$$\mathbf{Z}_0 = [\mathbf{x}_0^{(\text{Class})}, \mathbf{x}_0^{(1)}, \mathbf{x}_0^{(2)}, \dots, \mathbf{x}_0^{(N)}] + \mathbf{E}_{\text{pos}} \quad (5)$$

As shown in the left part of Figure 9, a ViT typically consists of M layers, denoted by $\{L_m\}_{m=1}^M$. The input \mathbf{Z}_0 mentioned above is fed into the first layer L_1 , producing the output $\mathbf{Z}_1 = L_1(\mathbf{Z}_0) = [\mathbf{x}_1^{(\text{Class})}, \mathbf{x}_1^{(1)}, \dots, \mathbf{x}_1^{(N)}] \in \mathbb{R}^{D \times (1+N)}$, which maintains the same size as \mathbf{Z}_0 . Namely, \mathbf{Z}_1 comprises $1 + N$ feature tokens, and each corresponds to the same column in \mathbf{Z}_0 . Similarly, for $m = 2, \dots, M$, each layer L_m takes the output of the previous layer as input and generates the output, $\mathbf{Z}_m = L_m(\mathbf{Z}_{m-1})$. Finally, the ‘‘Class’’ vector $\mathbf{x}_M^{(\text{Class})}$ in \mathbf{Z}_M serves as the image feature for prediction. When dealing with classification tasks, the predicted label $\hat{y} = \text{Head}(\mathbf{x}_M^{(\text{Class})})$ is generated through a linear head (*i.e.*, a fully-connected layer).

Details of each Transformer layer. As shown in the right part of Figure 9, each Transformer layer consists of a Multi-head Self-Attention (MSA) block, a Multi-Layer Perceptron (MLP) block, and two Layer Normalization (LN) layers Ba et al. (2016). Formally, a Transformer layer L_m can be defined as

$$\begin{aligned}\mathbf{Z}'_m &= \text{MSA}(\text{LN}(\mathbf{Z}_{m-1})) + \mathbf{Z}_{m-1} \\ \mathbf{Z}_m &= \text{MLP}(\text{LN}(\mathbf{Z}'_m)) + \mathbf{Z}'_m\end{aligned}\tag{6}$$

where $\mathbf{Z}_{m-1} = [\mathbf{x}_{m-1}^{(\text{Class})}, \mathbf{x}_{m-1}^{(1)}, \dots, \mathbf{x}_{m-1}^{(N)}] \in \mathbb{R}^{D \times (1+N)}$ is the output of the preceding $(m - 1)$ -th Transformer layer. The MLP is applied to each column vector of \mathbf{Z}'_m independently.

In order to encapsulate multiple complex relationships amongst different elements in the sequence, the MSA block comprises N_h single-head self-attention blocks. For the i^{th} single-head self-attention block, an generic input \mathbf{Z} is first projected into three matrices, namely Query $\mathbf{Q}^{(i)}$, Key $\mathbf{K}^{(i)}$, and Value $\mathbf{V}^{(i)}$

$$\mathbf{Q}^{(i)} = \mathbf{W}_Q^{(i)} \mathbf{Z}, \quad \mathbf{K}^{(i)} = \mathbf{W}_K^{(i)} \mathbf{Z}, \quad \mathbf{V}^{(i)} = \mathbf{W}_V^{(i)} \mathbf{Z},\tag{7}$$

where $\mathbf{W}_{Q/K/V}^{(i)} \in \mathbb{R}^{D_h \times D}$ ³ where D_h is the embedding dimension for a single head self-attention block and typically set to D/N_h . The i^{th} self-attention head in MSA is formulated as

$$\text{Attn}^{(i)}(\mathbf{Z}) = \mathbf{V}^{(i)} \times \text{Softmax}\left(\frac{\mathbf{K}^{(i)\top} \mathbf{Q}^{(i)}}{\sqrt{D_h}}\right) \in \mathbb{R}^{D_h \times (1+N)}\tag{8}$$

The outputs of all heads are concatenated and linearly projected by a fully connected layer (FC_{attn}) with weight $\mathbf{W}_O \in \mathbb{R}^{D \times (D_h \cdot N_h)}$ as the output of the MSA block.

$$\text{MSA}(\mathbf{Z}) = \mathbf{W}_O [\text{Attn}^0(\mathbf{Z}), \dots, \text{Attn}^{N_h}(\mathbf{Z})]\tag{9}$$

The MLP block can be defined as

$$\text{MLP}(\mathbf{Z}) = \text{GELU}(\mathbf{Z}\mathbf{W}_1 + b_1)\mathbf{W}_2 + b_2\tag{10}$$

where $\mathbf{W}_1 \in \mathbb{R}^{D \times 4D^4}$, $\mathbf{W}_2 \in \mathbb{R}^{4D \times D}$, $b_1 \in \mathbb{R}^{4D}$, $b_2 \in \mathbb{R}^D$ are weights and biases for two FC layers (FC_1 and FC_2) respectively.

Since PETL methods often entail incorporating additional components to modify the intermediate features within or between Transformer layers, we adopt the notation $\{h_1, \dots, h_{10}\}$ to denote the intermediate features in the unravelled view of a Transformer layer (as depicted in Figure 9) to facilitate a clearer illustration of the PETL methods discussed in the subsequent section.

³For brevity, we ignore the layer index m for the projection matrices $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V$, but each layer has its own projection matrices.

B.2 EVALUATED METHODS

In this section, we dive into the details of 12 State-Of-The-Art PETL approaches, categorized into three groups: Prompt-based, Adapter-based and Selective Parameter Tuning. We will describe the distinctions and tradeoffs between them. A consolidated overview of these approaches is summarized in Table 6.

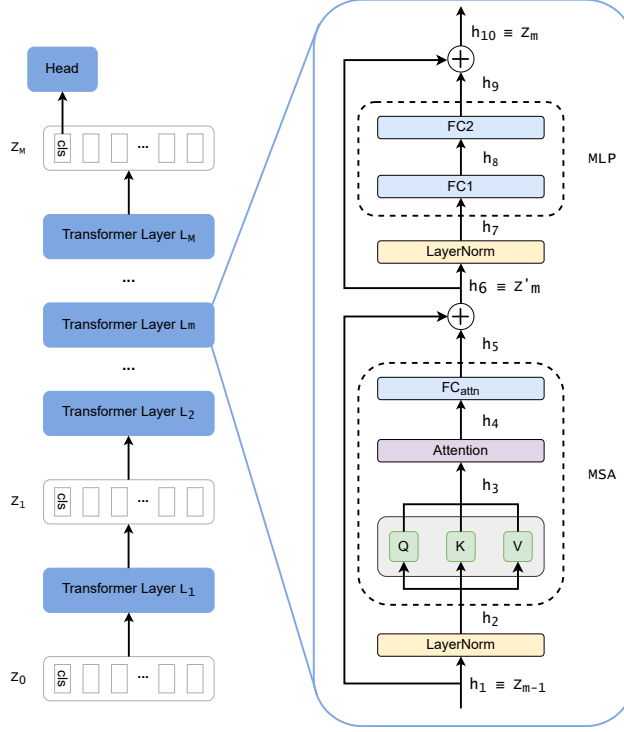


Figure 9: An overview of a Transformer block in ViT. We adopt the notation $\{h_1, \dots, h_{10}\}$ to denote the intermediate features within a Transformer block to facilitate a clearer illustration of the PETL methods discussed in subsection B.2.

B.2.1 PROMPT-BASED METHODS

Prompt-based learning emerged in NLP as an effective approach to adapt pre-trained models for downstream tasks Liu et al. (2023); Lialin et al. (2023). The core concept involves augmenting the model input with task-specific hints (prompts), which aid the pre-trained model in addressing novel tasks with its existing knowledge. Hard prompts are human-interpretable natural language hints, encompassing task instructions, in-context examples, or supporting information. Alternatively, soft prompts are continuous vector hints that are incorporated into the input embeddings of the input layers or hidden states of other layers. Soft prompts are updated during the fine-tuning process using gradient-based methods, guided by the downstream task-specific loss functions, while the pre-trained model itself remains fixed. The splendid success of prompts in NLP has sparked a growing interest in adopting it in computer vision Yu et al. (2023b); Tu et al. (2023) and multi-modal domains Gu et al. (2023).

In this paper, we investigate a prominent and strong prompt-based method called **Visual Prompt Tuning (VPT)** Jia et al. (2022), which represents one of the early endeavours in introducing prompts to computer vision. Specifically, VPT-Shallow adds l prompts $P_0 \in \mathbb{R}^{l \times D}$ to the input of the first Transformer layer Z_0 and the output \tilde{P}_0 of P_0 serves as the input for the next layer as depicted in Equation 11. VPT-Shallow can be perceived as the addition of learnable pixels to the original images. On the other hand, VPT-Deep inserts l prompts $\{P_m \in \mathbb{R}^{l \times D}\}_{m=0}^M$ to the input of every Transformer layer Z_m but their outputs are discarded at the end of the layer as illustrated in Equation 12.

⁴4 is the MLP ratio in ViT-B

Table 6: PETL Methods Summary: Prompt-based and adapter-based methods incorporate additional parameters to modify features while keeping the original backbone intact. However, these added parameters introduce additional inference overhead. In contrast, selective tuning methods modify the backbone by updating selective parameters, thereby incurring no additional inference overhead.

Method	What	Tunable Parameters	Hyper Parameters	Modified Type	Inference Efficient
VPT-Deep	$h_1 = [h_1, P]$	$P \in \mathbb{R}^{l \times D}$	l : Number of prompts	Feature	✗
AdaptFormer	$h_9 = h_9 + \text{Adapter}(h_7)$	$W_{\text{down/up}} \in \mathbb{R}^{r \times D/D \times r}$ in Adapter	s : Scale factor in Adapter r : Bottleneck dimension	Feature	✗
Pfeif. Adapter	$h_9 = \text{Adapter}(h_9)$	$W_{\text{down/up}} \in \mathbb{R}^{r \times D/D \times r}$ in Adapter	s : Scale factor in Adapter r : Bottleneck dimension	Feature	✗
Houl. Adapter	$h_5 = \text{Adapter}_1(h_5)$ $h_9 = \text{Adapter}_2(h_9)$	$W_{\text{down/up}}^1 \in \mathbb{R}^{r \times D/D \times r}$ in Adapter ₁ $W_{\text{down/up}}^2 \in \mathbb{R}^{r \times D/D \times r}$ in Adapter ₂	s : Scale factor in Adapter r : Bottleneck dimension	Feature	✗
Convpass	$h_5 = \text{Convpass}_1(h_2) + h_5$ $h_9 = \text{Convpass}_2(h_7) + h_9$	$W_{\text{conv2d}}^1 \in \mathbb{R}^{r \times r \times k \times k}$ $W_{\text{down/up}}^1 \in \mathbb{R}^{r \times D/D \times r}$ in Convpass ₁ $W_{\text{conv2d}}^2 \in \mathbb{R}^{r \times r \times k \times k}$ $W_{\text{down/up}}^2 \in \mathbb{R}^{r \times D/D \times r}$ in Convpass ₂	s : Scale factor in Convpass r : Bottleneck dimension k : Kernel size of conv2d	Feature	✗
RepAdpater	$h_2 = \text{RepAdapter}_1(h_2)$ $h_7 = \text{RepAdapter}_2(h_7)$	$W_{\text{conv1d}}^1 \in \mathbb{R}^{r \times D}$ $b^1 \in \mathbb{R}^r$ in RepAdapter ₁ $W_{\text{conv1d}}^2 \in \mathbb{R}^{D \times G}$ $b^2 \in \mathbb{R}^G$ in RepAdapter ₂	s : Scale factor in RepAdapter r : Bottleneck dimension G : Number of groups	Feature	✗
LayerNorm	$h_2 = \text{LayerNorm}_1(h_1)$ $h_7 = \text{LayerNorm}_2(h_6)$	$W^{1(2)}, b^{1(2)} \in \mathbb{R}^D$ in LayerNorm ₁₍₂₎	N/A	Backbone	✓
BitFit	Fine-tune all bias terms in the network	$b^{1(2)} \in \mathbb{R}^D$ in LayerNorm ₁₍₂₎ $b^{Q/K/V} \in \mathbb{R}^D$ in Q/K/V $b^{FC_{\text{attn}}} \in \mathbb{R}^D$ in FC _{attn} $b^1 \in \mathbb{R}^{4D}$, in FC ₁ , $b^2 \in \mathbb{R}^D$ in FC ₂	N/A	Backbone	✓
DiffFit	• LayerNorm + BitFit • $h_5 = \gamma_1 \cdot h_5$ • $h_9 = \gamma_2 \cdot h_9$	• All tunable parameters in LayerNorm & BitFit • $\gamma_1, \gamma_2 \in \mathbb{R}^D$	N/A	Backbone	✓
SSF	$h_2 = \text{SSF}_2(h_2)$, $h_3 = \text{SSF}_3(h_3)$ $h_5 = \text{SSF}_5(h_5)$, $h_7 = \text{SSF}_7(h_7)$ $h_8 = \text{SSF}_7(h_8)$, $h_9 = \text{SSF}_9(h_9)$	$W^{2,5,7,9} \in \mathbb{R}^D$, $b^{2,5,7,9} \in \mathbb{R}^D$ $W^3 \in \mathbb{R}^{3D}$, $b^3 \in \mathbb{R}^{3D}$ $W^8 \in \mathbb{R}^{4D}$, $b^8 \in \mathbb{R}^{4D}$	N/A	Backbone	✓
LoRA	$h_3 = \text{LoRA}(h_2) + h_3$	$W_{\text{down/up}}^{Q/K/V} \in \mathbb{R}^{r \times D/D \times r}$ in LoRA	r : Bottleneck dimension	Backbone	✓
FacT _{TT(TK)}	$h_3 = \text{FacT}_{\text{TT(TK)}}(h_2) + h_3$ $h_5 = \text{FacT}_{\text{TT(TK)}}(h_4) + h_5$ $h_8 = \text{FacT}_{\text{TT(TK)}}(h_7) + h_8$ $h_9 = \text{FacT}_{\text{TT(TK)}}(h_8) + h_9$	$U \in \mathbb{R}^{D \times r}$, $V \in \mathbb{R}^{D \times r}$, $\Sigma \in \mathbb{R}^{12L \times r \times r}$ in FacT _{TT} $U \in \mathbb{R}^{D \times r}$, $V \in \mathbb{R}^{D \times r}$, $A \in \mathbb{R}^{12L \times r}$, $B \in \mathbb{R}^{r \times r \times r}$ in FacT _{TK}	s : Scale factor in FacT _{TT(TK)} r : Bottleneck dimension	Backbone	✓

$$\begin{aligned} [\tilde{P}_1, Z_1] &= L_m([P_0, Z_0]) \\ [\tilde{P}_m, Z_m] &= L_m([\tilde{P}_{m-1}, Z_{m-1}]) \quad m = 2, 3, \dots, M \end{aligned} \quad (11)$$

$$[_, Z_m] = L_m([P_{m-1}, Z_{m-1}]) \quad m = 1, 2, 3, \dots, M \quad (12)$$

Throughout the adaptation process, the pre-trained model is frozen and no additional weights are introduced to the model, thereby preserving the model’s original behaviour. During the forward pass, the output Z_m of layer m is changed because of the interaction between Z_{m-1} and P_{m-1} (or \tilde{P}_{m-1}) in the MSA block. Thus, the output feature is adapted to the downstream tasks by iteratively tuning the prompts through gradient descent.

B.2.2 ADAPTER-BASED METHODS

Adapter-based methods typically introduce additional trainable parameters into a frozen pre-trained model to facilitate learning of downstream tasks Lialin et al. (2023). Initially developed for multi-domain adaptation Rebuffi et al. (2017; 2018) and continual learning Rosenfeld & Tsotsos (2018); Mai et al. (2022), the idea of Adapters is subsequently embraced by Houlsby *et al.* Houlsby et al.

(2019) in the NLP domain to adapt Transformer-based networks for downstream tasks, and it also has garnered increasing interest in the computer vision field Yu et al. (2023b). In this comparative analysis, we concentrate on five popular Adapter-based methods, encompassing the original Adapter, along with variants focusing on adjusting the positions of Adapters Chen et al. (2022b); Pfeiffer et al. (2021), introducing visual inductive biases Jie & Deng (2022), as well as employing re-parameterization to reduce the number of trainable parameters and inference latency Luo et al. (2023).

Houl. Adapter Houlsby et al. (2019) inserts two lightweight bottleneck-structured modules into each Transformer layer: one after the MSA block and the other after the MLP block. As depicted in Figure 10a, the Adapter is composed of a down-projection layer with $\mathbf{W}_{\text{down}} \in \mathbb{R}^{r \times D}$, a nonlinear activation function σ , an up-projection layer with $\mathbf{W}_{\text{up}} \in \mathbb{R}^{D \times r}$, a scaling factor s and a skip-connection. To limit the number of trainable parameters, the bottleneck dimension is much smaller than the feature dimension $r \ll D$. Formally, Houl. Adapter can be defined as:

$$h_5 = \text{Adapter}_1(h_5) \quad h_9 = \text{Adapter}_2(h_9) \quad (13)$$

$$\text{Adapter}(h) = s \cdot \mathbf{W}_{\text{up}} \sigma(\mathbf{W}_{\text{down}} h) + h \quad (14)$$

Pfeif. Adapter Pfeiffer et al. (2021) is a more efficient variant that introduces the Adapter solely after the MLP block, a strategy that has demonstrated effectiveness in recent studies Hu et al. (2021). Pfeif. Adapter can be defined formally as $h_9 = \text{Adapter}(h_9)$ where Adapter follows Equation 14.

AdaptFormer Chen et al. (2022b) proposed to insert the Adapter in parallel with the MLP block, which differs from the sequential design of Houl. and Pfeif. Adapter. The rationale behind this parallel design lies in the belief that the domain-specific features generated by the Adapter can complement the domain-agnostic features derived from the original MLP block, leading to an improved feature ensemble Szegedy et al. (2015). Formally, AdaptFormer can be defined as $\mathbf{h}_9 = \mathbf{h}_9 + \text{Adapter}(\mathbf{h}_7)$ where Adapter follows Equation 14.

ConvPass (Convolutional By-Passes) Jie & Deng (2022) addresses the concern that many existing Adapters lack visual inductive bias, potentially limiting their performance for downstream vision tasks with limited data. To this end, the authors introduce a convolutional bottleneck module, running in parallel with the MSA or(and) MLP block. This module encompasses a 1×1 convolution reducing the channel with $\mathbf{W}_{\text{down}} \in \mathbb{R}^{r \times D}$, a 3×3 convolution with the same input and output channel, a 1×1 convolution expanding the channel $\mathbf{W}_{\text{up}} \in \mathbb{R}^{D \times r}$, two nonlinear functions σ and a scaling factor s , as shown in Figure 10b. The authors argue that Convpass is more efficient at capturing visual information in low-data scenarios due to its hard-coded locality of convolutional layers. The formal definition of Convpass is shown in Equation 15.

$$\begin{aligned} h_5 &= \text{Convpass}_1(h_2) + h_5 \quad h_9 = \text{Convpass}_2(h_7) + h_9 \\ \text{Convpass}(h) &= s \cdot \mathbf{W}_{\text{up}} \sigma(\text{Conv2d}(\sigma(\mathbf{W}_{\text{down}} h))) \end{aligned} \quad (15)$$

RepAdapter Luo et al. (2023) found that the removal of the nonlinear function in the Adapter does not result in performance degradation for vision tasks. In light of this finding, the authors propose a linear Adapter with group-wise transformation Luo et al. (2022) and sequentially added two of these linear Adapters to both MSA and MLP blocks. Owing to the sequential placement of the RepAdapter and its inherent linearity, the additional parameters can be re-parameterized to the original MSA or MLP block after training, thereby incurring zero additional costs during inference. RepAdapter is illustrated in Figure 10c and formally defined in Equation 16.

$$\begin{aligned} h_5 &= \text{RepAdapter}_1(h_2) \quad h_7 = \text{RepAdapter}_2(h_7) \\ \text{RepAdapter}(h) &= s \cdot \phi_{\text{up}}(\phi_{\text{down}}(h)) + h \\ \tilde{h} &= \phi_{\text{down}}(h) = \mathbf{W}_{\text{down}} h \\ \phi_{\text{up}}(\tilde{h}) &= [\mathbf{W}_{g1} \tilde{h}_{g1}, \dots, \mathbf{W}_{gG} \tilde{h}_{gG}] \end{aligned} \quad (16)$$

where $\mathbf{W}_{\text{down}} \in \mathbb{R}^{r \times D}$, $\tilde{h}_{g(1, \dots, G)} \in \mathbb{R}^{\frac{r}{G} \times (N+1)}$ is the features splitted from $\tilde{h} \in \mathbb{R}^{r \times (N+1)}$ and G is the number of groups in group-wise transformation Luo et al. (2022). $\mathbf{W}_{g(1, \dots, G)} \in \mathbb{R}^{\frac{D}{G} \times \frac{r}{G}}$ is the projection weight matrix.

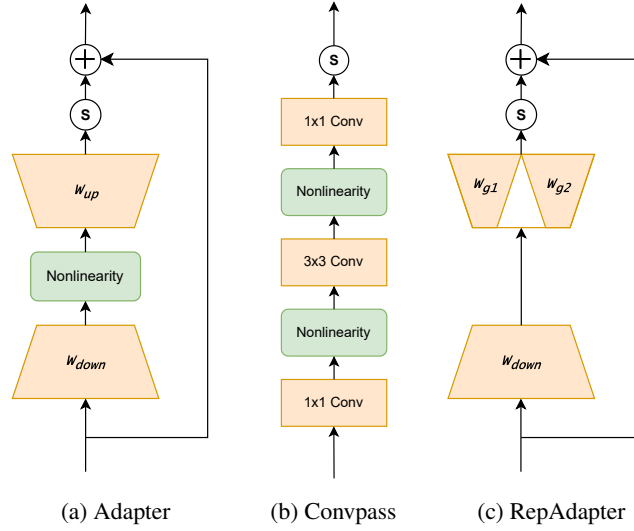


Figure 10: Comparison of three Adapter structures.

B.2.3 SELECTIVE PARAMETER TUNING METHODS

The methods falling within this category aim to selectively update the parameters of a pre-trained model for downstream tasks. Within transfer learning, two prominent strategies, namely *full fine-tuning* and *linear probing* Kornblith et al. (2019); Zhuang et al. (2020), represent the two extremes of this category. *Full fine-tuning* updates all the model parameters end-to-end based on the new dataset while *linear probing* treats the pre-trained model as a feature extractor and only updates the prediction heads while keeping the backbone frozen. Although *full fine-tuning* generally exhibits superior performance compared to *linear probing* Zhai et al. (2019), it possesses certain limitations that may hinder its practicality in real-world production settings. Firstly, it requires running gradient descent for all parameters and necessitates storing a separate fine-tuned model for each task, incurring significant computational, memory, and storage overhead. These challenges become more salient with Transformer-based models whose parameters grow exponentially. Secondly, *full fine-tuning* may distort pre-trained features and underperform *linear probing* in out-of-distribution (OOD) scenarios Kumar et al. (2021).

To cope with the above issues, a cohort of PETL methods has emerged under this category. In addition to the two common approaches mentioned above, our investigation encompasses seven methods that can be further categorized into two groups: direct selective tuning Zaken et al. (2022); Basu et al. (2023); Xie et al. (2023), which involves the direct modification of selective weights, and efficient selective tuning Hu et al. (2021); Jie & Deng (2023); Lian et al. (2022), which approximates the weight updates with low-rank factors.

Notably, an extra advantage of methods in this category is that they introduce **no additional inference latency**, making them particularly favourable when inference efficiency is a priority. Methods within the direct selective tuning group abstain from introducing any new modules, thus inherently avoiding extra inference latency. Meanwhile, for methods in the efficient selective tuning group, the added modules can often be seamlessly integrated into weights of the pre-trained models through the re-parameterization techniques Jacob et al. (2018); Ding et al. (2021), thereby ensuring the absence of increased inference latency as well.

Direct Selective Tuning

BitFit Zaken et al. (2022) is a simple yet effective method that only tunes the bias parts of the pre-trained model. For each Transformer layer in ViT, BitFit updates the bias terms in the QKV projections and the FC layer in the MSA block, two FC layers in the MLP block and two LN blocks. It also updates the bias in the projection for patch embedding. The original authors underscore

BitFit’s capability to achieve performance comparable to full fine-tuning or even surpass it under low and medium-data scenarios in BERT models Kenton & Toutanova (2019).

LayerNorm Basu et al. (2023) represents another simple but strong baseline that solely tunes the two LN blocks in each Transformer layer - one before the MSA block and another before the MLP block. Given that each LN block contains merely two trainable parameters $\{\mathbf{W}_{LN}, \mathbf{b}_{LN}\} \in \mathbb{R}^D$, LN-tune stands out as an exceedingly light-weight approach compared to other PETL methods. For instance, ViT-B/16 ($\sim 86\text{M}$ parameters) has only $\sim 38\text{k}$ LN parameters, accounting for $\sim 0.04\%$ of the total parameters.

DiffFit Xie et al. (2023) is a recently proposed PETL strategy designed for adapting large pre-trained diffusion models to the new domains. DiffFit exclusively fine-tunes the bias terms and the LN blocks within the network. Furthermore, it inserts learnable scale factors γ to shift the features after the MSA and the MLP blocks, as shown in Equation 17. Consequently, DiffFit can be regarded as a combination of the BitFit and Ln-Tune, incorporating additional feature shift factors.

$$\begin{aligned} h_5 &= \gamma_1 \cdot h_5 \\ h_9 &= \gamma_2 \cdot h_9 \end{aligned} \quad (17)$$

Efficient Selective Tuning

LoRA (Low-Rank Adaptation) Hu et al. (2021) drew inspiration from recent investigations demonstrating that the learned over-parametrized models in fact reside on a low intrinsic dimension Li et al. (2018); Aghajanyan et al. (2021). Building upon this insight, the authors hypothesize that the change in weights during model adaptation also exhibits a low intrinsic rank and injects trainable low-rank decomposition matrices to approximate the weight updates. The LoRA update methodology is strategically applied to the Query/Value projection weights $\mathbf{W}_{Q/V} \in \mathbb{R}^{D \times D}$ within the MSA block. Concretely, the weight updates are approximated as $\mathbf{W}_{Q/V} + \Delta\mathbf{W}_{Q/V} = \mathbf{W}_{Q/V} + \mathbf{W}_{\text{down}}^{Q/V} \mathbf{W}_{\text{up}}^{Q/V}$ where $\mathbf{W}_{\text{down/up}}^{Q/V} \in \mathbb{R}^{D \times r/r \times D}$ and rank $r \ll D$. The authors use a random Gaussian initialization for $\mathbf{W}_{\text{up}}^{Q/V}$ and zero for $\mathbf{W}_{\text{down}}^{Q/V}$ so that $\Delta\mathbf{W}_{Q/V} = \mathbf{W}_{\text{down}}^{Q/V} \mathbf{W}_{\text{up}}^{Q/V}$ is zero at the beginning of training. The formal definition of LoRA is articulated in Equation 18, utilizing the notations delineated in Figure 9.

$$\begin{aligned} h_3 &= \text{LoRA}(h_2) + h_3 \\ h_3 &= [\mathbf{Q}, \mathbf{K}, \mathbf{V}] \\ \text{LoRA}(h_2) &= [\mathbf{W}_{\text{down}}^Q \mathbf{W}_{\text{up}}^Q h_2, 0, \mathbf{W}_{\text{down}}^V \mathbf{W}_{\text{up}}^V h_2] \end{aligned} \quad (18)$$

FacT (Factor Tuning) Jie & Deng (2023) is inspired by the recent advances in Transformer compression Wang et al. (2022); Zhang et al. (2022a) and exploited the low-rank update paradigm (e.g., LoRA) to the extreme. While LoRA posits that the update for an individual weight matrix manifests a low-rank characteristic during fine-tuning, FacT advances the proposition that the weight updates spanning different matrices can also be effectively approximated using low-rank decomposition matrices. Specifically, FacT encapsulates the four weight matrices $\mathbf{W}_{Q/K/V/O} \in \mathbb{R}^{D \times D}$ in the MSA block and the two weight matrices $\mathbf{W}_1 \in \mathbb{R}^{D \times 4D}$, $\mathbf{W}_2 \in \mathbb{R}^{4D \times D}$ in the MLP block into a single $\mathbf{W}_{\text{FacT}} \in \mathbb{R}^{12M \times D \times D}$ tensor where M is the number of Transformer layer. The update of \mathbf{W}_{FacT} , $\Delta\mathbf{W}_{\text{FacT}}$, can be decomposed into several factors to promote parameter efficiency. To this end, the authors leverage the well-established Tensor-Train (TT) Oseledets (2011) and the Tucker (TK) De Lathauwer et al. (2000) format to decompose $\Delta\mathbf{W}_{\text{FacT}}$. FacT_{TT} and FacT_{TK} are used to denote different decomposition formats for FacT and their formal definitions can be found in Equation 19.

$$\text{FacT}_{\text{TT}} : \Delta\mathbf{W}_{\text{FacT}} = s \cdot \Sigma \times_2 \mathbf{U}^\top \times_3 \mathbf{V}^\top \quad (19)$$

$$\text{FacT}_{\text{TK}} : \Delta\mathbf{W}_{\text{FacT}} = s \cdot \mathbf{A} \times_1 \mathbf{B}^\top \times_2 \mathbf{U}^\top \times_3 \mathbf{V}^\top \quad (20)$$




Figure 11: Performance gain for PETL methods by turning drop-path-rate on.

where $\mathbf{U} \in \mathbb{R}^{D \times r}$, $\mathbf{V} \in \mathbb{R}^{D \times r}$, $\mathbf{\Sigma} \in \mathbb{R}^{12L \times r \times r}$, $\mathbf{B} \in \mathbb{R}^{12L \times r}$, $\mathbf{A} \in \mathbb{R}^{r \times r \times r}$ and the \times_j denotes mode- j product and s is the scaling factor.

Since $\Delta \mathbf{W}_{FacT}$ contains the updates for $\mathbf{W}_{Q/K/V/O}$, $\mathbf{W}_{1/2}$, the modified forward pass inherently influences h_3, h_5, h_8, h_9 . Let’s consider h_5 for elucidation. Once the weight update $\Delta \mathbf{W}_{FacT}$ is calculated with $\text{FacT}_{TT(TK)}$ in Equation 19, the corresponding update for \mathbf{W}_O , $\Delta \mathbf{W}_O$, is extracted from $\Delta \mathbf{W}_{FacT}$. Similar to the modified forward pass of LoRA, $h_5 = h_4 \Delta \mathbf{W}_O + h_5$.

SSF (Scale & Shift deep Features) Lian et al. (2022) employs linear transformations to adapt the intermediate features extracted by a pre-trained model. Motivated by the feature modulation methods Huang & Belongie (2017); Perez et al. (2018), SSF is designed to accommodate the distribution difference between the upstream and downstream datasets. Specifically, SSF modulates the features residing at $h_2, h_3, h_5, h_7, h_8, h_9$ by incorporating scale and shift factors. To demonstrate the mechanism of SSF, let’s consider $h_5 \in \mathbb{R}^{(N+1) \times D}$ as an illustrative example and other features can similarly undergo the same transformative process. Formally, the modulated h_5 is formulated as follows.

$$h_5 = \text{SSF}_5(h_5) = \mathbf{w}^5 \odot h_5 + \mathbf{b}^5 \quad (21)$$

where $\mathbf{w}^5 \in \mathbb{R}^D$, $\mathbf{b}^5 \in \mathbb{R}^D$ are the scale and shift factors affiliated with the SSF module attributed to h_5 , and \odot is the dot product. It is noteworthy that each modulated feature has its own SSF module with corresponding scale and shift factors. The modification details for other features are summarized in Table 6.

C MORE DETAILED RESULTS

Drop-path-rate. Learning with low-shot data is prone to over-fitting. We find that if the drop path rate — which stochastically drops a transformer block per sample Huang et al. (2016) — is set not as default (*i.e.*, nonzero), all the methods can benefit from such a regularization. Figure 11 shows the performance gain by tuning the drop-path-rate on compared with the default 0.

More results on prediction similarity analysis. Figure 13 shows the prediction analysis discussed in section 4 for all the datasets in VTAB-1K. It is expected that their predictions are similar for datasets with very high accuracy, such as Flowers102 (avg 99.1%) and Caltech101 (avg 91.4%). Beyond them, we find that most PETL methods show diverse predictions in other datasets in VTAB-1K.

Prediction similarity within the same PETL group. To verify if methods within the same PETL group share more prediction similarity, we plotted the prediction overlap for adapter-based methods, selective-tuning methods, and methods from different groups. As shown in Figure 12, methods within the same group share slightly more prediction similarity than those from different groups, but they still exhibit distinct predictions. Figure 3 in the main paper also supports this observation. Methods are grouped based on the categories defined in subsection 2.2. If methods within the same group had very high similarities, we would see bright squares, which are only slightly evident around BitFit, DiffFit, LayerNorm, and SSF.

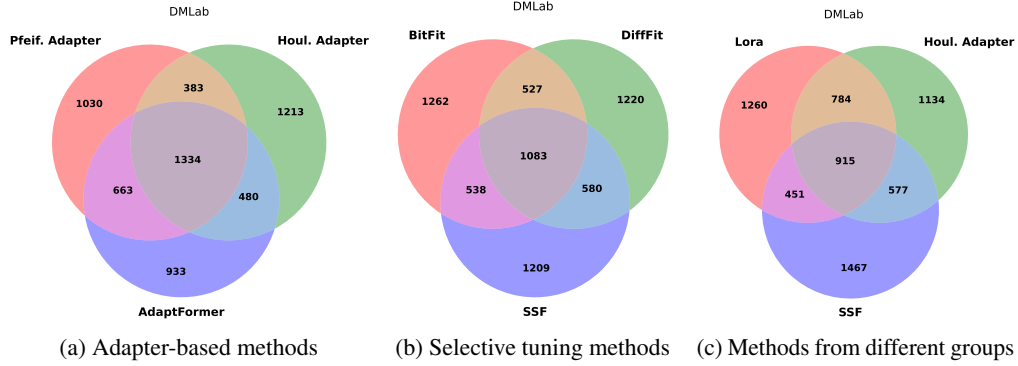


Figure 12: Prediction overlap for the 5K most confident samples. Although methods from the same group share slightly more prediction overlap than methods from other groups, they still have quite different predictions

WiSE PETL results for all distribution shift datasets. We provide detailed WiSE PETL performance for each distribution shift dataset in Figure 14. WiSE improves both the robustness and the in-distribution performance of PETL methods. Interestingly, even though full fine-tuning is generally less robust than PETL methods, applying WiSE allows it to achieve better performance in both target distribution and distribution shift data.

Figure 1a details. This figure illustrates the relative performance compared to linear probing (\times) on VTAB-1K. The range between the highest and lowest accuracy across 14 PETL methods is represented by $\bullet\text{--}\bullet$, while \blacksquare denotes the performance of full fine-tuning.

Figure 1b details. The X-axis represents the accuracy on ImageNet-1K, while the Y-axis shows the distribution shift accuracy (averaged across ImageNet-V2, ImageNet-S, ImageNet-R, and ImageNet-A). The cyan squares (\blacksquare) represent the zero-shot performance of the CLIP model, and stars (\star) denote the performance of fine-tuned models. Each curve corresponds to the WiSE+PETL method, with dots \bullet indicating different mixing coefficients α as described in section 7.

Figure 2 details. For each dataset in VTAB-1K, 15 methods (14 PETL methods plus linear probing) are ranked by accuracy. Within each dataset group (e.g., Natural), the element (i, j) in the ranking frequency matrix indicates how often method i ranks j^{th} . For instance, in the Natural group matrix, the entry $(1, 3)$ equals 2, meaning DiffFit ranked 3rd in two datasets within this group. The row sums correspond to the total number of datasets in each group (e.g., 7 datasets for the Natural group). Methods are sorted by their average rank (shown in brackets), and the parameters column indicates the number of trainable parameters in millions.

Figure 3 details. Each entry (i, j) in the prediction similarity matrix represents the percentage of test samples for which methods i and j made the same prediction. The diagonal entries are always 1, indicating perfect agreement with themselves. To compute (i, j) , predictions from models fine-tuned by methods i and j are compared, with (i, j) equaling the number of matching predictions divided by the total test samples.

Figure 4 details. The Venn diagrams are generated by fine-tuning a pre-trained model on CIFAR100 (VTAB-1K) using LoRA, SSF, and Adapter methods. For Figure 4(a), we selected the correct predictions from the top 5K most confident samples for each method and visualized the overlap among the three methods. For Figure 4(b), we did the same for the wrong predictions, selecting from the 5K least confident samples.

Figure 5 details. For each VTAB-1K dataset, the worst-performing PETL method serves as the baseline (\times). Each \bullet represents the relative performance of other PETL methods compared to this baseline. An ensemble prediction (\blacktriangle) is generated based on the average logits of all PETL methods for each test sample.

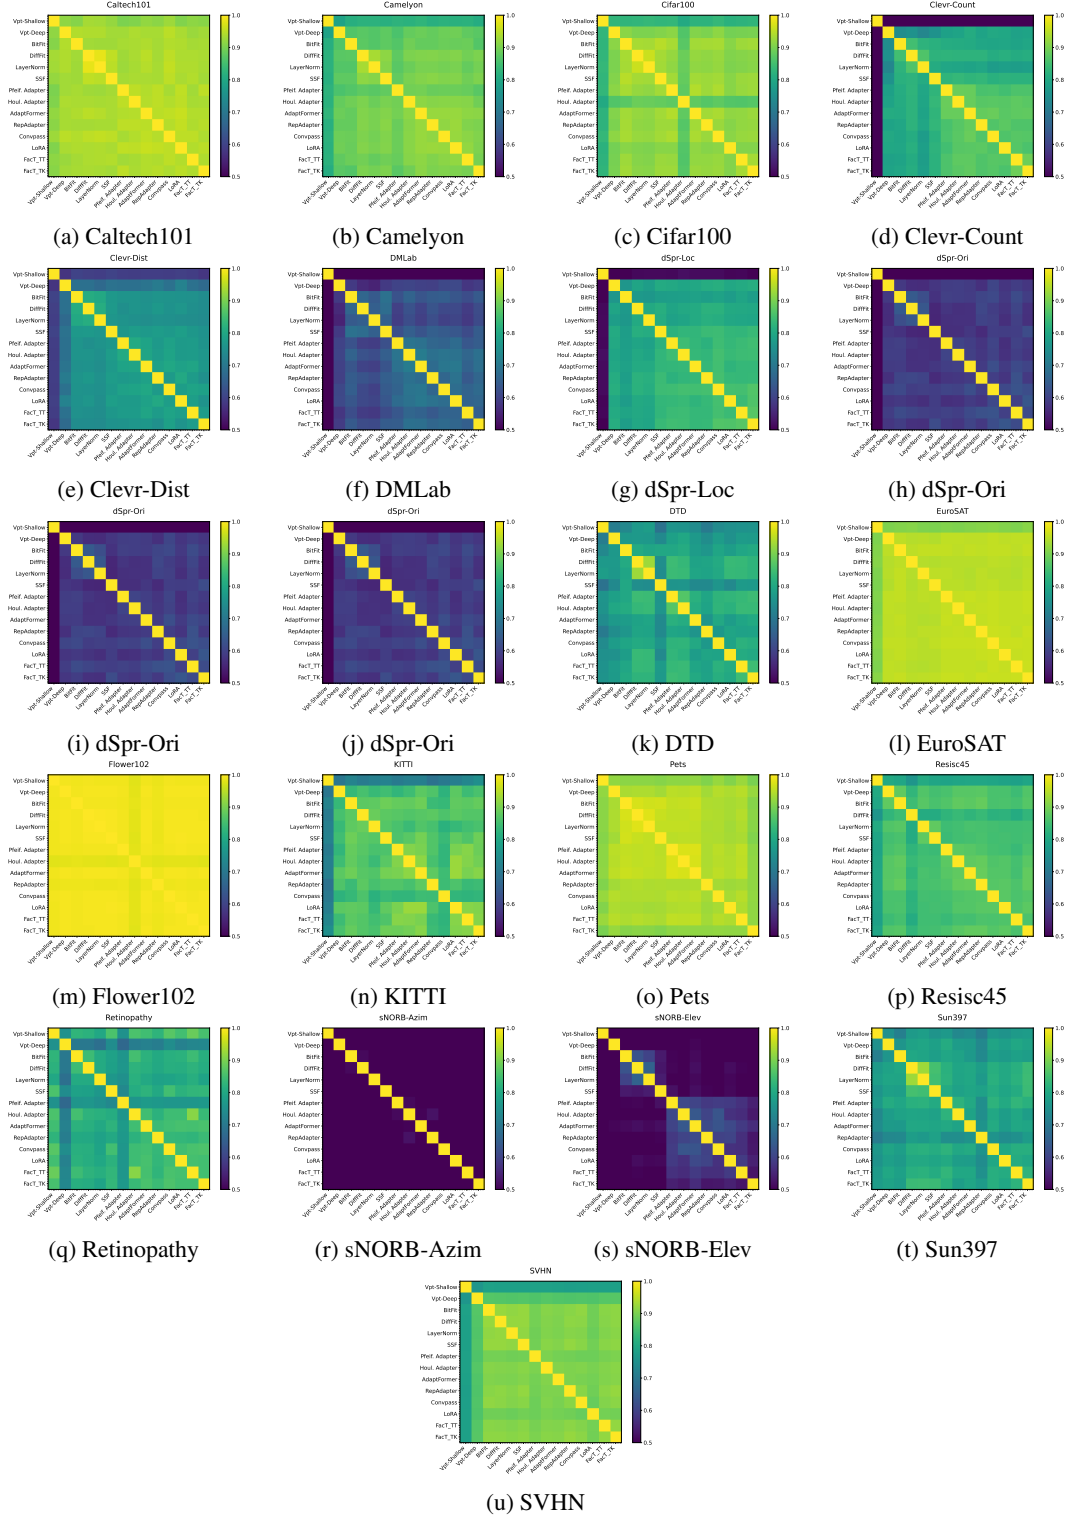


Figure 13: Prediction similarity analysis on other datasets.

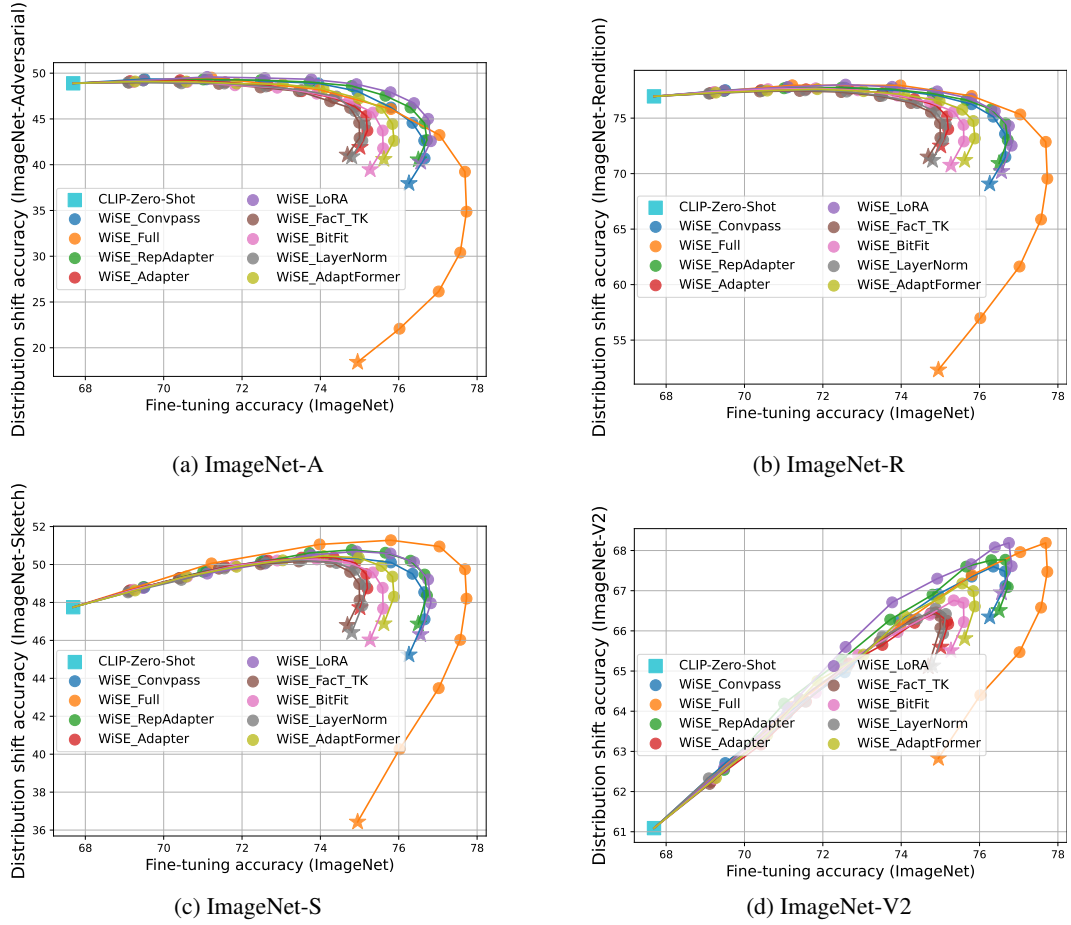


Figure 14: WISE PETL performance on all distribution shift datasets. Target distribution vs. distribution shifts

Figure 6 details. Different colors represent various PETL methods. Each • along a curve (corresponding to a single PETL method) indicates the accuracy at a specific tunable parameter size, allowing us to observe how the size of tunable parameters impacts accuracy.

Figure 7 details. Each sub-figure displays accuracy on the Y-axis, with columns representing linear probing (left), the best PETL methods (middle), and full fine-tuning (right). Sub-figures (a) and (b) correspond to VTAB-1K (low-shot), while (c) and (d) correspond to many-shot settings. Different colors represent distinct datasets.

D BROADER IMPACTS

Our study provides a unifying study of PETL in visual recognition. We expect it to serve as a valuable practical user guide to benefit society. Specifically, fine-tuning large models needs significant computation. A unifying study of PETL will ease end-users to apply more parameter-efficient and computation-efficient ways for fine-tuning. To our knowledge, our paper does not introduce any additional negative societal impacts compared to existing papers on PETL.