

FACING THE ELEPHANT IN THE ROOM: VISUAL PROMPT TUNING OR FULL FINETUNING?

**Cheng Han^{1,2}, Qifan Wang³, Yiming Cui⁴, Wenguan Wang⁵, Lifu Huang⁶,
Siyuan Qi⁷, Dongfang Liu^{1*}**

Rochester Institute of Technology¹, University of Missouri - Kansas City², Meta AI³,
University of Florida⁴, Zhejiang University⁵, Virginia Tech⁶, BIGAI^{7†}
{ch7858, dongfang.liu}@rit.edu, wqfcr@fb.com, cuiyiming@ufl.edu,
lifuh@vt.edu, wenguanwang.ai@gmail.com, syqi@bigai.ai

ABSTRACT

As the scale of vision models continues to grow, the emergence of Visual Prompt Tuning (VPT) as a parameter-efficient transfer learning technique has gained attention due to its superior performance compared to traditional full-finetuning. However, the conditions favoring VPT (the “when”) and the underlying rationale (the “why”) remain unclear. In this paper, we conduct a comprehensive analysis across 19 distinct datasets and tasks. To understand the “when” aspect, we identify the scenarios where VPT proves favorable by two dimensions: task objectives and data distributions. We find that VPT is preferable when there is 1) a substantial disparity between the original and the downstream task objectives (*e.g.*, transitioning from classification to counting), or 2) a similarity in data distributions between the two tasks (*e.g.*, both involve natural images). In exploring the “why” dimension, our results indicate VPT’s success cannot be attributed solely to overfitting and optimization considerations. The unique way VPT preserves original features and adds parameters appears to be a pivotal factor. Our study provides insights into VPT’s mechanisms, and offers guidance for its optimal utilization. The code is available at <https://github.com/ChengHan111/VPT-or-FT>.

1 INTRODUCTION

Recent advancements in artificial intelligence [38; 4; 93; 33], especially large models in natural language processing [113; 133; 58; 77; 131], have led to the development of large-scale vision models (*e.g.*, BiT [53], ViT [24], Swin [66], Florence [124]) that have revolutionized various tasks. These models are typically trained on extensive datasets (*e.g.*, ImageNet-21k [89], Open Images [55]) and then finetuned to adapt to specific downstream tasks [47] (*e.g.*, Coco-stuff [7], FGVC [49], VTAB-1k [127]). As models drastically scale up to boost performance, full finetuning that unfreezes all parameters for the update becomes computationally and storage-intensive [111; 14] and impractical. This traditional method can also lead to the loss of valuable knowledge acquired during pretraining, thereby limiting the models’ generalizability [49; 134].

To address this issue, the search for parameter-efficient finetuning methods in vision is an ongoing endeavor. These parameter-efficient approaches freeze most parts of the pretrained model, and only tune the rest or insert customized learnable modules, which significantly reduce the number of learnable parameters compared to full finetuning [16; 48; 69; 85; 129]. Among all these methods, visual prompt tuning [49], which is again inspired by language-domain prompting works [68; 102; 65; 64; 34], stands out as one of the most prominent techniques in this field. VPT introduces a small number of extra trainable parameters (typically less than 1% of model parameters) in the input space of the transformer layers while keeping the backbone frozen. For the first time, it achieves superior performance over full finetuning in image recognition, and is quickly adapted into other visual tasks such as image segmentation [62; 120; 130; 110], image captioning [135], *etc.*

With VPT’s growing popularity [3; 99; 112], a research question naturally arises: *when and why does visual prompt tuning (VPT) outperform full finetuning (FT) as transfer learning paradigm?*

*Corresponding author

†National Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence

To address this question, we conducted extensive experiments on 19 diverse datasets and tasks, wherein VPT outperformed FT in 16 instances. To discern **when** VPT is preferred, we categorized transfer learning scenarios into four groups based on the disparity between the original and downstream tasks, based on task objectives and data distributions (Figure 1). We found that VPT is advantageous when (*Case 1*) a substantial gap exists between the original and downstream tasks (e.g., transitioning from classification to counting) or (*Case 2*) the data distributions are notably similar between the tasks (e.g., both deal with natural images). The size of the downstream task dataset also influences this preference, with FT becoming increasingly favorable as the data size grows.

We further investigate **why** VPT excels. While one plausible hypothesis suggests that FT is more prone to overfitting due to its numerous tunable parameters, our experiments reveal that this is only part of the story. Overfitting is primarily observed in *Case 1* (high task disparity), while in *Case 2*, both methods show no trend for overfitting. A possible explanation is that the additional parameters introduced by VPT offer additional dimensions to escape the local minima of the pretrained model. However, empirical results do not support this assumption when we compare methods with additional dimensions. Our exploration of various tuning method variations underscores the importance of preserving the original feature for achieving superior performance, albeit through a non-obvious pathway. Our results suggest that VPT preserves features and add parameters in a unique manner that is pivotal in this process.

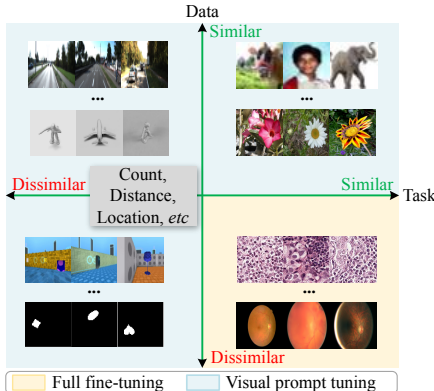


Figure 1: VPT is identified to be preferable in 3 out of 4 transfer learning scenarios when downstream data is limited.

Overall, this study aims to provide a thorough reassessment of the effectiveness of VPT compared to FT. The primary **contributions** of this paper can be summarized as follows:

- We identify the transfer learning scenarios where VPT proves advantageous by considering two critical dimensions: data distributions and task objectives. Notably, VPT is preferable in 3 out of 4 quadrants, with FT gradually closing the performance gap as the downstream data size increases.
- We uncover that overfitting alone does not account for all of VPT’s success. Moreover, the advantage of VPT cannot be attributed solely to additional parameters aiding the model in escaping local minima. The unique introduction of extra parameters plays a key role in VPT’s performance.
- To showcase the efficacy of prompt tuning over full finetuning, we provide attention map visualizations for both methods and demonstrate that visual prompts could enhance feature learning.

2 RELATED WORK

Full finetuning. The emergence of convolutional neural networks (CNNs) (e.g., VGG [96], ResNet [40], MobileNet [45; 90; 44]) makes a significant contribution to the rapid advancements in computer vision. The trained CNN backbones (knowledge) on large-scale datasets can be easily transferred to adapt downstream vision tasks through finetuning (normally full finetuning), which is termed as “pretrain-then-finetune”. This approach allows for a flexible adaptation, standing in contrast to earlier methods that relied on hard-designed task-specific models [19; 5; 20]. Instead of training each task from scratch, it initializes the model with pre-trained weights, updating all the parameters from the original backbone with separate instances for different tasks. Full finetuning is a common practice under this paradigm, which keeps being a powerful yet affordable approach [47; 134; 105; 13] with broad applicability. The subsequent prominent transformer-based architectures in vision (e.g., ViT [24], Swin [66]) inherited the same pipeline for adaptation as training these backbones from scratch can hardly learn global features [24; 12] on small datasets (i.e., Transformers extract features with less inductive bias, which contributes to continuously improve accuracy by increasing training data [26; 24]). However, the limited training data can not fully exploit the capacity). Recent breakthroughs in language for scaling up models in size led to stronger generality [4; 136; 83; 6]. Following this tendency, upcoming vision models (e.g., Florence [124], CoCa [123]) are heavily parameterized, when following the common approach under “pretrain-then-finetune” paradigm, becomes impractical to fully finetune these models. This is primarily due to the

inherent parameter-inefficient nature and the storage-wise expensive requirements of full finetuning. Therefore, parameter-efficient finetuning methods become a critical research area.

Visual Prompt Tuning. With the significant growth in the scale of current vision models, the development of parameter-efficient finetuning methods under “pretrain-then-finetune” paradigm becomes increasingly critical. Current methods can be generally categorized into partial tuning [16; 48; 69], extra module [85; 129; 8] and prompt tuning [49; 50; 23; 118; 125], while partial tuning and extra module face several limitations that hinder their application: First, they generally cannot reach competitive results to full finetuning [49; 16; 48; 69]; Second, some require to insert specific architecture/block design [129; 85; 8] during tuning, rendering them non-universal solutions when considering various backbones. Prompt tuning [59], on the other hand, provides a straightforward solution with strong performance gains during finetuning, which is originally proposed for language-domain [68; 42; 63; 82]. The emergence of visual-related prompt tuning [36; 30; 115] has only recently come to the fore, signaling a new paradigm in parameter-efficient finetuning techniques in the field of computer vision. Generally, prompt tuning prepends extra sets of learnable parameters to the input sequence of backbones, and only updates these parameters during finetuning. While appearing to be simplistic, the paradigm of visual prompt tuning has exhibited satisfactory performance enhancements, as opposed to other parameter-efficient finetuning methods which require specific module designs (*i.e.*, extra module) [85; 129; 8] or coercive freezing of certain portions of the backbone network (*i.e.*, partial tuning) [16; 48; 69]. Moreover, these methods show a substantial performance gap to full finetuning, thereby rendering their widespread implementation under “pretrain-then-finetune” paradigm. Although current visual prompt tuning presents promising results, it remains unexplored in the comprehension of the underlying mechanisms that are responsible for their resounding success. In contrast, in the field of NLP, there are several works [13; 107] that investigate various aspects of prompt tuning (*e.g.*, data scale, evaluation, stability). In light of this view, this paper endeavors to explore and elucidate diverse facets pertaining to visual prompt tuning.

3 METHODOLOGY

Full finetuning is a widely used finetuning method under “pretrain-then-finetune” paradigm. In particular, given a pre-trained model with parameter denoted by θ and downstream dataset D , full finetuning aims to learn a new set of parameters θ' under the same model architecture as in pretraining. The objective of full finetuning is to minimize the corresponding loss function $L(\theta')$ between the predictions $f(x; \theta')$ and the ground truth y by optimizing the parameters θ' over the entire network. While full finetuning obtains good performance with sufficient data, it is computationally expensive and requires the storage of all parameters.

Visual Prompt Tuning, in contrast, keeps the pretrained backbone model frozen while finetuning only a small set of task-specific visual prompts [49; 65; 34; 59; 119; 108], defined as $P = \{P_0, P_1, \dots, P_{N-1}\}$, where P_{i-1} is learnable vectors which are prepended to the input sequence of the i -th layer (as shown in Figure 2(b)). The optimization process of prompt tuning involves updating only the newly added prompts while keeping the pretrained model θ fixed. As a consequence, prompt tuning is particularly beneficial in terms of limiting the overall parameter storage and updating, and preserving the learned knowledge without forgetting during the finetuning phase. In general, for a Vision Transformer [24] (ViT) with N layers L , prompt tuning can be formulated as:

$$\begin{aligned} Z_1 &= L_1(P_0, E), \\ Z_i &= L_i(P_{i-1}, Z_{i-1}) \quad i = 2, 3, \dots, N, \\ y &= \text{Head}(Z_N) \end{aligned} \quad (1)$$

Z_i is the contextual embeddings computed by the i -th encoder layer. E is the image patch embeddings initialized with frozen embedding projection from the backbone. A classification head is then used to map the final layer’s output embedding to the prediction¹. **Trainable** and **frozen** parameters are represented in different colors, respectively.

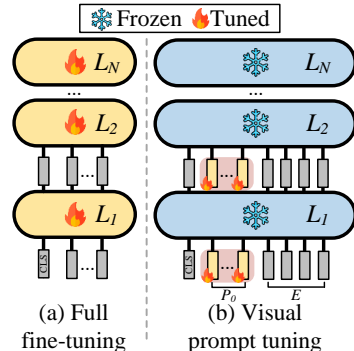


Figure 2: **Full finetuning vs. visual prompt tuning.** Visual prompt tuning only learns a small set of prompts.

¹The mapping of [CLS] token is a general procedure in ViT. However, there are some vision transformer architectures (*e.g.*, Swin [66], CaiT [101]) do not use [CLS] and utilize global pooled Z_N as input for Head.

4 WHEN SHOULD WE CHOOSE VPT?

Our investigation starts with experimental analysis to identify the conditions that VPT are preferable. We first introduce our experiment settings and then present our key findings.

4.1 EXPERIMENT SETUP

Dataset Our experiments are conducted on VTAB-1k [127] image classification benchmark, which encompasses 19 visual tasks categorized into three groups: *Natural*, *Specialized*, and *Structured* (*i.e.*, for image examples, refer to §D). The *Natural* group consists of seven datasets comprising natural images captured by standard cameras. The *Specialized* group includes four datasets that cover images taken by specialized equipment. The *Structured* group contains eight tasks that require geometric comprehensions, such as counting and distance measurement. Unless stated otherwise, we follow the standard experimental setup of [49; 127; 37] and use an 800-200 split for hyperparameter tuning, where 800 samples are used for training and 200 for validation. We evaluate the final performance on the full testing set. For the one-shot experiments, please refer to Appendix §G.4.

Baselines To ensure fair comparisons, we adopt the approach in VPT [49] and focus on the widely-used Vision Transformer (ViT) [24] for image classification. For the pre-training stage, we follow the default procedure outlined in [24; 86] and pre-train the vision Transformer on the ImageNet-21k dataset [89]. Our implementation and experiments are conducted using the publicly available model². It should be noted that there are other backbone models (*e.g.*, Swin [66]) and training objectives (*e.g.*, MAE [41]) that can be used for this study. We have observed similar patterns.

Implementation Details Following [49; 37; 69; 15], we conduct grid search to match the best tuning hyperparameters, learning rate (*i.e.*, [50, 25, 10, 5, 2.5, 1, 0.5, 0.25, 0.1, 0.05]), and weight decay (*i.e.*, [0.01, 0.001, 0.0001, 0.0]) on val set for each task. Notably, we follow [37] and *do not cover* specific-designed large learning rate in [49] for general propose. In the training of all models, a cosine decay policy is employed to schedule the learning rate. The total number of training epochs is set to be 100 (including 10 warm-up epochs). We follow the same batch size setting: 64/128 for ViT-B/16. Our method is implemented in Pytorch [80]. Experiments are conducted on NVIDIA A100-80GB GPUs. For reproducibility, our full implementation will be publicly released.

4.2 INITIAL EXPERIMENTAL RESULTS

The initial results of FT and VPT are presented in the first two rows of Table 1 and Table 2 (*i.e.*, the last two rows: Mixed and FT-then-PT, will be introduced and compared to FT and VPT in §5.2). Out of all 19 datasets/tasks, VPT performs better on 16 instances, including 6/7 Natural datasets, 2/4 Specialized datasets, and 8/8 Structured datasets.

Table 1: **Image classification accuracy on various training strategies on VTAB-1k [127] *Natural* and *Specialized*** for ViT-B/16 [24] pretrained on supervised ImageNet-21k [89]. Underlines indicate the better results between FT and VPT, and **bolds** indicates the best results among all variants. We report five runs to test average accuracy instead of three in [49] to consider more randomness, and “Number of Wins” in (·) compared to full finetuning (Full) [47].

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Natural</i> [7]							Mean	VTAB-1k [127] <i>Specialized</i> (4)				Mean
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397		Patch	Camelyon	EuroSAT	Resisc45	
FT [47]	64.5	88.1	65.0	97.3	84.6	87.3	39.5	75.19	81.2	95.7	83.2	<u>73.3</u>	83.48
VPT [49]	<u>77.7</u>	90.2	68.8	98.1	88.4	82.5	51.0	79.53(6)	77.9	96.2(9)	83.4	73.1	82.65(1)
Mixed	61.5	86.2	60.2	95.6	85.0	85.7	35.4	72.80(1)	77.8	96.2(7)	74.0	71.0	79.78(1)
FT-then-PT	85.9	89.0	66.3	95.9	86.3	87.1	36.9	78.20(4)	79.4	95.6	81.8	74.0	82.70(1)

Table 2: **Image classification accuracy on various training strategies on VTAB-1k [127] *Structured*** for ViT-B/16 [24] pretrained on supervised ImageNet-21k [89].

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Structured</i> [8]							Mean	
	Clevr/cnt.	Clevr/dist.	DMLab KITTI/dist.	dSprites/loc.	dSprites/ori.	SmallNORB/az.	SmallNORB/elev.		
FT [47]	55.4	58.2(2)	40.4	74.7	54.0	47.0	26.2	28.6	48.07
VPT [49]	68.2	<u>58.2(3)</u>	45.3	77.7	78.4	48.4	31.2	41.3	56.09(8)
Mixed	57.9	54.6	37.7	69.2	20.9	45.2	28.5	29.5	42.94(3)
FT-then-PT	63.6	61.3	39.1	73.0	55.1	51.6	27.6	31.7	50.38(6)

²https://github.com/google-research/vision_transformer

4.3 CHOOSE VPT FOR THREE QUADRANTS OF TRANSFER LEARNING

Upon initial examination, these results do not reveal a distinct pattern indicating when VPT outperforms FT. However, upon deeper reflection on the nature of transfer learning, the effectiveness of finetuning methods under the “pretrain-then-finetune” paradigm can be significantly impacted by the disparities between pretraining and finetuning. These disparities can be attributed to two key aspects: 1) the data distributions in the pretraining and finetuning datasets are dissimilar, and 2) the nature of the downstream task objective is different from the pretraining one.

Hence we analyze the datasets by categorizing them into four quadrants along these two dimensions. To measure the discrepancy in image distribution, we adopt the Fréchet Inception Distance (FID) [17; 56], which is known to correspond well with human judgments of image quality and diversity [116; 79]. In terms of task disparity, tasks in the *Natural* and *Specialized* subsets are closely related to image classification and thus have low disparities, while tasks in *Structured*, such as counting and distance, are regarded as distinct from image classification.

The FID scores with respect to the better finetuning method of all 19 tasks under the default train-val split are shown in Figure 3. The figure shows that VPT reaches superior performance than FT in two different scenarios: (1) when the disparity between the task objectives of the original and downstream tasks is high, or (2) when the data distributions are similar. Only for those tasks with low disparity (the left 11), a relatively larger FID score (dissimilar datasets) typically leads to higher performance of full finetuning compared to prompt tuning.³

4.4 GAP NARROWS AS DOWNSTREAM DATASETS EXPAND

The above observation holds for scenarios when we have a limited amount of data for the downstream tasks. However, the effectiveness of different finetuning methods may vary in high-resource scenarios. To investigate this, we conducted a comprehensive evaluation of FT and VPT across different tuning set sizes. Specifically, we re-create the training and validation sets that keep a consistent 4:1 ratio, and vary the number of training samples from 400 to 20,000 to explore the behavior of the FT and VPT when the tuning data size grows. It is important to note that not all datasets support all combinations of splits due to insufficient data samples (e.g., VTAB-1k *Natural* Oxford Flowers102 [73] has only 1020 images for both `train` and `val`).

The performance of full finetuning and prompt tuning across different training set sizes and datasets is presented in Figure 4. The figure shows that the performance gap between full finetuning and prompt tuning decreases as the amount of training data increases, with full finetuning even outperforming prompt tuning under high-resource settings. These results demonstrate the robustness and generality of prompt tuning in the face of limited finetuning examples. Although full finetuning generally achieves higher accuracy when rich data examples are available, the prompt tuning still

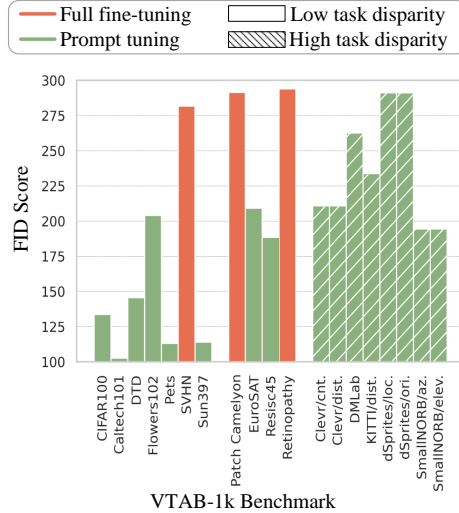


Figure 3: **Overall FID score with respect to win/loss of visual prompt tuning on VTAB-1k benchmark** categorized into *Natural*, *Specialized* and *Structured*, respectively. ■ ■ colors represent the method with higher accuracy. Under the same train-val split and low task disparity, a higher FID score might potentially lead to relatively higher accuracy on full finetuning, and vice versa. Solid filled represents that the disparity between the target task and the pre-trained task is small while slash filled means a large disparity (e.g., distance, azimuth, counting). The FID scores show significant robustness in repeat runs (i.e., $std < 0.5\%$), we therefore do not present error bars here.

³Here the range of FID scores reported appears to be larger than existing approaches [99; 17; 79; 128]. This is because current methods typically report the distance between real and generated images, whereas our analysis involves the computation of distances between real datasets with significantly different distributions.

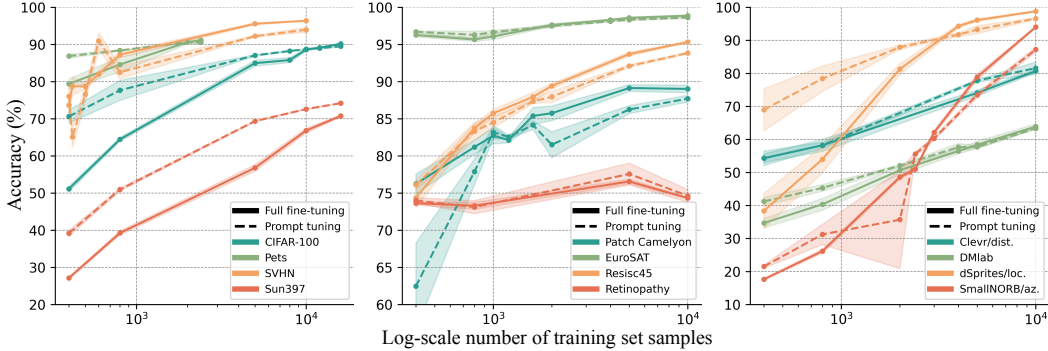


Figure 4: **Analysis of dataset capacity** on VTAB-1k *Natural* (left), *Specialized* (middle) and *Structured* (right), respectively. For each group, we select four datasets and plot accuracy plots on FT (*i.e.*, solid lines) and VPT (*i.e.*, dotted lines). We take the log scale of training data samples for better separation and each color stands for an individual classification task. Each point is given by average over five runs. In general, with the dataset increasing in size, the performance gap between FT and VPT becomes narrow. FT even surpasses VPT in 9 of 12 cases in this plot with the increasing of data samples (the same tendency takes place in other datasets). For per-task accuracy tables among different dataset scales and detailed FT and VPT accuracy plots, see the supplementary material §A.

Table 3: **One-shot image classification accuracy on VTAB-1k [127] *Natural* and *Specialized*** for ViT-B/16 [24] pretrained on supervised ImageNet-21k [89]. Experiments show that VPT outperforms the FT in **13 of 19** instances under the one-shot learning setting, which further suggests the advantage of VPT in limited data cases. “Tuned/Total” is the average percentage of tuned parameters required by 19 tasks.

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Natural</i> [7]							Mean	VTAB-1k [127] <i>Specialized</i> [4]					Mean
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397		Patch Camelyon	EuroSAT	Resisc45	Retinopathy		
FT [47]	37.4	67.9	25.3	80.2	37.8	12.7	30.3	41.66	58.2	44.1	42.3	32.6	44.30	
VPT [49]	54.6	78.4	37.4	94.9	70.5	11.6	51.2	56.94(6)	68.4	45.6	41.1	34.7	47.45(3)	
- Tuned / Total	0.20%	0.20%	0.15%	0.10%	0.04%	0.54%	0.41%	0.23%	1.06%	1.07%	0.15%	0.02%	0.57%	

reaches a competitive performance with much fewer parameters. This observation aligns with current research in NLP prompt tuning [13; 22] and suggests that prompt tuning techniques have great potential in both low- and high-resource scenarios.

To further test the effectiveness of the two methods in low-resource settings, we conducted one-shot learning experiments where each class is learned from a single labeled example [104; 98]. These experiments represent an extreme case in which we seek to evaluate the methods’ ability to adapt to novel tasks after having been exposed to only a small amount of training data. Table 3 and Table 4 show the experimental results, with prompt tuning outperforming full finetuning in **13 out of 19** tasks and achieving substantially higher accuracy in some cases (*i.e.*, 41.66% *vs* **56.94%** in VTAB-1k *Natural*). These results favor our assumption that prompt tuning is more effective than full finetuning for the aforementioned three quadrants when only limited finetuning examples are available. We observe that in only half of the Structured cases, VPT performs better. Overall, the accuracies for one-shot experiments are comparatively low. Hence it should be noted that randomness could play a role in one-shot experiments, and the results might not reliably conclusive.

5 WHY DOES VPT OUTPERFORM FT?

5.1 OVERFITTING IS PART OF THE REASON

After identifying the conditions that VPT would be favorable over FT, we dive deeper to explore the underlying reasons. One hypothesis is that FT optimizes all the parameters in the backbone, which may easily overfit to the downstream task, especially when limited data is available [121; 21; 39].

Figure 5 presents training and testing loss curves, up to 100 epochs, for six representative tasks in VTAB-1k [127], comprising of 2 tasks for from each kind (*Natural*, *Specialized*, and *Structured*). Comprehensive results covering all 19 tasks are available in the supplementary material §A.

We observe that for 8 out of 9 *Structured* tasks, both FT and VPT overfit (*i.e.*, testing error increases as training progresses). These tasks, such as counting and angular measurement, are very different

Table 4: **One-shot image classification accuracy on VTAB-1k [127] Structured** for ViT-B/16 [24] pretrained on supervised ImageNet-21k [89].

ViT-B/16 [24] (85.8M)	VTAB-1k [127] Structured [8]								Mean
	Clevr/cnt.	Clevr/dist.	DMLab KITT/dist.	dSprites/loc.	dSprites/ori.	SmallNORB/az.	SmallNORB/elev.		
FT [47]	17.8	21.0	17.1	40.5	9.0	11.4	7.4	12.7	17.11
VPT [49]	21.5	22.1	18.7	37.0	6.7	7.5	6.2	14.9	16.83(4)
- Tuned / Total	0.54%	2.11%	1.07%	0.54%	0.12%	0.55%	2.12%	2.11%	1.14%

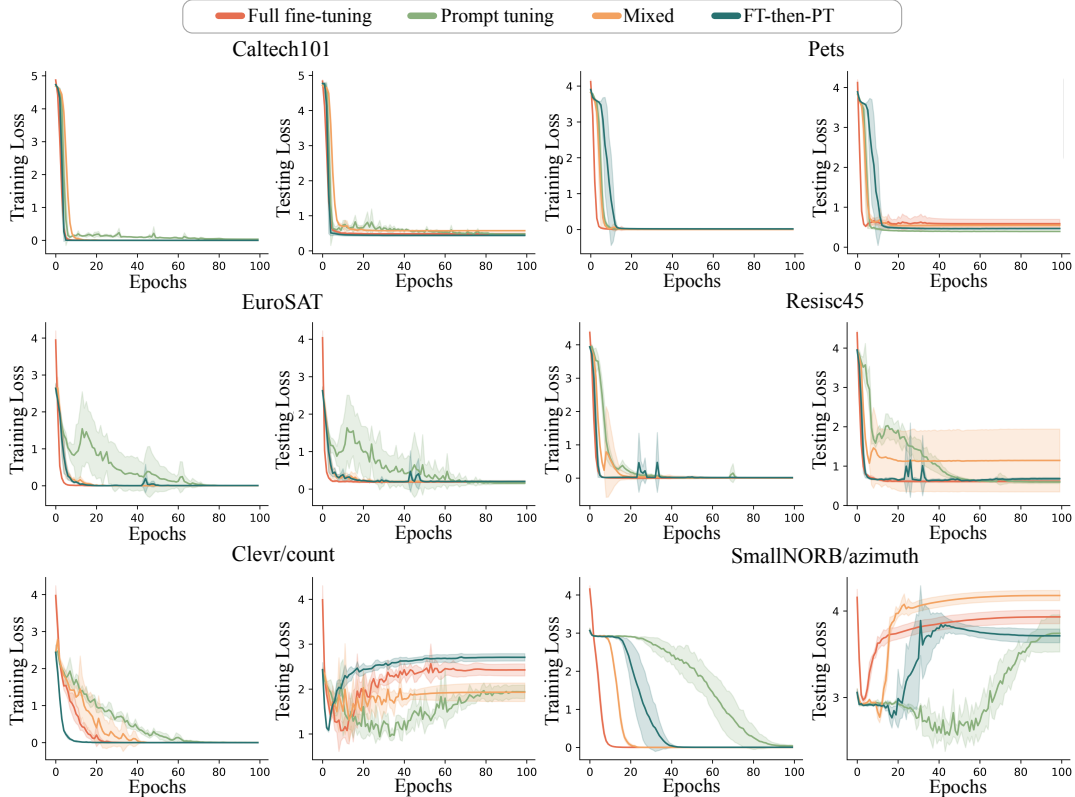


Figure 5: **Training/testing loss curves** of six datasets from VTAB-1k. █ █ █ █ colors represent four training strategies: full finetuning, prompt tuning, mixed, and FT-then-PT, respectively. We show two representative tasks from VTAB-1k *Natural*, *Specialized*, and *Structured*, respectively. Full results and log scale results are presented in the supplementary material §A.

from image classification, and may require more training data for model adaptation. We conducted additional experiments by enlarging the training set (to 5000 and 10000) of these tasks in Figure 6. It can be observed that with the increase of training samples, the phenomenon of overfitting is mitigated in both methods.

However, **among the remaining 10 tasks within the categories of *Natural* and *Specialized*, overfitting behavior is observed in only 1 task**. In other words, training and testing losses consistently decrease for both methods in scenarios where the task objectives are similar between the original and downstream tasks. In such cases, overfitting is not the cause of performance degradation of FT.

5.2 OPTIMIZATION DOES NOT PARTICULARLY FAVOR VPT

Another possible explanation for why prompt tuning can achieve superior performance is that it is the additional dimensions introduced by the prompts that help the loss to be further optimized during the tuning process. This hypothesis is illustrated by Figure 7.

To test our hypothesis in depth, we perform experiments on two additional variants of tuning strategies besides full finetuning and prompt tuning. The first method is noted *Mixed*, where we add visual prompts and perform full finetuning on the entire network instead of focusing only on network parameters or prompt vectors in isolation. The second method is called *FT-then-PT*, where

we first perform full finetuning on the network for the downstream task and then prepend additional prompts and perform prompt tuning for the same task. This approach applies 200 epochs in total.

We compare the above methods with full finetuning and visual prompt tuning in Figure 5. Per-task performance is also detailed in Table 1 and Table 2. We observe that in general, *VPT* demonstrates the highest performance. The *FT* approach and the *FT-then-PT* approach yield similar results, albeit with a noticeable gap compared to *VPT*. In contrast, the *Mixed* approach exhibits the lowest performance among the evaluated methods.

These results indicate several key findings. (1) The presence of additional dimensions **does not** significantly aid the optimization process in escaping local minima. If this were the case, we would expect either the *Mixed* or the *FT-then-PT* method to outperform *FT*, but this is not observed. (2) Preserving the original feature proves to be crucial for transferring to downstream tasks, as evidenced by the superior performance of *VPT* compared to *FT-then-PT*. (3) Maintaining a fixed feature space in *VPT* may compel it to learn more valuable embeddings with the trainable prompts, as it significantly outperforms the *Mixed* method. These observations lead to an important insight: pretrained parameters play a pivotal role in capturing general features, while the added parameters are important for potentially encoding task information in the context of transfer learning.

5.3 FURTHER OBSERVATIONS

Overall, our observations indicate that the preservation of initial features is important for *VPT*'s success, but in a very sophisticated manner. Notably, *VPT* surpasses more straightforward methods of feature preservation, such as retraining the final linear probe layer or incorporating a multi-layer probe [37], even though all these methods maintain the originally learned features.

This potentially explains cases in which data distributions exhibit similarity, but not in scenarios characterized by distinct data distributions and high task disparity. In these instances, we hypothesize that the substantial task divergence requires a stronger need for additional finetuning examples to achieve comprehensive adaptation, whereas prompt tuning demonstrates better generalization. The introduced prompts serve as guiding mechanisms for the model to enhance task-specific feature learning, a pivotal aspect of task encoding, particularly in situations with limited finetuning data.

It's also worth discussing scenarios in which *FT* outperforms *VPT*, specifically in situations involving similar tasks with varying data distributions. Our hypothesis is that the feature representation captured by the pretrained model may not be well-suited for the downstream task due to significant differences in data distribution. Consequently, updating all model parameters through full finetuning becomes more effective for learning the feature representation within

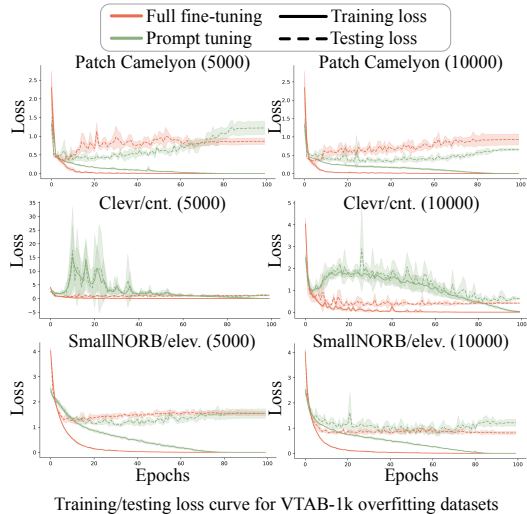


Figure 6: **Training/testing loss curves for VTAB-1k Structured datasets with larger training sets.** █ █ colors represent full finetuning and prompt tuning, respectively. Solid lines are the training loss curves and dashed lines show the corresponding testing curves. We select three tasks that suffer from overfitting and enlarge the dataset train-val split to 5000-1250 (left) and 10000-2500 (right), respectively. Full results are shown in the supplementary material §A.

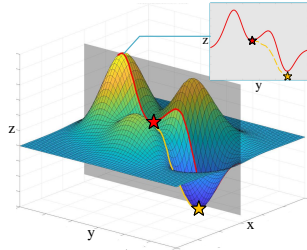


Figure 7: Hypothesis of extra dimensions helping the pre-trained model to be further optimized, *e.g.*, by escaping local minima. In this illustration, pre-training the original 1D function (the red line) might get stuck in the middle (the red star). When a second dimension is introduced during finetuning, it is easier to be better optimized (from the yellow line to the yellow star). However, we found that this hypothesis does not hold.

the new distribution. Conversely, when finetuning data is similar, there may be no need to modify the backbone model extensively; instead, a minimal set of prompts can be inserted for adaptation.

5.4 VISUALIZING THE EFFECT OF VPT ON FEATURE LEARNING

To the best of our knowledge, the visual explanations of visual prompt tuning are found to be rare [60; 75; 71]. In light of this view, we seek to investigate whether prompt tuning can offer actual visualization meanings and support stronger visualization explanations, particularly in cases where it outperforms full finetuning. In this experiment, we apply gradient-weighted class activation mapping (GradCAM⁴) [92], which is a popular technique for producing visual explanations for decisions [61; 109; 2; 10; 35; 29; 11]. In general, it utilizes the gradients of a target concept and channels them through the final layer of the network to generate a coarse localization map that accentuates the salient regions in the image which contribute significantly towards predicting the concept of interest.

Figure 8 presents examples that full finetuning fails to recognize while prompt tuning makes correct classification during testing. It can be seen from the results that a clear concept of interest can be observed from both methods through a heat map. In these prompt tuning success cases, we can see straightforward visual explanation differences (*e.g.*, take the bicycle image in Figure 8 upper right as an example, when full finetuning fails to make a correct decision, prompt tuning instead recognize the bicycle with its structural features successfully), which suggests that prompt tuning effectively guides the model to pay more attention to the important regions of the images, and thus is capable of enhancing the learning of stronger features by leveraging domain-specific information and patterns that might not be well captured in full finetuning. An additional visual explanation technique (*i.e.*, Integrated Gradients [100]) and more visual inspection examples on GradCAM are also included in the supplementary material §B.

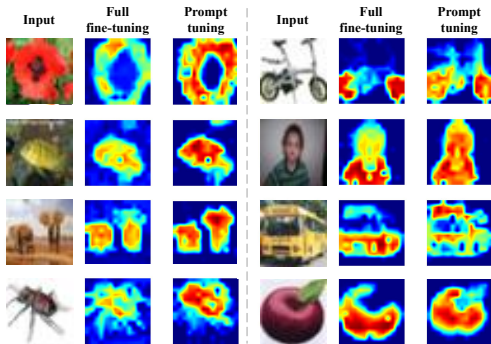


Figure 8: **Visual inspection of full finetuning and prompt tuning** using GradCAM [92]. Note that red regions correspond to a high score for the class. From left to right are input image after standard data augmentation, GradCAM results for full finetuning and GradCAM results for prompt tuning. Figure best viewed in color.

6 CONCLUSION AND DISCUSSION

When models scale up, how shall we face the elephant in the room? Driven by this question, we aim to provide a justifiable recommendation regarding the choice between full finetuning or prompt tuning during training. In this pursuit, we conduct an extensive study across 19 datasets, comparing full finetuning with visual prompt tuning, to examine various hypotheses regarding the effectiveness of visual prompt tuning. Our experimental results *show* that overfitting is not the root cause of full finetuning’s inferior performance to prompt tuning. We further notice that dataset disparities might be a potential factor in the behavior of full finetuning and prompt tuning. Also, prompt tuning shows its strong generality with limited data, while full finetuning catches up or surpasses prompt tuning with more data presented. Attention map visualization suggests that the visual prompts are capable of enhancing the learning of features that might not be well captured in full finetuning. We *suggest* that prompt tuning should be applied consciously under “pretrain-then-finetune” paradigm. Task disparities and dataset scale are two main factors that determine the most suitable way for downstream finetuning. In spite of our comprehensive coverage of various datasets in image recognition, it is noteworthy that other visual-related tasks for finetuning (*e.g.*, semantic segmentation [62; 28]) have not been fully explored. This observation motivates us for future investigation of such scenarios.

⁴<https://github.com/pytorch/captum>

REPRODUCIBILITY STATEMENT

To help readers reproduce our results, we have described the implementation details in §4.1 and §G. We will release our source code after acceptance. All the datasets we use are publicly available.

ACKNOWLEDGMENT

This research was supported by the National Science Foundation under Grant No. 2242243.

REFERENCES

- [1] Sercan Ö Arik and Tomas Pfister. Tabnet: Attentive interpretable tabular learning. In *AAAI*, 2021.
- [2] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, 58, 2020.
- [3] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Visual prompting: Modifying pixel space to adapt pre-trained models. *arXiv preprint arXiv:2203.17274*, 2022.
- [4] 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.
- [5] Leo Breiman. Random forests. *Machine learning*, 45, 2001.
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, 2020.
- [7] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *CVPR*, 2018.
- [8] Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. Tinytl: Reduce memory, not parameters for efficient on-device learning. In *NeurIPS*, 2020.
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021.
- [10] Diogo V Carvalho, Eduardo M Pereira, and Jaime S Cardoso. Machine learning interpretability: A survey on methods and metrics. *Electronics*, 8(8):832, 2019.
- [11] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *CVPR*, 2021.
- [12] Bin Chen, Ran Wang, Di Ming, and Xin Feng. Vit-p: Rethinking data-efficient vision transformers from locality. *arXiv preprint arXiv:2203.02358*, 2022.
- [13] Guanzheng Chen, Fangyu Liu, Zaiqiao Meng, and Shangsong Liang. Revisiting parameter-efficient tuning: Are we really there yet? *arXiv preprint arXiv:2202.07962*, 2022.
- [14] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. *arXiv preprint arXiv:2205.13535*, 2022.
- [15] Xiang Chen, Ningyu Zhang, Xin Xie, Shumin Deng, Yunzhi Yao, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. Knowprompt: Knowledge-aware prompt-tuning with synergistic optimization for relation extraction. In *ACM Web Conference*, 2022.

- [16] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021.
- [17] Min Jin Chong and David Forsyth. Effectively unbiased fid and inception score and where to find them. In *CVPR*, 2020.
- [18] Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. What does bert look at? an analysis of bert’s attention. *arXiv preprint arXiv:1906.04341*, 2019.
- [19] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20, 1995.
- [20] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1), 1967.
- [21] Tom Dietterich. Overfitting and undercomputing in machine learning. *ACM Computing Surveys*, 27(3), 1995.
- [22] 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*, 2023.
- [23] Shaohua Dong, Yunhe Feng, Qing Yang, Yan Huang, Dongfang Liu, and Heng Fan. Efficient multimodal semantic segmentation via dual-prompt learning. *arXiv preprint arXiv:2312.00360*, 2023.
- [24] 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 *ICLR*, 2021.
- [25] Frédéric Dufaux. Grand challenges in image processing, 2021.
- [26] Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Lev-ent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *ICML*, 2021.
- [27] Dumitru Erhan, Yoshua Bengio, Aaron Courville, and Pascal Vincent. Visualizing higher-layer features of a deep network. *University of Montreal*, 2009.
- [28] Marc Fischer, Alexander Bartler, and Bin Yang. Prompt tuning for parameter-efficient medical image segmentation. *arXiv preprint arXiv:2211.09233*, 2022.
- [29] Shang-Hua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip Torr. Res2net: A new multi-scale backbone architecture. *IEEE TPAMI*, 2019.
- [30] Yunhe Gao, Xingjian Shi, Yi Zhu, Hao Wang, Zhiqiang Tang, Xiong Zhou, Mu Li, and Dimitris N Metaxas. Visual prompt tuning for test-time domain adaptation. *arXiv preprint arXiv:2210.04831*, 2022.
- [31] Timnit Gebru, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, and Li Fei-Fei. Fine-grained car detection for visual census estimation. In *AAAI*, 2017.
- [32] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012.
- [33] Yolanda Gil, Daniel Garijo, Deborah Khider, Craig A Knoblock, Varun Ratnakar, Maximiliano Osorio, Hernán Vargas, Minh Pham, Jay Pujara, Basel Shbita, et al. Artificial intelligence for modeling complex systems: taming the complexity of expert models to improve decision making. *ACM Transactions on Interactive Intelligent Systems*, 2021.
- [34] Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*, 2021.

- [35] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 2018.
- [36] Meng-Hao Guo, Cheng-Ze Lu, Zheng-Ning Liu, Ming-Ming Cheng, and Shi-Min Hu. Visual attention network. *arXiv preprint arXiv:2202.09741*, 2022.
- [37] Cheng Han, Qifan Wang, Yiming Cui, Zhiwen Cao, Wenguan Wang, Siyuan Qi, and Dongfang Liu. E²vpt: An effective and efficient approach for visual prompt tuning. In *ICCV*, 2023.
- [38] Antoine L Harfouche, Daniel A Jacobson, David Kainer, Jonathon C Romero, Antoine H Harfouche, Giuseppe Scarascia Mugnozza, Menachem Moshelion, Gerald A Tuskan, Joost JB Keurentjes, and Arie Altman. Accelerating climate resilient plant breeding by applying next-generation artificial intelligence. *Trends in Biotechnology*, 37(11), 2019.
- [39] Douglas M Hawkins. The problem of overfitting. *Journal of Chemical Information and Computer Sciences*, 44(1), 2004.
- [40] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [41] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022.
- [42] 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 *ICML*, 2022.
- [43] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 2017.
- [44] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *ICCV*, 2019.
- [45] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [46] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [47] Eugenia Iofinova, Alexandra Peste, Mark Kurtz, and Dan Alistarh. How well do sparse imagenet models transfer? In *CVPR*, 2022.
- [48] Menglin Jia, Zuxuan Wu, Austin Reiter, Claire Cardie, Serge Belongie, and Ser-Nam Lim. Exploring visual engagement signals for representation learning. In *ICCV*, 2021.
- [49] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, 2022.
- [50] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *ECCV*, 2022.
- [51] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. Novel dataset for fine-grained image categorization: Stanford dogs. In *CVPR Workshop*, 2011.
- [52] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *ICML*, 2017.

- [53] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *ECCV*, 2020.
- [54] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical report*, 2009.
- [55] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.
- [56] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of imagenet classes in fr\`echet inception distance. In *ICLR*, 2023.
- [57] Thibault Laugel, Marie-Jeanne Lesot, Christophe Marsala, Xavier Renard, and Marcin Detryniecki. The dangers of post-hoc interpretability: Unjustified counterfactual explanations. In *IJCAI*, 2019.
- [58] Franck Le, Mudhakar Srivatsa, Raghu Ganti, and Vyas Sekar. Rethinking data-driven networking with foundation models: challenges and opportunities. In *ACM Workshop*, 2022.
- [59] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *EMNLP*, 2021.
- [60] Aodi Li, Liansheng Zhuang, Shuo Fan, and Shafei Wang. Learning common and specific visual prompts for domain generalization. In *ACCV*, 2022.
- [61] Pantelis Linardatos, Vasilis Papastefanopoulos, and Sotiris Kotsiantis. Explainable ai: A review of machine learning interpretability methods. *Entropy*, 23(1), 2020.
- [62] Lingbo Liu, Bruce XB Yu, Jianlong Chang, Qi Tian, and Chang-Wen Chen. Prompt-matched semantic segmentation. *arXiv preprint arXiv:2208.10159*, 2022.
- [63] 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*, 2023.
- [64] Xiao Liu, Kaixuan Ji, Yicheng Fu, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*, 2021.
- [65] 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 *ACL*, 2022.
- [66] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021.
- [67] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *NeurIPS*, 2017.
- [68] Fang Ma, Chen Zhang, Lei Ren, Jingang Wang, Qifan Wang, Wei Wu, Xiaojun Quan, and Dawei Song. Xprompt: Exploring the extreme of prompt tuning. In *EMNLP*, 2022.
- [69] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *ECCV*, 2018.
- [70] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In *CVPR*, 2015.

- [71] Chengzhi Mao, Revant Teotia, Amrutha Sundar, Sachit Menon, Junfeng Yang, Xin Wang, and Carl Vondrick. Doubly right object recognition: A why prompt for visual rationales. *arXiv preprint arXiv:2212.06202*, 2022.
- [72] John Merrill, Geoff Ward, Sean Kamkar, Jay Budzik, and Douglas Merrill. Generalized integrated gradients: A practical method for explaining diverse ensembles. *arXiv preprint arXiv:1909.01869*, 2019.
- [73] M-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*, 2008.
- [74] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*, 2008.
- [75] Changdae Oh, Hyeji Hwang, Hee-young Lee, YongTaek Lim, Geunyoung Jung, Jiyoung Jung, Hosik Choi, and Kyungwoo Song. Blackvip: Black-box visual prompting for robust transfer learning. *arXiv preprint arXiv:2303.14773*, 2023.
- [76] Samet Oymak, Ankit Singh Rawat, Mahdi Soltanolkotabi, and Christos Thrampoulidis. On the role of attention in prompt-tuning. *arXiv preprint arXiv:2306.03435*, 2023.
- [77] Gerhard Paaß and Sven Giesselbach. Foundation models for natural language processing—pre-trained language models integrating media. *arXiv preprint arXiv:2302.08575*, 2023.
- [78] Geondo Park, June Yong Yang, Sung Ju Hwang, and Eunho Yang. Attribution preservation in network compression for reliable network interpretation. In *NeurIPS*, 2020.
- [79] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *CVPR*, 2022.
- [80] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- [81] Zheyun Qin, Cheng Han, Qifan Wang, Xiushan Nie, Yilong Yin, and Xiankai Lu. Unified 3d segmenter as prototypical classifiers. In *NeurIPS*, 2023.
- [82] Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 2020.
- [83] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 2020.
- [84] Sahana Ramnath, Preksha Nema, Deep Sahni, and Mitesh M Khapra. Towards interpreting bert for reading comprehension based qa. *arXiv preprint arXiv:2010.08983*, 2020.
- [85] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. In *NeurIPS*, 2017.
- [86] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021.
- [87] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- [88] Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong. Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistic Surveys*, 16:1–85, 2022.

- [89] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *IJCV*, 2015.
- [90] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018.
- [91] Maximilian Seitzer. pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid>, August 2020. Version 0.3.0.
- [92] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 2017.
- [93] Bhavin J Shastri, Alexander N Tait, Thomas Ferreira de Lima, Wolfram HP Pernice, Harish Bhaskaran, C David Wright, and Paul R Prucnal. Photonics for artificial intelligence and neuromorphic computing. *Nature Photonics*, 15(2), 2021.
- [94] Baifeng Shi, Siyu Gai, Trevor Darrell, and Xin Wang. Toast: Transfer learning via attention steering. *arXiv preprint arXiv:2305.15542*, 2023.
- [95] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *ICML*, 2017.
- [96] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [97] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [98] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *NeurIPS*, 2017.
- [99] Kihyuk Sohn, Yuan Hao, José Lezama, Luisa Polania, Huiwen Chang, Han Zhang, Irfan Essa, and Lu Jiang. Visual prompt tuning for generative transfer learning. *arXiv preprint arXiv:2210.00990*, 2022.
- [100] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *ICML*, 2017.
- [101] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *ICCV*, 2021.
- [102] Lifu Tu, Caiming Xiong, and Yingbo Zhou. Prompt-tuning can be much better than fine-tuning on cross-lingual understanding with multilingual language models. *arXiv preprint arXiv:2210.12360*, 2022.
- [103] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *CVPR*, 2015.
- [104] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *NeurIPS*, 2016.
- [105] David Vos, Till Döhmen, and Sebastian Schelter. Towards parameter-efficient automation of data wrangling tasks with prefix-tuning. In *NeurIPS Workshop*, 2022.
- [106] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [107] Chaozheng Wang, Yuanhang Yang, Cuiyun Gao, Yun Peng, Hongyu Zhang, and Michael R Lyu. No more fine-tuning? an experimental evaluation of prompt tuning in code intelligence. In *ACM ESEC/FSE*, 2022.

- [108] Qifan Wang, Yuning Mao, Jingang Wang, Hanchao Yu, Shaoliang Nie, Sinong Wang, Fuli Feng, Lifu Huang, Xiaojun Quan, Zenglin Xu, et al. Aprompt: Attention prompt tuning for efficient adaptation of pre-trained language models. In *EMNLP*, 2023.
- [109] Wenguan Wang, Cheng Han, Tianfei Zhou, and Dongfang Liu. Visual recognition with deep nearest centroids. In *ICLR*, 2023.
- [110] Wenxuan Wang, Jiachen Shen, Chen Chen, Jianbo Jiao, Yan Zhang, Shanshan Song, and Jiangyun Li. Med-tuning: Exploring parameter-efficient transfer learning for medical volumetric segmentation. *arXiv preprint arXiv:2304.10880*, 2023.
- [111] Zhen Wang, Rameswar Panda, Leonid Karlinsky, Rogerio Feris, Huan Sun, and Yoon Kim. Multitask prompt tuning enables parameter-efficient transfer learning. *arXiv preprint arXiv:2303.02861*, 2023.
- [112] Jingyuan Wen, Yutian Luo, Nanyi Fei, Guoxing Yang, Zhiwu Lu, Hao Jiang, Jie Jiang, and Zhao Cao. Visual prompt tuning for few-shot text classification. In *ICCL*, 2022.
- [113] Walter F Wiggins and Ali S Tejani. On the opportunities and risks of foundation models for natural language processing in radiology. *Radiology: Artificial Intelligence*, 4(4), 2022.
- [114] Cheng-En Wu, Yu Tian, Haichao Yu, Heng Wang, Pedro Morgado, Yu Hen Hu, and Linjie Yang. Why is prompt tuning for vision-language models robust to noisy labels? In *ICCV*, 2023.
- [115] Yinghui Xing, Qirui Wu, De Cheng, Shizhou Zhang, Guoqiang Liang, and Yanning Zhang. Class-aware visual prompt tuning for vision-language pre-trained model. *arXiv preprint arXiv:2208.08340*, 2022.
- [116] Qiantong Xu, Gao Huang, Yang Yuan, Chuan Guo, Yu Sun, Felix Wu, and Kilian Weinberger. An empirical study on evaluation metrics of generative adversarial networks. *arXiv preprint arXiv:1806.07755*, 2018.
- [117] Wanqi Xue and Wei Wang. One-shot image classification by learning to restore prototypes. In *AAAI*, 2020.
- [118] Liqi Yan, Cheng Han, Zenglin Xu, Dongfang Liu, and Qifan Wang. Prompt learns prompt: exploring knowledge-aware generative prompt collaboration for video captioning. In *IJCAI*, 2023.
- [119] Li Yang, Qifan Wang, Jingang Wang, Xiaojun Quan, Fuli Feng, Yu Chen, Madian Khabza, Sinong Wang, Zenglin Xu, and Dongfang Liu. Mixpave: Mix-prompt tuning for few-shot product attribute value extraction. In *ACL*, 2023.
- [120] Senqiao Yang, Jiarui Wu, Jiaming Liu, Xiaoqi Li, Qizhe Zhang, Mingjie Pan, and Shanghang Zhang. Exploring sparse visual prompt for cross-domain semantic segmentation. *arXiv preprint arXiv:2303.09792*, 2023.
- [121] Xue Ying. An overview of overfitting and its solutions. In *Journal of Physics*, volume 1168, pp. 022022, 2019.
- [122] Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. Gnnexplainer: Generating explanations for graph neural networks. In *NeurIPS*, 2019.
- [123] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- [124] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- [125] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and language prompt learning. *arXiv preprint arXiv:2210.07225*, 2022.

- [126] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014.
- [127] 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.
- [128] Gongjie Zhang, Kaiwen Cui, Tzu-Yi Hung, and Shijian Lu. Defect-gan: High-fidelity defect synthesis for automated defect inspection. In *WACV*, 2021.
- [129] Jeffrey O Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, and Jitendra Malik. Side-tuning: a baseline for network adaptation via additive side networks. In *ECCV*, 2020.
- [130] Yabo Zhang, Zihao Wang, Jun Hao Liew, Jingjia Huang, Manyu Zhu, Jiashi Feng, and Wang-meng Zuo. Associating spatially-consistent grouping with text-supervised semantic segmentation. *arXiv preprint arXiv:2304.01114*, 2023.
- [131] Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. How does nlp benefit legal system: A summary of legal artificial intelligence. *arXiv preprint arXiv:2004.12158*, 2020.
- [132] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016.
- [133] Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, et al. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt. *arXiv preprint arXiv:2302.09419*, 2023.
- [134] Han Zhou, Xingchen Wan, Ivan Vulić, and Anna Korhonen. Autopeft: Automatic configuration search for parameter-efficient fine-tuning. *arXiv preprint arXiv:2301.12132*, 2023.
- [135] Peipei Zhu, Xiao Wang, Lin Zhu, Zhenglong Sun, Wei-Shi Zheng, Yaowei Wang, and Changwen Chen. Prompt-based learning for unpaired image captioning. *IEEE Transactions on Multimedia*, 2023.
- [136] Liu Zhuang, Lin Wayne, Shi Ya, and Zhao Jun. A robustly optimized bert pre-training approach with post-training. In *CNCCL*, 2021.
- [137] Luisa M Zintgraf, Taco S Cohen, Tameem Adel, and Max Welling. Visualizing deep neural network decisions: Prediction difference analysis. *arXiv preprint arXiv:1702.04595*, 2017.

SUMMARY OF THE APPENDIX

This supplementary contains additional details for the twelfth International Conference on Learning Representations submission, titled “Facing the Elephant in the Room: Prompt-Tuning or Full finetuning?”. The supplementary is organized as follows:

- §A shows **comprehensive training/testing curve** for VTAB-1k benchmark and cover comprehensive experiments on overfitting datasets.
- §B presents **visualization inspections** on a new-added explanation method — Integrated Gradients, and more results on GradCAM.
- §C shows the per-task results of full finetuning and visual prompt tuning on different **backbone and pretraining objectives**.
- §D presents the **image examples from the VTAB-1k image classification benchmark**.
- §E shows the optimal **prompt length** for each task in VTAB-1k image classification benchmark.
- §F further presents the **FID scores from a new-added image classification benchmark — FGVC [37], and compare the performance between visual prompt tuning and full fine-tuning**.
- §G is the **discussion** of legal/ethical considerations, reproducibility, social impact, limitations and future work.

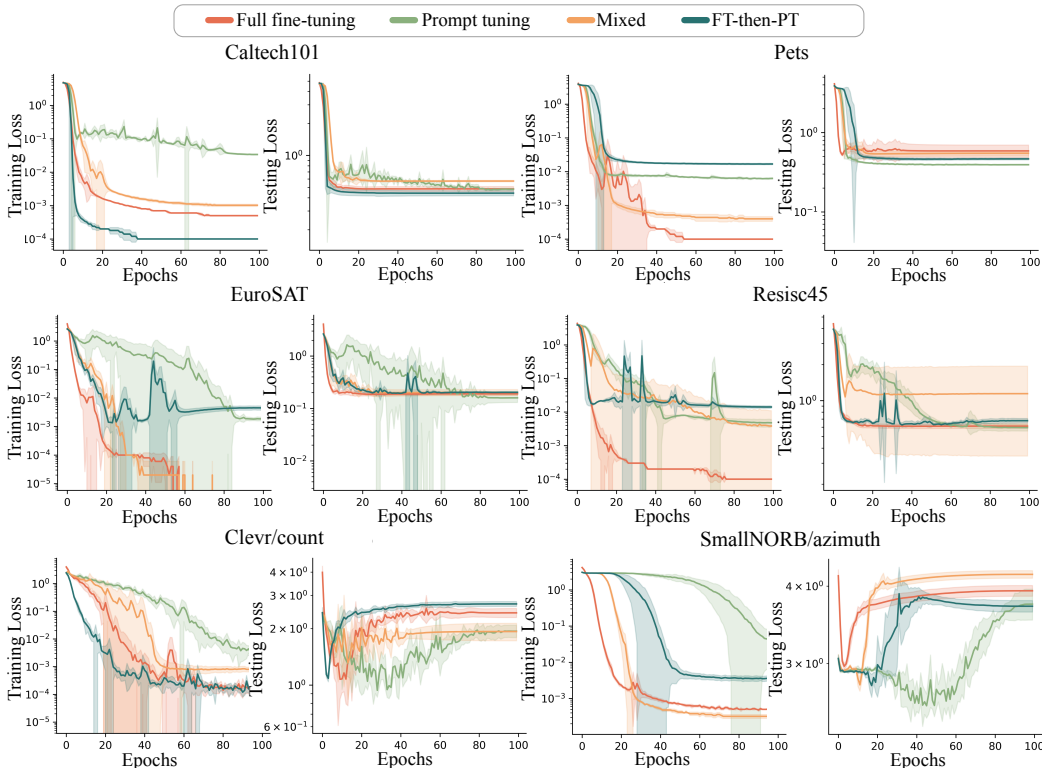


Figure 9: **Log-scale version of Figure 5.** Consistent to our paper, ■■■■ colors represent four training strategies: full finetuning, prompt tuning, mixed and FT-then-PT, respectively. Same for Figure 10, Figure 11 and 12.

A PER-TASK TRAINING/TESTING CURVE

In Figure 9, we further provide the log-scale version of Figure 5.

In Figure 10, 11 and 12, we comprehensively present per-task training/testing curve on VTAB-1k [127] *Natural*, *Specialized* and *Structured*, respectively. We do not observe overfitting for full

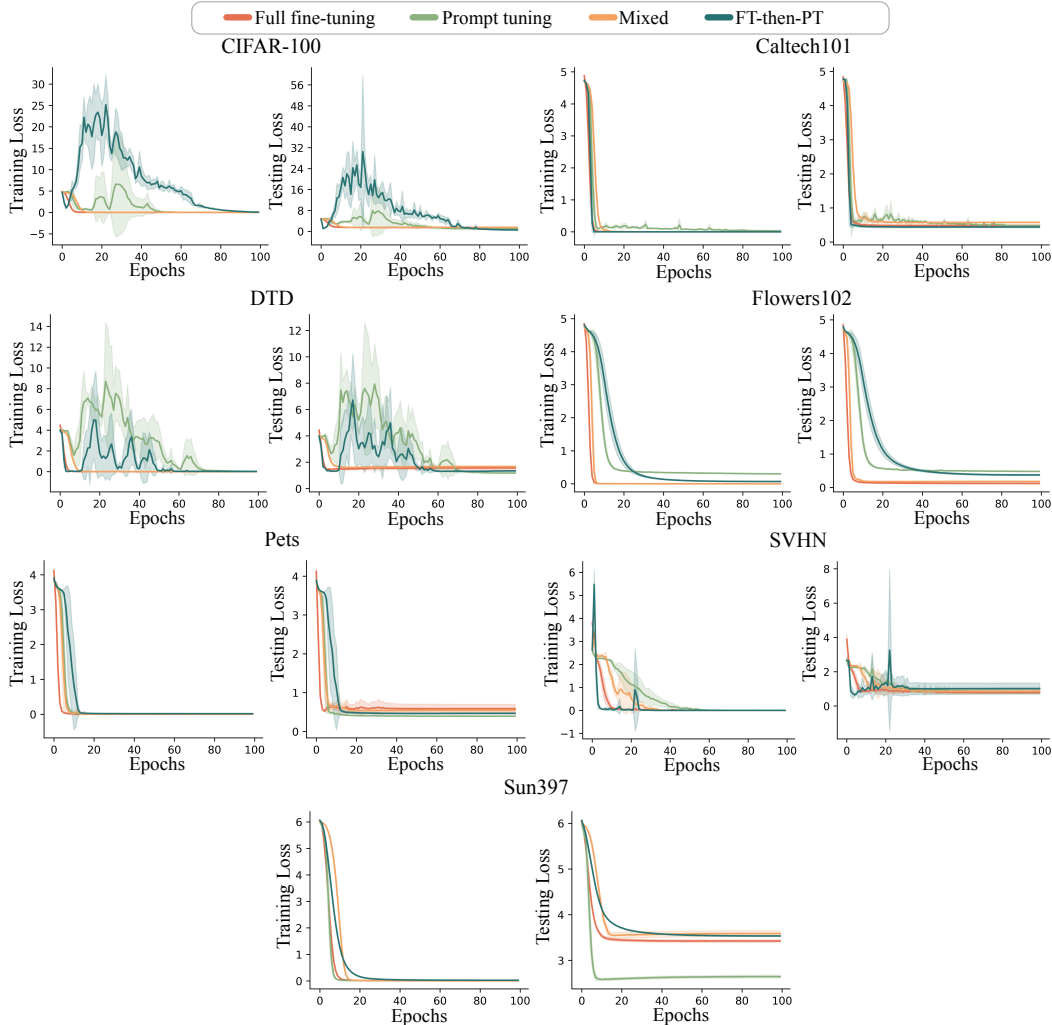


Figure 10: Per-task training/testing loss curve for VTAB-1k [127] *Natural*.

finetuning in 10 of 19 cases, while 9 of 19 datasets suffer from overfitting for both full finetuning and prompt tuning. Further experiments in Figure 13 on increasing training samples reduce overfitting for both methods. Overall, it can be inferred that overfitting is not the underlying cause of performance degradation observed in full finetuning when compared to prompt tuning.

In Table 5 and 6, we provide per-task results in accuracy with different dataset scales (See Figure 2, 4 and §5.1 in our paper. Not all datasets are provided in the same scale due to the limit number of data samples). We further provide comprehensive accuracy curves with different dataset scales in Figure 14, 15 and 16 for VTAB-1k *Natural*, *Specialized* and *Structured*, respectively, which covers more combinations that are not showed in Table 5 and 6. The figures and Tables are consistent with our observations and support our assumption in the main paper that the performance gap between full finetuning and prompt tuning decreases as the amount of training data increases, with full finetuning even outperforming prompt tuning under high-resource settings, demonstrating the robustness and generality of prompt tuning in the face of limited finetuning examples. Though full finetuning generally achieves higher accuracy when rich data examples are available, the prompt tuning still has a competitive performance with much fewer parameters.

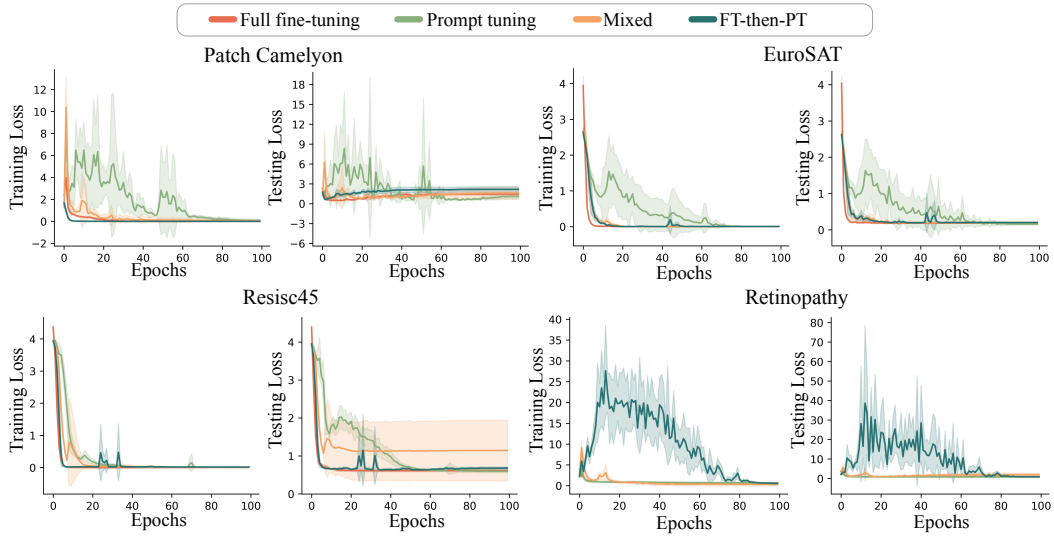


Figure 11: Per-task training/testing loss curve for VTAB-1k [127] *Specialized*.

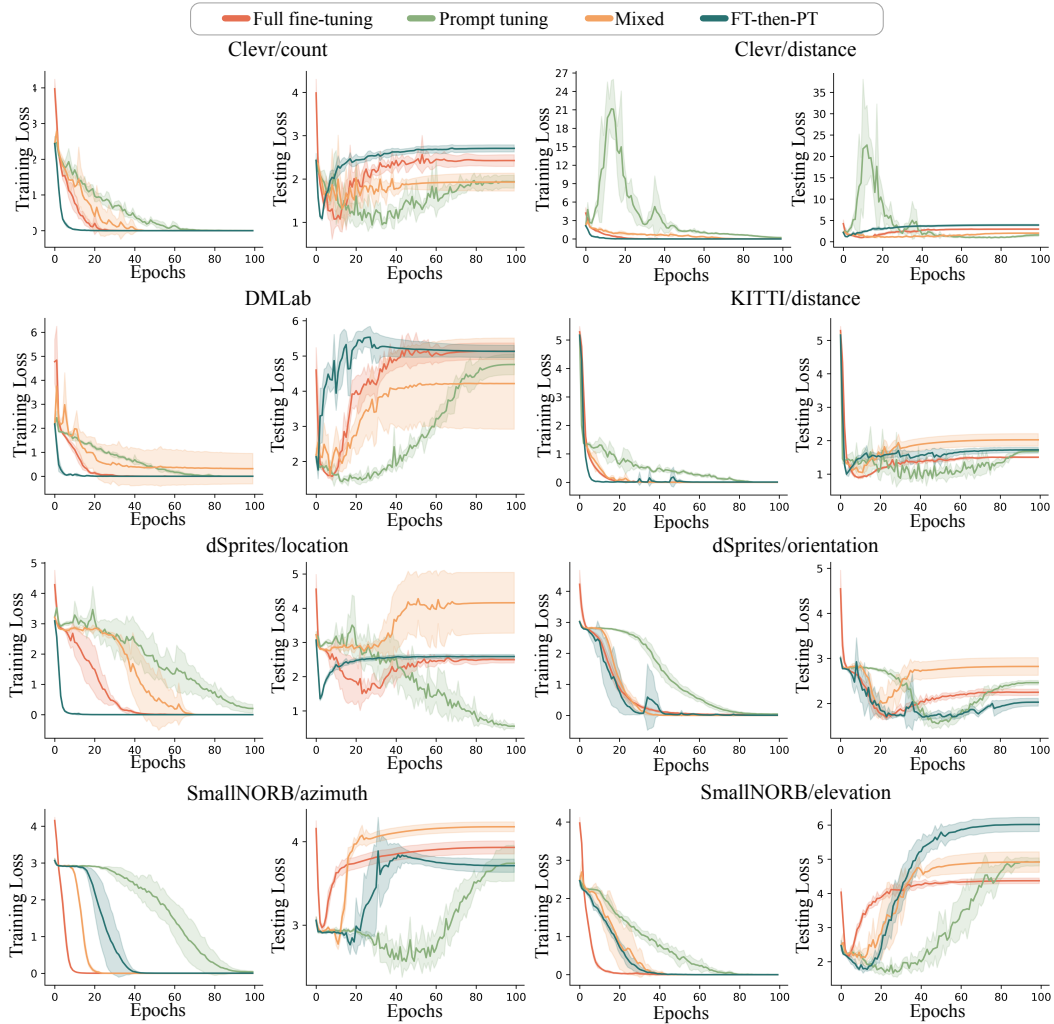


Figure 12: Per-task training/testing loss curve for VTAB-1k [127] *Structured*.

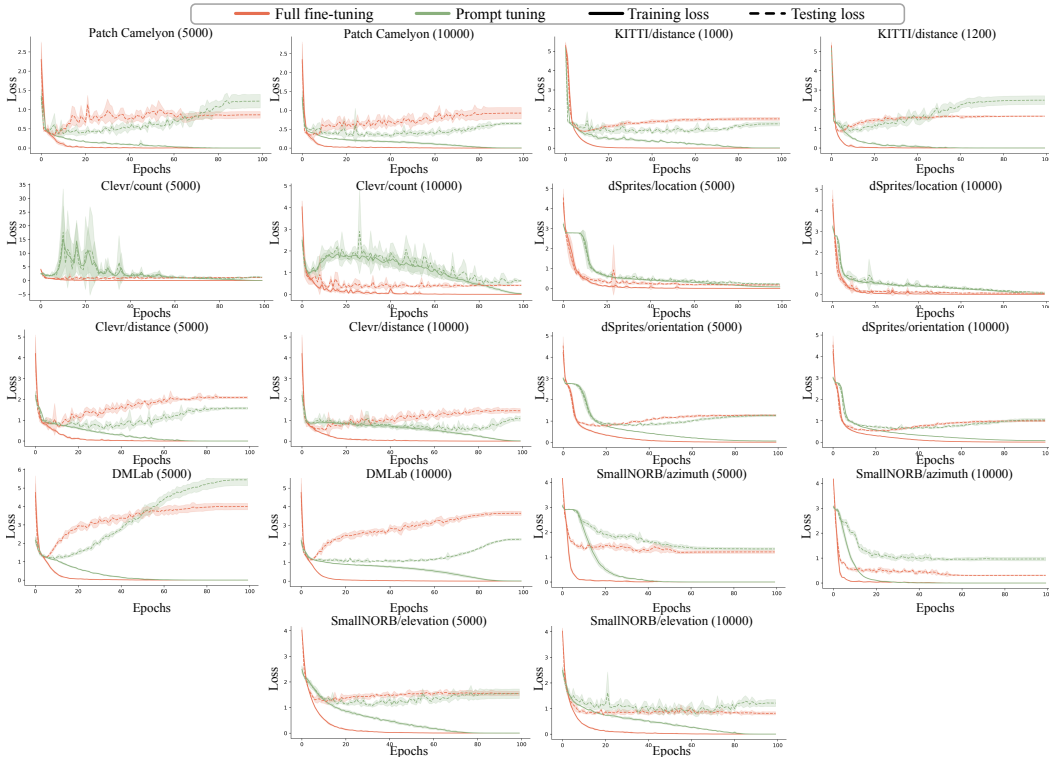


Figure 13: **Training/testing loss curve for VTAB-1k [127] when increasing training samples.** Note that the listed 9 cases are observed in overfitting on both full finetuning and prompt tuning under default numbers of training samples. We further increase the number of training samples and find that both methods are mitigated from overfitting. Further enlarging in training data results in unacceptable training time and fails to cover some cases in this plot (*i.e.*, not having enough training samples). (·) represents the number of training samples. For most cases, we apply 5000/10000 number of images for training. VTAB-1k [127] *Structured* KITTI/distance [32], on the other hand, does not have sufficient data for training, we present results for 1000/1200.

Table 5: **Image classification accuracy on different scales on VTAB-1k [127] *Natural* and *Specialized*** for ViT-B/16 [24] pretrained on supervised ImageNet-21k [89]. Note that not all combinations in Figure 4 are listed in this table. “# Samples” represent the number of training samples during finetuning. We pick 400, 800, 5000 and 10000 number of training data as the general number of samples across all datasets. All results are averaged on five runs. Same for Table 6.

ViT-B/16 [24] (85.8M)	# Samples	VTAB-1k [127] <i>Natural</i> [7]							Mean	VTAB-1k [127] <i>Specialized</i> (4)				Mean
		CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397		Patch Camelyon	EuroSAT	Resisc45	Retinopathy	
FT [47]	400	51.1	76.1	55.1	92.9	79.4	73.6	27.1	65.04	76.2	96.3	74.4	73.6	80.13
VPT [49]	400	70.6	84.9	60.6	97.5	86.9	76.0	39.2	73.67	62.5	96.7	76.1	74.0	77.33
FT [47]	800	64.5	88.1	65.0	97.3	84.6	87.3	39.5	75.19	81.2	95.7	83.2	73.3	83.48
VPT [49]	800	77.7	90.2	68.8	98.1	88.4	82.5	51.0	79.53	77.9	96.2(9)	83.3	73.1	82.65
FT [47]	5000	85.0	-	-	-	-	95.6	56.8	79.13	89.1	98.6	93.7	76.6	89.50
VPT [49]	5000	87.1	-	-	-	-	92.2	69.4	82.9	86.3	98.3	92.1	74.3	87.75
FT [47]	10000	88.6	-	-	-	-	96.4	66.8	83.93	89.0	98.9	95.3	77.6	90.20
VPT [49]	10000	88.7	-	-	-	-	94.0	72.6	85.10	87.7	98.6	93.9	74.6	88.70

B VISUALIZATION INSPECTIONS

We further introduce Integrated Gradients [100] (IG), and present visualization inspections in Figure 17(a) to support the advantages of visual prompts discussed in our paper. In general, IG tries to attribute the prediction of a deep network to its input features, and it is commonly applied [61; 2; 4] in the research of explainable Artificial Intelligence (AI). IG gains widespread adoption as an interpretability technique, owing to its versatility in explaining any differentiable model, including

Table 6: **Image classification accuracy on different scales on VTAB-1k [127] *Structured*** for ViT-B/16 [24] pretrained on supervised ImageNet-21k [89].

ViT-B/16 [24] (85.8M)	# Samples	VTAB-1k [127] <i>Structured</i> [8]								Mean
		Clevr/cnt.	Clevr/dist.	DMLab KITTI/dist.	dSprites/loc.	dSprites/ori.	SmallNORB/az.	SmallNORB/elev.		
FT [47]	400	45.3	54.3(1)	34.7	72.2	38.4	32.6	17.7	23.3	39.82
VPT [49]	400	60.8	54.3(1)	41.2	67.0	69.0	38.4	21.5	29.4	47.70
FT [47]	800	55.4	58.2(2)	40.4	74.7	54.0	47.0	26.2	28.6	48.07
VPT [49]	800	68.2	58.2(3)	45.3	77.7	78.4	48.4	31.2	41.3	56.09
FT [47]	5000	84.4	74.1	58.4	-	96.1	73.1	78.9	67.95	66.62
VPT [49]	5000	76.2	77.8	57.8	-	93.3	70.2	73.5	68.02	73.83
FT [47]	10000	93.4	80.7	63.9	-	98.8	79.4	94.0	83.4	84.80
VPT [49]	10000	84.4	81.6	63.4	-	96.6	76.7	87.2	80.3	81.46

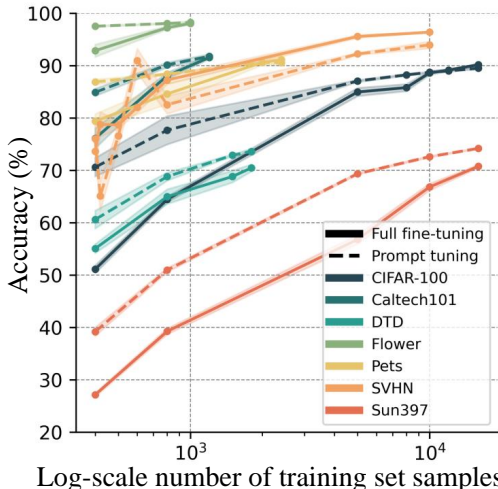


Figure 14: **Analysis of dataset capacity on VTAB-1k [127] *Natural***. The solid lines stand for full finetuning and the dotted lines represent prompt tuning. We take the log-scale of training data samples for better separation and each color stands for an individual classification task. Each point is given by average over five runs. In general, with the increasing of dataset in size, the performance gap between full finetuning and prompt tuning becomes closer. Same for Figure 15 and 16.

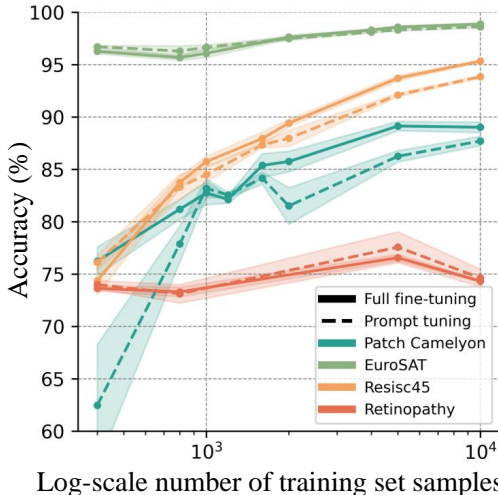


Figure 15: **Analysis of dataset capacity on VTAB-1k [127] *Specialized***.

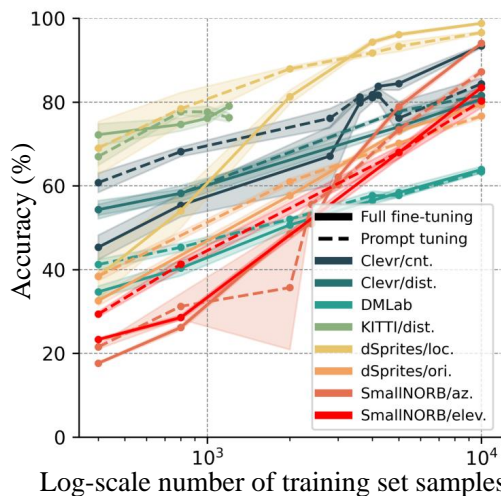
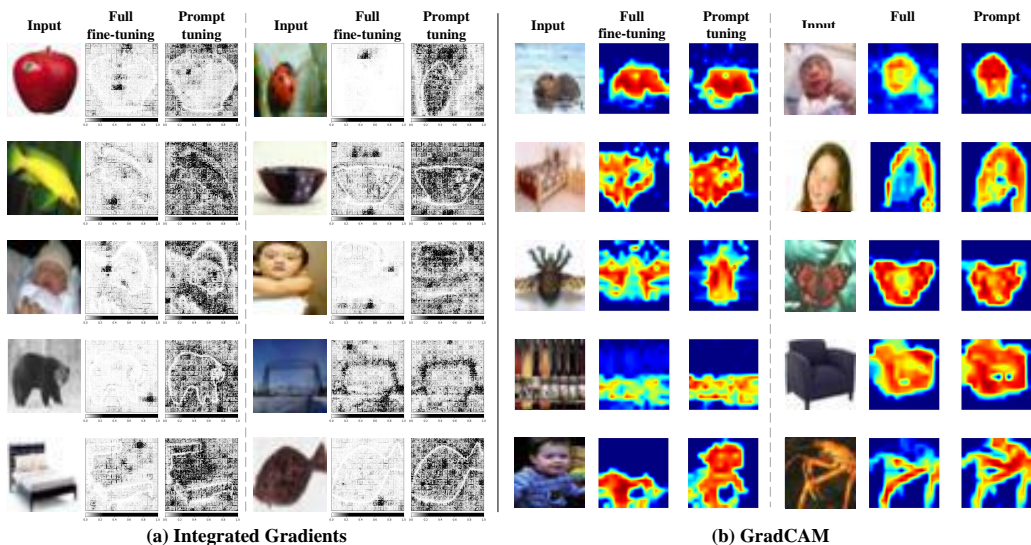
Figure 16: Analysis of dataset capacity on VTAB-1k [127] *Structured*.

Figure 17: (a) **Integrated Gradients (IG) [100] visual inspection of full finetuning and prompt tuning.** Note that the darker regions responds to high score for class (In light of the observation that certain images exhibit a consistently negative gradient, it becomes necessary to take the absolute value in order to ensure consistency and save successfully across all images within the testing set. This approach results in the emergence of darker boundaries within the resulting images). From left to right are input images after standard data augmentation, IG results for full finetuning and IG results for prompt tuning. (b) **More visual inspection of full finetuning and prompt tuning using GradCAM [92].** Consistent to our paper, the red regions correspond to high score for class. From left to right are input image after standard data augmentation, GradCAM results for full finetuning and GradCAM results for prompt tuning. Figure best viewed in color.

images [100; 4], text [18; 84], and structured data [72; 122; 1]. In Figure 17(a), we can observe straightforward visual explanation differences (*e.g.*, take the *Coccinella septempunctata* image as an example, when full finetuning fails to provide high gradients to make a correct decision, prompt tuning instead recognize it with significantly higher gradients), showing consistency with our paper.

In Figure 17(b), we present more visualization inspection results for full finetuning and prompt tuning using GradCAM [92]. Overall, we present additional visual evidence to support the notion that prompt tuning encompasses actual visual explanations throughout the training process.

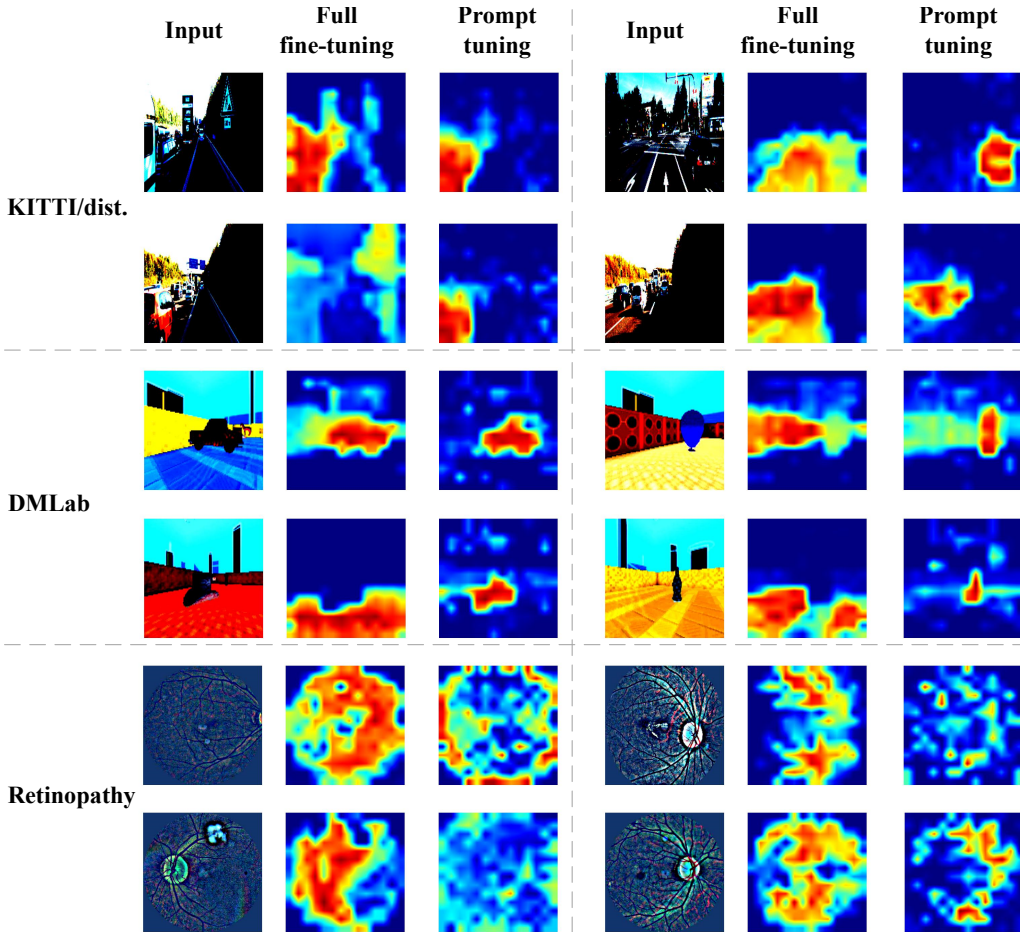


Figure 18: **Visual inspection of full finetuning and prompt tuning in other 3 groups.** For KITTI-dist. and DMLab, we present the cases where full finetuning gets inferior performance to visual prompt tuning and vice versa for Retinopathy. KITTI-dist. lies in the 2nd quadrant, DMLab lies in the 3rd quadrant and Retinopathy lies in the 4th quadrant with regard to Figure 1. Figure best viewed in color.

Table 7: **VTAB-1k [127] Natural per-task results for Swin-Base [66]** pretrained on supervised ImageNet-21k. Best results among full finetuning and prompt tuning are **bold**. We report the “number of wins” in [·] compared to full finetuning. Same for Table 8 to 15.

Swin-Base [66] (86.7M)	VTAB-1k [127] Natural (7)							Mean
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397	
FT [47]	72.2	88.0	71.2	98.3	89.5	89.4	45.0	79.10
VPT [49]	79.6	90.8	78.0	99.5	91.4(3)	86.4	51.7	78.78 (6)
- Tuned / Total (%)	0.13	0.13	0.07	0.13	0.06	0.70	0.48	0.28

We further present visualization inspection results based on Figure 1. In Figure 8 and 17, we primarily focus on images with similar data distribution and task, we thus further posit visualization inspection results under the setting of other three groups in Figure 18. As seen, for KITTI-dist. and DMLab from 2nd and 3rd quadrant respectively, visual prompt tuning presents clear dense visual evidence and focuses on correct positions. On the other hand, for Retinopathy lies in the 4th quadrant, full finetuning presents more stable results, focusing on cells’ characteristics.

Table 8: **VTAB-1k [127] *Specialized* per-task results for Swin-Base [66]** pretrained on supervised ImageNet-21k.

Swin-Base [66] (86.7M)	VTAB-1k [127] <i>Specialized</i> [4]				Mean
	Patch Camelyon	EuroSAT	Resisc45	Retinopathy	
FT [47]	86.6	96.9	87.7	73.6	86.21
VPT [49]	80.1	96.2	85.0	72.0	83.33 (0)
- Tuned / Total (%)	0.07	0.13	0.19	0.02	0.10

Table 9: **VTAB-1k [127] *Structured* per-task results for Swin-Base [66]** pretrained on supervised ImageNet-21k.

Swin-Base [66] (86.7M)	VTAB-1k [127] <i>Structured</i> [8]								Mean
	Clevr/cnt.	Clevr/dist.	DMLab	KITTI/dist.	dSprites/loc.	dSprites/ori.	SmallNORB/az.	SmallNORB/elev.	
FT [47]	75.7	59.8	54.6	78.6	79.4	53.6	34.6	40.9	59.65
VPT [49]	67.6	59.4	50.1	61.3	74.4	50.6	25.7	25.7	51.85 (0)
- Tuned / Total (%)	0.70	0.70	0.14	0.69	0.15	0.09	0.16	0.02	0.38

C PER-TASK RESULTS ON DIFFERENT PRETRAINING OBJECTIVES.

We also report the per-task results on VTAB-1k [127] Swin-Base [66] (*i.e.*, Table 7, 8 and 9), MAE [41] (*i.e.*, Table 10, 11 and 12) and MoCo v3 [16] (*i.e.*, Table 13, 14 and 15), respectively. Overall, the empirical results consistently reveal the attainment of superior or competitive performance gains through prompt tuning, in comparison to full finetuning, across diverse tasks within the default train-val split. Notably, these performance gains are achieved while maintaining a substantially reduced number of model parameters.

D IMAGE EXAMPLES

In Figure 19, we include image examples from VTAB-1k [127] image classification benchmark, including 19 visual tasks categorized into three groups: *Natural*, *Specialized*, and *Structured*.

E PROMPT LENGTH

In Table 16, 17 and 18, we posit the per-task prompt length, which is consistent with the original VPT [37] approach.

F FID ON FGVC

In Table 19, we further present the corresponding FID scores on FGVC [37] image classification benchmark. Specifically, it contains 5 Fine-Grained Visual Classification, including CUB-200-2011 [106], NABirds [103], Oxford Flowers [74], Stanford Dogs [51] and Stanford Cars [31]. As seen, a higher FID score might potentially lead to relatively higher accuracy on full fine-tuning, and vice versa. This is consistent with the hypothesis claimed in §4.3.

G DISCUSSION

G.1 ASSET LICENSE AND CONSENT

The majority of Visual Prompt Tuning (VPT) [49] is licensed under CC-BY-NC 4.0. Portions of [49] are available under separate license terms: google-research/task_adaptation and huggingface/transformers are licensed under Apache-2.0; Swin-Transformer [66] and ViT-pytorch [24] are licensed under MIT; and MoCo-v3 [16] and MAE [41] are licensed under CC BY 4.0. Fréchet Inception Distance (FID) [43; 91] is licensed under Apache License 2.0.

Table 10: **VTAB-1k [127] *Natural* per-task results for ViT-B/16 [24] pretrained on MAE [41].** Though prompt tuning shows inferior performance to full finetuning, its parameter-efficient nature makes it applicable in parameter-sensitive scenes. Same for MOCO [16] applications.

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Natural</i> [7]							Mean
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397	
FT [47]	24.6	84.2	56.9	72.7	74.4	86.6	15.8	59.31
VPT [49]	8.2	55.2	58.0	39.3	45.2	19.4	21.9	35.31(2)
- Tuned / Total (%)	0.20	0.20	0.15	0.10	0.04	0.54	0.41	0.23

Table 11: **VTAB-1k [127] *Specialized* per-task results for ViT-B/16 [24] pretrained on MAE [41].**

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Specialized</i> [4]				Mean
	Patch Camelyon	EuroSAT	Resisc45	Retinopathy	
FT [47]	81.8	94.0	72.3	70.6	79.68
VPT [49]	77.9	94.9	45.4	73.6	72.95(2)
- Tuned / Total (%)	1.06	1.07	0.15	0.02	0.57

G.2 REPRODUCIBILITY

This paper is implemented in Pytorch [80]. Experiments are conducted on NVIDIA A100-80GB GPUs. For reproducibility, our full implementation shall be publicly released upon paper acceptance.

G.3 MORE DISCUSSION ON DATASET CAPACITY

Previous works [107] demonstrate that prompt tuning in language also shows same trend that prompt tuning achieves better performance when lacking task-specific data.

Similarly in vision perspective, VPT [37] partially explore this question by claiming a consistently advanced performance to full fine-tuning across training data sizes, which is proved to be mistaken in our paper (see §4.4) when the dataset scale continue to expand. Our research undertakes a subsequent update of the claims made in the preceding research, rather than merely adopting them in their original form. Also, the results in [37] are shown on a different dataset benchmark (*i.e.*, FGVC [37]). Our paper contributes an additional piece to the puzzle of addressing gaps in common image recognition benchmarks.

G.4 CHOOSING SINGLE TRAINING SAMPLE FOR ONE-SHOT LEARNING

We select one image per class for `train` and `val`, respectively. The choosing of the representative image for each class under the one-shot learning scene is important (see §4.4). We want to ensure that the chosen image is not an outlier [117]. Instead, it should share common features with other images from the same class. We thus apply a method called iterative testing. Specifically, we experiment with different images from the same class and evaluate our model. If the model’s performance falls in a reasonable range, we state that the chosen images are representative. Also, data augmentation methods are applied, which helps to fix the variations of the image.

G.5 SOCIAL IMPACT

This work systematically investigates several hypotheses to demystify the mechanisms behind visual prompt tuning’s success. We carefully study whether its success is provided by its enhanced resilience against overfitting, flexibility in task transfer, efficient learning on small datasets, or improved optimization due to additional dimensions. Through extensive experiments on 19 diverse datasets and tasks, we reveal that surprisingly, overfitting is not the root cause of full tuning’s inferior performance. Prompt tuning demonstrates better generalization when there is limited data, while full tuning catches up or even surpasses prompt tuning when more data is presented. However, prompt tuning is still a competitive approach considering its parameter-efficient nature. Overall, our paper suggests that researchers should have a clear view of the task disparities and dataset scale of down-

Table 12: VTAB-1k [127] *Structured* per-task results for ViT-B/16 [24] pretrained on MAE [41].

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Structured</i> [8]								Mean
	Clevr/cnt.	Clevr/dist.	DMLab	KITTI/dist.	dSprites/loc.	dSprites/ori.	SmallNORB/az.	SmallNORB/elev.	
FT [47]	67.0	59.8	45.2	75.3	72.5	47.5	30.2	33.0	53.82
VPT [49]	39.0	40.9	30.6	53.9	21.0	12.1	11.0	14.88	27.91(0)
- Tuned / Total (%)	0.54	2.11	1.07	0.54	0.12	0.55	2.12	2.11	1.14

Table 13: VTAB-1k [127] *Natural* per-task results for ViT-B/16 [24] pretrained on MOCO [16].

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Natural</i> [7]							Mean
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397	
FT [47]	57.6	91.0	64.6	91.6	79.9	89.8	29.1	71.95
VPT [49]	70.1	88.3	65.9	88.4	85.6	57.8	35.7	70.26(4)
- Tuned / Total (%)	0.20	0.20	0.15	0.10	0.04	0.54	0.41	0.23

stream finetuning, and carefully select appropriate way for finetuning under “pretrain-then-finetune” paradigm.

G.6 LIMITATIONS AND FUTURE WORK

Although we investigate several interesting and compelling doubts between full finetuning and visual prompt tuning, which are of great significance to the research community, it also comes with new challenges and unveils some intriguing questions. For example, the examination of visual prompt tuning reveals similarities to prompt tuning approaches in NLP [13; 107] and vision from different perspectives (*e.g.*, [114] discusses that prompt tuning is highly robust to label noises; [76] demonstrates how prompt-tuning enables the model to attend to context-relevant information), thereby suggesting the need for further investigation to facilitate a comprehensive and unified study. This topic can be extended to parameter-efficient methods other than prompt tuning techniques. For example, we build limited experiments on LoRa [46], showing that there is similar trend on dataset capacity when comparing with full fine-tuning (see Table 20). Another essential future direction deserving of further investigation is the the design and analysis of network’s attention position and *ad-hoc* interpretability. In our paper, we follow common practice and propose network interpretability through various visualization explanation methods (*i.e.*, GradCAM, IG) at the final layer of the network. However, we highlight that visualizing the attention from cls token to other tokens [9] can be an alternative approach to understand the effect of VPT on feature. We share include different visualization position in future research. We further highlight [94] a very important work in discussing the relation between attention visual evidence and performance. During our current research, the visual evidence can hardly provide an intuitive/straightforward explanation for the performance gain, while some current works [94] pave a path for such the connection. Also, the discussed visualization approaches can be generally categorized into the *post-hoc* explainability (*i.e.*, producing explanations for trained networks by importance values [97; 126; 132; 27; 70; 95] or sensitivities of inputs [67; 52; 137]). Such approaches, however, might suffer from possible misleading [87; 88; 57; 2]. Therefore, it is worth further investigating the *ad-hoc* explainability of full finetuning and prompt tuning [81; 109] (*i.e.*, case/concept-based reasoning), particularly in decision-critical and safety-sensitive scenarios [78; 25]. The investigation of prompt tuning’s *ad-hoc* explainability in research is still relatively scarce and necessitates further exploration.

Table 14: **VTAB-1k [127] *Specialized* per-task results for ViT-B/16 [24] pretrained on MOCO [16].**

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Specialized</i> [4]				Mean
	Patch Camelyon	EuroSAT	Resisc45	Retinopathy	
FT [47]	85.1	96.4	83.1	74.3	84.72
VPT [49]	83.1	91.0	81.2	74.0	82.33(0)
- Tuned / Total (%)	1.06	1.07	0.15	0.02	0.57

Table 15: **VTAB-1k [127] *Structured* per-task results for ViT-B/16 [24] pretrained on MOCO [16].**

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Structured</i> [8]								Mean
	Clevr/cnt.	Clevr/dist.	DMLab	KITTI/dist.	dSprites/loc.	dSprites/ori.	SmallNORB/az.	SmallNORB/elev.	
FT [47]	55.2	56.9	44.6	77.9	63.8	49.0	31.5	36.9	51.98
VPT [49]	48.5	55.8	37.2	64.6	52.3	26.5	19.4	34.8	42.39(0)
- Tuned / Total (%)	0.54	2.11	1.07	0.54	0.12	0.55	2.12	2.11	1.14

Table 16: **VTAB-1k [127] *Natural* per-task prompt length for ViT-Base [24] pretrained on supervised ImageNet-21k.**

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Natural</i> (7)							Mean
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397	
VPT [49]	100	5	1	200	50	200	1	79.58

Table 17: **VTAB-1k [127] *Specialized* per-task prompt length for ViT-Base [24] pretrained on supervised ImageNet-21k.**

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Specialized</i> [4]				Mean
	Patch Camelyon	EuroSAT	Resisc45	Retinopathy	
VPT [49]	5	50	50	10	28.75

Table 18: **VTAB-1k [127] *Structured* per-task prompt length for ViT-Base [66] pretrained on supervised ImageNet-21k.**

ViT-B/16 [24] (85.8M)	VTAB-1k [127] <i>Structured</i> [8]								Mean
	Clevr/cnt.	Clevr/dist.	DMLab	KITTI/dist.	dSprites/loc.	dSprites/ori.	SmallNORB/az.	SmallNORB/elev.	
VPT [49]	100	200	100	100	100	100	200	200	137.5

Table 19: **FGVC [49] per-task results for ViT-Base/16 [24] pretrained on supervised ImageNet-21k.**

ViT-Base/16 [24] (85.8M)	FGVC [49] [5]					Mean
	CUB-200-2011	NAbirds	Oxford Flowers	Stanford Dogs	Stanford Cars	
FULL [47]	87.3	82.7	98.8	89.4	84.5	88.54
VPT [49]	88.5	84.2	99.0	90.2	83.6	89.11 (4)
- Tuned / Total (%)	0.29	1.02	0.14	1.17	2.27	0.98
FID score	146.845	169.117	167.690	143.840	182.487	-

Table 20: **Cifar-100 [54] for ViT-Base [24] (bottom-up) pretrained on supervised ImageNet-21k under the settings of [94].**

# training sample	400	800	10000
Full fine-tuning	44.5%	70.2%	87.9%
LoRa [46]	69.6%	83.6%	90.7%
Performance gap	25.1%	13.4%	2.8%

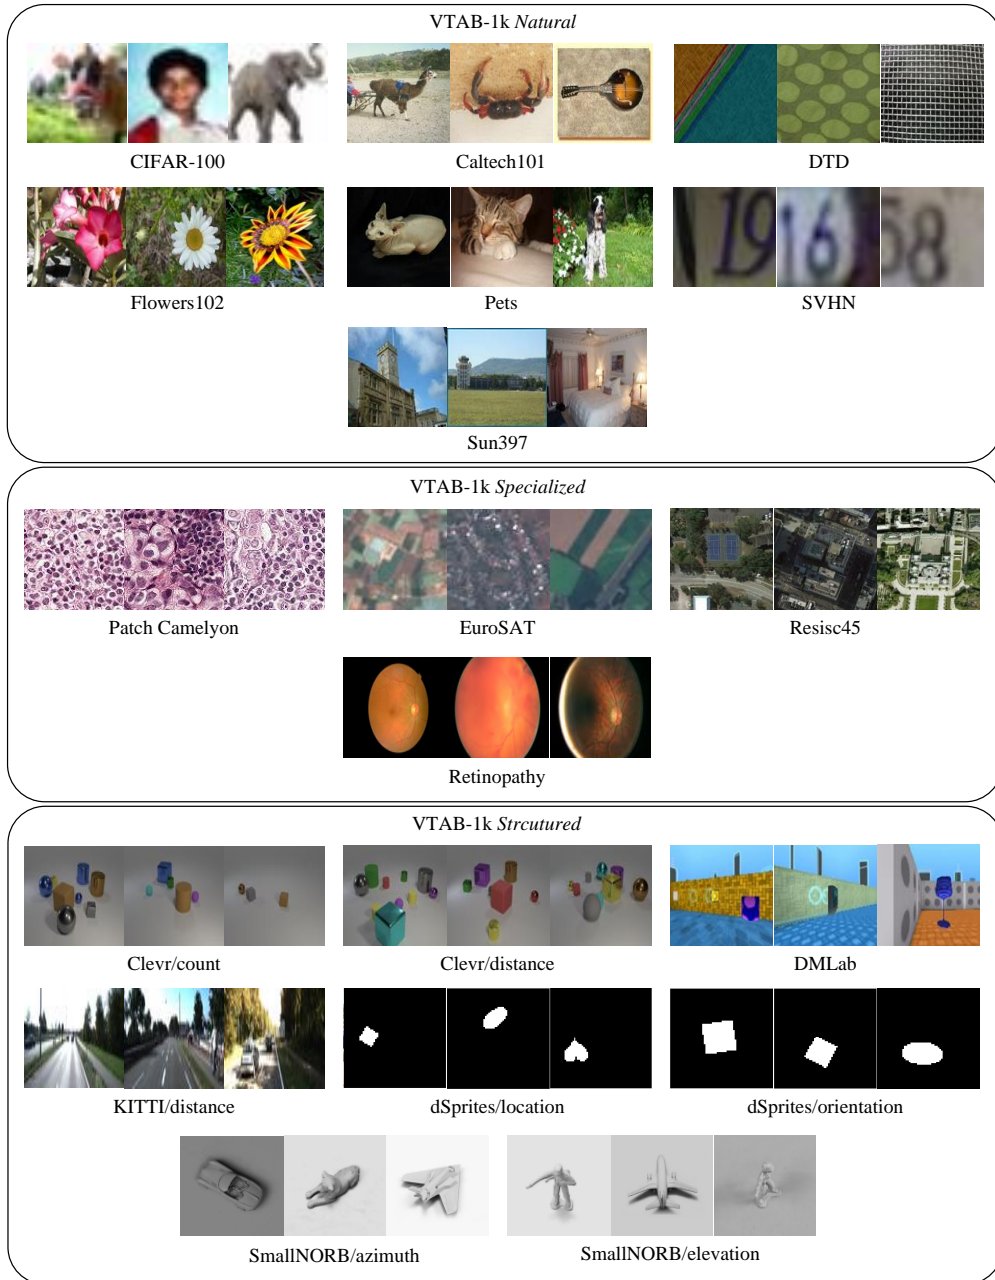


Figure 19: Dataset examples from VTAB-1k [127] image classification benchmark.