This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Faster Parameter-Efficient Tuning with Token Redundancy Reduction

Kwonyoung Kim¹ Jungin Park^{1*} Jin Kim¹ Hyeongjun Kwon¹ Kwanghoon Sohn^{1,2*}

¹Yonsei University, ²Korea Institute of Science and Technology (KIST)

{kyk12, newrun, kimjin928, kwonjunn01 khsohn}@yonsei.ac.kr

Abstract

Parameter-efficient tuning (PET) aims to transfer pretrained foundation models to downstream tasks by learning a small number of parameters. Compared to traditional fine-tuning, which updates the entire model, PET significantly reduces storage and transfer costs for each task regardless of exponentially increasing pre-trained model capacity. However, most PET methods inherit the inference latency of their large backbone models and often introduce additional computational overhead due to additional modules (e.g. adapters), limiting their practicality for compute-intensive applications. In this paper, we propose Faster Parameter-Efficient Tuning (FPET), a novel approach that enhances inference speed and training efficiency while maintaining high storage efficiency. Specifically, we introduce a plug-and-play token redundancy reduction module delicately designed for PET. This module refines tokens from the self-attention layer using an adapter to learn the accurate similarity between tokens and cuts off the tokens through a fully-differentiable token merging strategy, which uses a straight-through estimator for optimal token reduction. Experimental results prove that our FPET achieves faster inference and higher memory efficiency than the pre-trained backbone while keeping competitive performance on par with state-of-the-art PET methods. The code is available at https://github.com/ kyk120/fpet.

1. Introduction

Pre-training on large-scale web-collected data followed by fine-tuning on specific downstream tasks is a foundational paradigm that leads to state-of-the-art performance across various vision-related tasks. Nonetheless, modifying all parameters for each distinct task is inefficient in terms of storage, as it requires updating and storing the whole parameters for every individual task. To address this issue, parameter-efficient tuning (PET) approaches [5, 7–9, 12–



Figure 1. Average accuracy vs. inference time and GPU memory usage on VTAB-1K [38]. Our FPET significantly surpasses all existing PET methods in terms of inference speed and computation cost. Note that the dotted line in (a) represents the *no inference latency* established in prior studies.

17, 21, 22, 29–31, 34, 37, 39, 40, 42] have proposed to utilize a minimal number of parameters to transfer pretrained models to downstream tasks. They have demonstrated significant storage efficiency while attaining comparable or even surpassing performance to full-tuning. For example, [16] outperforms full-tuning performance on VTAB-1K [38] by 8% with only less than 2% trainable parameters.

PET methods have continually evolved with the fundamental objective of conserving resources by redundancy reduction [16]. Specifically, they have accomplished redundancy reduction through rank decomposition with additional lightweight learnable modules (e.g. adapter) [5, 8, 9, 12, 13, 15–17, 29, 31, 40, 42], task-specific parameter reduction with additional learnable tokens [14, 21, 34, 41], or precision redundancy reduction with adapter quantiza-

^{*} Corresponding authors.

tion [16]. While these strategies have made significant advances in storage efficiency, they are still insufficient to apply to real-world applications (*e.g.* web platforms utilized by millions of users or edge devices that are constrained by resources) due to two main challenges: *inference latency* and *computing memory usage*. Most existing works often require higher inference latency and computation than fulltuning models during inference due to additional modules or parameters. Some works [13, 15, 22, 28, 37] have facilitated *no inference latency* compared to the backbone models, however, they inherit the original inference latency and computation requirement from the backbone.

To address the above limitations, some approaches [7, 39] have proposed to truncate the network architecture [7] or disentangle task-specific and pre-trained task-agnostic knowledge [39]. However, they still suffer from an increased number of learnable parameters and a slower inference speed [39], and an additional pre-training step [7]. Furthermore, they often exhibit degraded performance compared to state-of-the-art PET methods.

In this paper, we explore a straightforward approach to achieve both a feasible inference latency and a training computational efficiency. Motivated by the recent token reduction approaches [2, 3], which reduces model's input space, we present Faster Parameter-Efficient Tuning (FPET) that formulates PET with token redundancy reduction. Instead of directly utilizing the previous methods, which provides a sub-optimal solution due to their non-differentiable nature, we introduce a fully differentiable token redundancy reduction module. Specifically, we incorporate a straightthrough estimator (STE) [1] into token reduction to make the token selection process fully differentiable. In addition, our FPET performs token reduction once in the backbone model's intermediate layer contrary to [2, 3], which performs token merging based on token similarity in the early layer where the impact of the adapter is not fully manifest. Consequently, our FPET achieves faster inference speed, surpassing even the no inference latency approaches [13, 15, 22, 28, 37], while maintaining efficiency gains during training. Experiments on VTAB-1K [38] and extensive ablation study demonstrate the effectiveness and high efficiency of our FPET, achieving a faster inference speed and a lower computation requirement than existing PET methods while attaining comparable performance to state-of-the-art.

2. Related Work

2.1. Parameter-efficient tuning

Parameter-efficient tuning (PET) seeks to leverage large pre-trained models by tuning and storing only the minimal task-specific parameters, thereby minimizing storage requirements across a multitude of downstream tasks. Prompt-based methods [14, 19, 21, 26, 33, 34, 41, 43] introduce a small number of learnable tokens that are concatenated with input tokens of transformer. Through the selfattention layer, prompts modulate input tokens for adaptation, but the additional token count (1 to 200) results in a quadratic increase in computational complexity.

Adapter-based methods [5, 8, 9, 12, 13, 15–17, 25, 29– 31, 40], has made significant strides in decreasing rank redundancy employing rank decomposition matrices into transformer architecture. Especially, [9, 15, 17, 25] further improved the decomposition utilizing Kronecker products or butterfly factorization. Furthermore, [16] employs quantization on adapters, considerably diminishing the storage demands for each task by reducing numerical precision redundancy. However, their concern is limited to storage memory efficiency.

There have been attempts to maintain inference speed without incurring any increase from adapters [13, 15, 22, 37]. By employing low-rank adapters devoid of non-linear functions [13, 15], or utilizing a scaling and shifting module as an adapter [22], they enable the integration of learned parameters into the pre-trained model in advance. Furthermore, [37] maintains inference speed by implementing partial fine-tuning instead of additional adapters. Despite these attempts, the improvements have been limited in maintaining the original inference speed of pre-trained models.

Recent studies [7, 39] have proposed various approaches to enhance the efficiency of PET. These approaches include disentangled learning [39] and structural reduction methods [7]. While these strategies have shown potential for improving efficiency, the gains are often conditional and come with trade-offs. Specifically, they may enhance training or inference efficiency but at the expense of the need for an additional pre-training step [7], a significant increase in the number of learnable parameters [39], slower inference speed [39] or degraded accuracy [7, 39]. These limitations can affect their applicability in practical scenarios. In contrast, our approach is designed to enhance training, inference, and parameter efficiency simultaneously, offering a more practical and balanced solution for real-world PET implementations.

2.2. Token reduction

Vision Transformer (ViT) [6] has significantly advanced various vision tasks but faces challenges in resourceconstrained environments due to its high computational demands. Token pruning methods methods [18, 23, 24, 32, 35, 36] have been developed to accelerate ViT by selectively removing less critical tokens; however, these approaches can result in permanent information loss. To address this limitation, token merging [2] has been introduced as a technique that merges similar tokens, demonstrating improved performance over pruning methods. Unlike prior solutions such



Figure 2. Comparison between the bipartite soft matching [2] and the proposed bipartite differentiable matching. Our proposal is fully differentiable and refines the similarity between tokens, making token merging more optimal.

as *k*-means clustering [27] or graph cuts [4], token merging is a fast, non-iterative approach. However, it cannot be optimized as it is a parameter-free heuristic method, and its bipartite soft matching process is non-differentiable due to the use of *max* and top-*k* operations. Besides, to avoid sudden changes that could negatively affect the network, this method gradually merges tokens across all layers leading to less efficiency gain in GPU memory usage.

Recent advancements using learned thresholds [3] to merge tokens whose similarities exceed those learned thresholds. Despite learning thresholds, they did not consider refining the similarities between tokens. Additionally, their matching process remains non-differentiable, and the dynamic reduction strategy with thresholds does not guarantee efficiency gains unless the batch size is 1.

As these methods were designed for in-domain trainingfree or full-tuning scenarios, we have carefully adapted token merging for PET, thereby achieving significant efficiency improvements.

3. Preliminaries

3.1. Parameter-efficient tuning

In this section, we first lay out the framework for parameterefficient tuning. We start by concisely formulating the adapters that are commonly utilized.

LoRA [13] utilizes low-rank decomposition to approxi-

mate the variation in pre-trained parameters in response to changes in the input. Specifically, it learns the alterations in \mathbf{W}^q and \mathbf{W}^v within the self-attention layer to respectively modify the query and value projection weights. The adapter consist of two FC layers, $\mathbf{W}_{down} \in \mathbb{R}^{d \times h}$ and $\mathbf{W}_{up} \in \mathbb{R}^{h \times d}$ where, d is token dimension and $h \ll d$. Given an input $X \in \mathbb{R}^{N \times d}$, query and value formulations are

$$\mathbf{Q} = \mathbf{X}\mathbf{W}^{q} + s \cdot \mathbf{X}\mathbf{W}^{q}_{\text{down}}\mathbf{W}^{q}_{\text{up}},$$

$$\mathbf{V} = \mathbf{X}\mathbf{W}^{v} + s \cdot \mathbf{X}\mathbf{W}^{v}_{\text{down}}\mathbf{W}^{v}_{\text{up}},$$
(1)

where \mathbf{W}^q and \mathbf{W}^v are frozen and s is hyper-parameter.

AdaptFormer [5] employs low-rank decomposition as LoRA [13] but it incorporates non-linear layer such as ReLU activation between \mathbf{W}_{down} and \mathbf{W}_{up} . Additionally, AdaptFormer positions the adapter in parallel with the FFN layer in transformer layer. Given an input $X \in \mathbb{R}^{N \times d}$, the formulation is

$$\mathbf{X} = \mathbf{X} + FFN(\mathbf{X}) + s \cdot ReLU(\mathbf{X}\mathbf{W}_{down})\mathbf{W}_{up}, \quad (2)$$

where, the FFN module is frozen and s is a hyper-parameter.

3.2. Bipartite soft matching

ToMe [2] introduced bipartite soft matching as a noniterative, gradual method for merging tokens. The overall framework of bipartite soft matching is illustrated in Fig. 2a. Given tokens $\mathbf{X} = {\mathbf{x}_i} \in \mathbb{R}^{N \times d}$, where the total number of tokens is N and d is the feature dimension, tokens are assigned to two sets

$$\mathbf{X}_{\mathbf{A}} = \{ \mathbf{x}_i \mid i \mod 2 = 0 \}$$

$$\mathbf{X}_{\mathbf{B}} = \{ \mathbf{x}_i \mid i \mod 2 = 1 \}$$
(3)

To determine the optimal similarity and reduce complexity, the keys (**K**) from the self-attention layer are utilized instead of token features. Keys corresponding to each token from self-attention layer, $\mathbf{K} = {\mathbf{k}_i} \in \mathbb{R}^{N \times d'}$ are assigned to two sets

$$\mathbf{K}_{\mathbf{A}} = \{\mathbf{k}_i \mid i \mod 2 = 0\}$$

$$\mathbf{K}_{\mathbf{B}} = \{\mathbf{k}_i \mid i \mod 2 = 1\}$$
(4)

where d' is the key dimension and k_i corresponds to the x_i . The similarity matrix $\mathbf{C} \in \mathbb{R}^{N_A \times N_B}$ between two sets, $\mathbf{K}_A \in \mathbb{R}^{N_A \times d'}$ and $\mathbf{K}_B \in \mathbb{R}^{N_B \times d'}$, where N_A and N_B are the number of tokens in \mathbf{X}_A and \mathbf{X}_B respectively, is calculated as

$$\mathbf{C} = \mathbf{K}_{\mathrm{A}} \times \mathbf{K}_{\mathrm{B}}^{\mathrm{T}}$$
(5)

For each token in set \mathbf{K}_{A} , to find the most similar connection within \mathbf{K}_{B} , max operation along the last dimension of \mathbf{C} is applied to generate max values $\mathbf{C}_{max} \in \mathbb{R}^{N_{A}}$ and indices $\mathbf{I}_{max} \in \mathbb{R}^{N_{A}}$.

$$\mathbf{C}_{max}, \, \mathbf{I}_{max} = max(\mathbf{C}, \, dim = -1) \tag{6}$$

 I_{max} contains for each token in A, the most similar counterpart in B and C_{max} contains the corresponding maximum similarity. To match most similar k pairs among those, the top-k indices are extracted.

$$\mathbf{I}_{topk} = \operatorname{top-}k(argsort(\mathbf{C}_{max})) \tag{7}$$

Then, I_{max} and I_{topk} specify which tokens in A should be merged with which token in B. The merged tokens X_{merged} can be expressed as

$$\mathbf{X}_{\text{merged}} = merge(\mathbf{X}_{\text{A}}, \mathbf{X}_{\text{B}}, \mathbf{I}_{max}, \mathbf{I}_{topk})$$
(8)

where the *merge* is a function that merges tokens by averaging matched tokens according to the input indices. Finally, the merged tokens are concatenated with the remaining unmerged tokens and returned as the output of the token merging module. Note that due to the non-differentiable nature of their *max* and top-*k* operations, the matching process cannot be explicitly incorporated into the training objective.

4. Faster parameter-efficient tuning

We introduce token redundancy reduction into PET to surpass the original speed of backbone models and lower computational costs without compromising on accuracy. To achieve this goal, we have investigated optimal configuration of token redundancy reduction module suitable for faster PET.

4.1. Token redundancy reduction module for FPET

The recently proposed token merging method, ToMe [2], presents a straightforward token reduction strategy that mitigates concerns regarding information loss associated with token pruning. Rather than discarding tokens, ToMe utilizes a bipartite soft matching strategy to merge similar tokens, thereby preserving more information. To minimize information loss in our approach, we incorporate token merging as our chosen token reduction strategy.

ToMe [2], originally designed for in-domain trainingfree or full fine-tuning settings, measures similarities between tokens and gradually merges them across layers to prevent abrupt changes and minimize risk. In contrast, our approach implements the token merging module solely at the middle layer, merging half of the tokens. This modification addresses the potential for sub-optimal merging outcomes in the original ToMe methodology in the context of PET, where early-stage merging based on a similarity matrix may not fully reflect the impact of adapters. Unlike ToMe [2], which conducts layer-by-layer merging of 8 tokens per layer (totaling 96 tokens), our approach merges 98 tokens at the 6th layer of a 12-layer ViT-B/16. By implementing the merging module only once at the middle layer, where token redundancy is sufficiently high, we aim to enhance accuracy while also conserving computation memory by avoiding the repeated computation of the similarity matrix across layers.

4.2. Bipartite differentiable matching

To further enhance our method, we reformulate the matching process to be differentiable for optimization, as depicted in Fig. 2b. We begin by defining the token division sets, X_A and X_B , based on observations in Fig. 2 which indicate that a naïve bipartite soft matching split results in vertical stripe patterns when tokens are reorganized in their original image format. In this arrangement, tokens in X_B can only merge with adjacent side tokens, not with those above or below. To address this, we divide the keys (**K**) of the tokens into two sets using a checkerboard pattern:

$$\mathbf{K}_{A} = \{\mathbf{k}_{ij} \mid (i+j) \mod 2 = 0\}$$

$$\mathbf{K}_{B} = \{\mathbf{k}_{ij} \mid (i+j) \mod 2 = 1\}$$
(9)

Here, kij denotes the two-dimensional index when tokens are reorganized into their original image format. This configuration allows each token in **XB** to merge with adjacent tokens, supporting a more comprehensive merging strategy.

The bipartite soft matching [2] uses the keys (\mathbf{K}) from the self-attention layer, which encapsulates the information within each token, to measure similarity. However, these keys may not be the most optimal feature for similarity. To ensure a more optimal similarity matrix \mathbf{C} , we implement a differentiable matching process and refine the keys (\mathbf{K}) using learnable adapters to achieve a more optimized matching result. The refined keys (\mathbf{K}') are represented as:

$$\mathbf{K}' = \mathbf{K} + s \cdot \operatorname{ReLU}(\mathbf{K}\mathbf{W}_{\operatorname{down}})\mathbf{W}_{\operatorname{up}}$$
(10)

Using these refined keys (**K**'), we compute the similarity matrix $C' \in \mathbb{R}^{N_A \times N_B}$:

$$\mathbf{C}' = \mathbf{K}'_{\mathrm{A}} \times \mathbf{K}'_{\mathrm{B}}^{\dagger} \tag{11}$$

Instead of using the non-differentiable *max* and top-*k* operations, we leverage the matching matrix $\mathbf{C}_{AB} \in \mathbb{R}^{N_A \times N_B}$ where each row is a one-hot vector indicating the closest key in \mathbf{K}_B for each key in \mathbf{K}_A . To construct \mathbf{C}_{AB} , we compute the soft matching matrix $\hat{\mathbf{C}}_{AB} \in \mathbb{R}^{N_A \times N_B}$ as:

$$\hat{\mathbf{C}}_{\mathrm{AB}} = \sigma(\mathbf{C} - \bar{\mathbf{C}}) \tag{12}$$

where $\bar{\mathbf{C}} \in \mathbb{R}^{N_A \times 1}$ represents the average of the top-1 and top-2 values along the last dimension of \mathbf{C} and σ is a sigmoid function. Notably, the top-*k* operation serves only to shift values, without compromising differentiability. For each row of $\hat{\mathbf{C}}_{AB}$, only the top-1 value exceeds 0.5 since the sigmoid function outputs values greater than 0.5 for positive inputs. Thresholding $\hat{\mathbf{C}}_{AB}$ at 0.5, we derive the hard matching matrix \mathbf{C}_{AB} . The \mathbf{X}_{merged} is then expressed as:

$$\mathbf{X}_{\text{merged}} = average(\mathbf{C}_{\text{AB}}^{\top} \times \mathbf{X}_{\text{A}} + \mathbf{X}_{\text{B}})$$
(13)

Here, *average* is a function that adaptively averages tokens based on the number of matched tokens. Each token in X_A is combined with its closest token in X_B and each summed token is averaged by the number of its matched tokens. Since all tokens in X_A are matched to those in X_B , no unmerged tokens remain. Therefore, X_{merged} is returned as the output of our token merging module.

While we reformulate the merging process as a differentiable matrix calculation, the hard matching matrix C_{AB} is non-differentiable. To approximate gradients for C_{AB} , we use a straight-through estimator[1]. Specifically, we redefine the matching matrix as:

$$\tilde{\mathbf{C}}_{AB} = \hat{\mathbf{C}}_{AB} + const(\mathbf{C}_{AB} - \hat{\mathbf{C}}_{AB})$$
 (14)

where *const* is a function that extracts value as a constant from the tensor. Although C_{AB} and \tilde{C}_{AB} hold identical values, \tilde{C}_{AB} is differentiable since gradients are propagated through \hat{C}_{AB} .

Since the matching process is differentiable, gradients are propagated to the learnable adapter that refines the key(\mathbf{K}). This allows for the explicit learning and the optimization of the similarity matrix. However, allowing further backward propagation of these gradients can negatively impact accuracy due to the resulting unnecessary push-andpull effect among tokens, similar to what is observed in contrastive loss. While the push-and-pull effect is very beneficial in self-supervised learning for distinguishing features among different samples, however, such an effect among tokens is undesirable for our task. Therefore, we halt the further propagation of these gradients to the backbone. For more detailed understanding, our code implementation is included in the supplementary materials.

5. Experiments

In this section, we demonstrate the superiority of our proposed method in terms of both efficiency gain and accuracy. For our implementation, we integrate our token merging module at the 6th layer along with a quantized adapter [42] for key (\mathbf{K}) refinement. The training epoch is 100 using AdamW optimizer. Since our proposed method can be applied to off-the-shelf PET methods in plug-and-play manner, we evaluate our method on 5 state-of-the-art PET methods, RepAdapter [28], LoRA [13], AdaptFormer [5], Bi-LoRA [16] and Bi-AdaptFormer [16].

5.1. Datasets

To demonstrate the efficacy of our method, we conducted evaluations across a range of downstream tasks using VTAB-1K [38], which encompasses 19 diverse image classification tasks spanning different domains. VTAB-1K is divided into three categories: Natural, Specialized, and Structured. The Natural category includes classic vision tasks with images captured by standard cameras. In contrast, the Specialized category comprises images from natural scenes captured with specialized equipment, such as those used in satellite or medical imaging. The Struc*tured* category is centered on understanding scene structure, involving tasks like object counting or 3D depth prediction, often with images from simulated environments. Each dataset contains 1000 samples for training and validation. Following prior works [14-16, 22, 40], we train models using training and validation samples, and report the top-1 accuracy on test set.

5.2. Metrics

We assess the performance of our approach using several metrics: accuracy, inference time, FLOPs and GPU memory usage. Accuracy evaluations are conducted on the VTAB-1K dataset, with 'Average' referring to the mean accuracy across the three groups. GPU memory usage is recorded during the training phase with batch size 64, whereas inference time and FLOPs are measured during testing with batch size 256. The inference time metric represents the time taken to process a single image. All experiments are implemented on single NVIDIA A6000 GPUs.

5.3. Comparison to the state-of-the-art methods

We evaluate our method against various state-of-the-art approaches on VTAB-1K, which encompasses a diverse range

Method	Acc (%)	Time (ms)	FLOPs (G)	Mem (GB)	
Traditional Fine-Tuning					
Full	68.9	2.62	17.6	11.9	
Linear	57.6	2.62 (+0.0%)	17.6 (+0.0%)	2.7 (-77.3%)	
Parameter-Efficient Tunii	ıg				
VPT-Deep [14]	72.0	2.79 (+6.5%)	18.5 (+5.1%)	9.8 (-26.5%)	
BitFit [37]	65.2	2.62 (+0.0%)	17.6 (+0.0%)	8.0 (-33.0%)	
SSF [22]	75.7	2.62 (+0.0%)	17.6 (+0.0%)	10.9 (-8.2%)	
FacT-TT [15]	75.2	2.62 (+0.0%)	17.8 (+1.1%)	11.9 (-0.0%)	
RepAdapter [28]	76.1	2.62 (+0.0%)	17.6 (+0.0%)	8.9 (-25.6%)	
LoRA [13]	75.7	2.62 (+0.0%)	17.6 (+0.0%)	8.4 (-29.6%)	
AdaptFormer [5]	76.2	2.68 (+1.5%)	17.6 (+0.0%)	7.6 (-36.0%)	
Bi-LoRA [16]	76.7	2.62 (+0.0%)	17.6 (+0.0%)	8.4 (-29.4%)	
Bi-AdaptFormer [16]	77.0	2.77 (+5.7%)	17.7 (+0.6%)	7.6 (-35.9%)	
Efficiency-gained Parame	eter-Efficie	nt Tuning			
SynQT [39]	72.9	2.83 (+8.0%)	16.84 (-4.3%)	3.6 (-69.9%)	
Pruned RepAdapter [7]	74.8	[1.06, 2.62]	[7.1, 17.6]	[5.9, 8.9]	
Faster Parameter-Efficien	t Tuning (O	Ours)			
RepAdapter	76.1	2.10 (-19.8%)	13.3 (-24.4%)	7.4 (-38.0%)	
LoRA	75.6	2.10 (-19.8%)	13.3 (-24.4%)	7.1 (-40.3%)	
AdaptFormer	76.2	2.12 (-19.1%)	13.5 (-23.3%)	6.2 (-47.9%)	
Bi-LoRA	76.4	2.10 (-19.8%)	13.3 (-24.4%)	7.1 (-40.3%)	
Bi-AdaptFormer	77.0	2.17 (-17.2%)	13.5 (-23.3%)	6.2 (-47.8%)	

Table 1. Comparison with state-of-the-art methods in terms of accuracy, inference time per image, FLOPs, and GPU memory usage during training. In terms of efficiency, we present both the absolute values and the relative gap in comparison to the full finetuning method. For Pruned RepAdapter [7], the values are shown as ranges to indicate its inconsistent efficiency across datasets.

of downstream datasets. In this subsection, we provide detailed comparisons of PET methods and efficiency-gained PET methods, respectively.

5.3.1. PET methods.

In this section, we compare our method with full finetuning, linear probing which trains classification head only, VPT [14], NOAH [40], LoRA [13], SSF [22], Adapt-Former [5], BitFit [37], FacT-TT [15], Bi-LoRA and Bi-Adaptformer [16]. All baseline models utilize the ViT-B/16, pre-trained on ImageNet-21K in a supervised fashion, as their backbone. For LoRA [13] and AdaptFormer [5], we set the hidden dimension to 8. The settings for other baseline models follow the configurations reported in their respective original papers.

As shown in Tab. 1 and Fig. 1, our method demonstrates constant efficiency gain for both training and inference. Remarkably, we achieve faster inference speeds compared to PET methods, including traditional fine-tuning approaches that do not necessitate additional computation during inference. While certain PET methods [13, 15, 22, 37] have been developed with the goal of achieving no-latency inference, they have only succeeded in matching the speed of traditional fine-tuning methods, leaving inherent limitations. Our method, however, overcomes these constraints with a 19.8% increase in speed over the original pre-trained models, demonstrating the superiority of our approach. In

terms of GPU memory consumption, we achieve a 40% reduction compared to full fine-tuning. While linear probing exhibits greater memory efficiency owing to its simplicity, it significantly underperforms in accuracy, inference speed, and FLOPs compared to our method.

The higher FLOPs observed in VPT-Deep [14] compared to other methods suggest that an increase in the number of tokens leads to a quadratic rise in computational complexity. In contrast, our method succeed in achieving an average reduction of 24% in FLOPs by reducing token redundancy, compared to each original implementation. Our method attains efficiency gains with minimal impact on accuracy. Notably, when applied to Bi-AdaptFormer [16], compared to the original implementation, it results in negligible accuracy loss while achieving improvements of 21.7%, 23.7% and 18.4% in inference time, FLOPs and GPU memory usage respectively.

Tab. 2 presents the comprehensive results of our comparison with state-of-the-art models. Through our framework, while realizing efficiency gains, models still retain their superiority in both accuracy and parameter efficiency. As we employ only a very lightweight adapter to refine the similarity matrix, less than 0.005M, the increase in the number of trainable parameters is minimal.

5.3.2. Efficiency-gained PET methods.

We further compare our method with efficiency-gained PET methods including SynQT [39] and Pruned RepAdapter [7]. As shown in Tab. 1, while SynQT significantly reduces GPU memory usage during training, it introduces a much slower inference speed compared to the backbone model, and its FLOPs reduction is much smaller than ours. Additionally, as detailed in Tab. 2, SynQT employs $2.3 \times$ to $16 \times$ more learnable parameters than our approach, resulting in lower accuracy.

Pruned RepAdapter [7] exhibits varying levels of efficiency gains depending on the dataset, as it adaptively determines the number of layers to prune during the pre-training stage for each dataset. In Tab. 1, the maximum and minimum efficiency gain is presented. For the EuroSAT [10] dataset, only 5 layers are used leading to a 59%, 57%, and 50% reduction in inference time, FLOPs, and GPU memory usage respectively. However, for datasets such as dspr-Ori [11] and sNORB-Ele [20], all 12 layers are utilized, resulting in efficiency metrics nearly identical to those of the original RepAdapter [28]. Although Pruned RepAdapter [7] achieves average efficiency gains, its full-layer implementation is not suitable for resource-constrained environments making real-world applications challenging. In contrast, our method consistently achieves efficiency gains across different datasets, offering a more practical and feasible solution for real-world scenarios.

		Natural				Specialized				Structured											
	# param (M)	Average	CIFAR100	Caltech101	DTD	Flowers	Pets	NHAS	Sun397	Camelyon	EuroSAT	Resisc45	Retinopathy	Clevr-Count	Clevr-Dist	DMLab	KITTI-Dist	dSpr-Loc	dSpr-Ori	sNORB-Azim	sNORB-Ele
Traditional Fine-Tuning																					
Full Linear	85.8 0	68.9 57.6	68.9 64.4	87.7 85.0	64.3 63.2	97.2 97.0	86.9 86.3	87.4 36.6	38.8 51.0	79.7 78.5	95.7 87.5	84.2 68.5	73.9 74.0	56.3 34.3	58.6 30.6	41.7 33.2	65.5 55.4	57.5 12.5	46.7 20.0	25.7 9.6	29.1 19.2
Parameter-Efficient Tuni	ng																				
VPT-Shallow [14] VPT-DEEP [14] NOAH [40] BitFit [37] SSF [22] FacT-TT [15] RepAdapter [28] LoRA [13] AdaptFormer [5] Bi-LoRA [16] Bi-AdaptFormer [16]	$\begin{array}{c} 0.06\\ 0.53\\ 0.36\\ 0.10\\ 0.24\\ 0.04\\ 0.22\\ 0.29\\ 0.16\\ 1.17\\ 0.64 \end{array}$	67.8 72.0 75.5 65.2 75.7 75.3 76.1 75.7 76.2 76.7 76.2 76.7 77.0	77.7 78.8 69.6 72.8 69.0 71.3 72.4 70.9 72.0 72.1 74.1	86.9 90.8 92.7 87.0 92.6 89.6 91.6 93.0 92.7 91.7 92.4	62.6 65.8 70.2 59.2 75.1 70.7 71.0 70.0 70.2 71.2 72.1	97.5 98.0 99.1 97.5 99.4 98.9 99.2 99.1 99.3 99.1 99.3	87.3 88.3 90.4 85.3 91.8 91.0 91.4 91.2 91.0 91.4 91.6	74.5 78.1 86.1 59.9 90.2 87.8 90.7 86.0 87.5 90.2 89.0	51.2 49.6 53.7 51.4 52.9 54.6 55.1 55.8 54.8 55.8 55.8 56.3	78.2 81.8 84.4 78.7 87.4 85.2 85.3 87.1 87.4 87.0 88.2	92.0 96.1 95.4 91.6 95.9 95.5 95.9 94.6 95.2 95.4 95.2	75.6 83.4 83.9 72.9 87.4 83.4 84.6 82.4 85.2 85.5 86.0	72.9 68.4 75.8 69.8 75.5 75.7 75.9 74.7 75.2 75.5 76.2	50.5 68.5 82.8 61.5 75.9 82.0 82.3 83.5 83.6 83.1 83.9	58.6 60.0 68.9 55.6 62.3 69.0 68.0 63.1 62.6 64.1 63.6	40.5 46.5 49.9 32.4 53.3 49.8 50.4 51.7 52.1 52.2 53.0	67.1 72.8 81.7 55.9 80.6 80.0 79.9 79.3 81.0 81.3 81.4	68.7 73.6 81.8 66.6 77.3 79.2 80.4 85.2 86.4 86.2	36.1 47.9 48.3 40.0 54.9 48.4 49.2 51.8 53.1 53.5 54.8	20.2 32.9 32.8 15.7 29.5 34.2 38.6 33.5 34.5 36.7 35.2	34.1 37.8 44.2 25.1 37.9 41.4 41.0 43.6 40.3 44.4 41.3
Efficiency-gained Param	eter-Effi	cient Tuning																			
SynQT [39] Pruned RepAdapter [7]	2.72 0.18	72.9 74.8	70.9 71.4	89.7 87.3	68.8 68.1	98.5 96.0	89.6 89.9	77.8 89.3	50.6 53.4	82.3 85.0	96.7 95.3	83.5 81.9	75.2 75.2	71.8 80.9	62.7 69.8	48.5 50.5	75.4 80.7	74.1 80.5	49.0 47.1	31.7 35.7	36.1 41.0
Faster Parameter-Efficient	nt Tuning	g (Ours)																			
RepAdapter LoRA AdaptFormer Bi-LoRA Bi-Adaptformer	0.23 0.30 0.17 1.18 0.64	76.1 75.6 76.2 76.4 77.0	72.1 70.1 71.3 71.9 74.1	91.5 92.7 93.5 91.1 92.8	71.8 69.4 69.9 70.9 72.5	99.3 99.1 99.3 99.1 99.4	90.7 90.8 90.7 90.5 91.1	90.3 85.4 87.0 89.4 89.6	55.0 55.6 54.7 55.9 56.2	85.2 87.2 87.5 87.4 88.3	96.2 94.6 95.1 94.7 94.9	84.5 82.5 84.5 84.4 86.3	75.6 74.1 76.2 74.9 75.3	82.2 83.0 83.6 83.5 83.8	67.7 63.4 63.1 65.1 63.0	49.7 50.6 52.2 52.1 52.8	79.9 81.6 81.3 79.7 81.4	82.2 84.7 87.1 85.8 85.7	48.7 51.5 54.1 54.2 54.4	36.9 34.3 33.5 36.7 35.9	41.7 43.3 40.2 44.4 42.2

Table 2. Comparison with state-of-the-art methods on VTAB-1K [38] benchmark. Average indicates average accuracy over three groups. # param denotes the number of learnable parameters.

5.4. Ablation study

In this section, we present ablation studies to further investigate specific efficacy of our proposed methodology.

5.4.1. Comparison to other token reduction method.

As shown in Tab. 3, our proposed method achieves superior efficiency gains compared to existing token reduction techniques. Unlike ToMe [2], which merges a fixed number of tokens at every layer, LTMP [3] dynamically determines the number of tokens to prune using a learnable threshold. To enable gradient flow, LTMP masks rather than reduces the number of tokens, requiring a modified attention module with masked softmax, which incurs memory overhead as outlined in Tab. 3. Moreover, LTMP does not yield inference-time efficiency gains when the batch size exceeds one. While ToMe improves inference speed, it fails to reduce GPU memory usage due to bipartite soft matching across all layers. Both ToMe and LTMP incur larger accuracy drops than our method, highlighting the effectiveness of our design in balancing performance and efficiency within the PET framework.

5.4.2. Effectiveness of each component.

In Tab. 3, we further evaluate the effectiveness of each component through progressive integration. Applying bipartite soft matching (BSM)[2] at a middle layer improves both accuracy and efficiency over the original ToMe[2] implementation. Incorporating the checkerboard pattern yields a slight accuracy gain, demonstrating the benefit of spatial priors. Adding the key refinement adapter without gradient stopping, however, significantly degrades accuracy due to interference in token representations, inducing unnecessary push and pull effect among tokens. Lastly, applying gradient stopping leads to our best overall performance, achieving strong accuracy while maintaining efficiency.

5.4.3. Trade-offs between efficiency and accuracy.

In Fig. 3, we present the trade-off between efficiency and accuracy when applying our token reduction module at different transformer layers. The module can be flexibly applied at various layers to balance computational cost and model performance. When applied at layer 6, our method achieves the upper-bound accuracy of 76.22%, equivalent to the original PET method without token merging, while reducing FLOPs by 23.3%. At layer 4, a slight accuracy drop of 0.57% yields a 31.82% FLOPs reduction, and at layer 2, a 1.93% drop corresponds to a 40.34% reduction. Despite these trade-offs, our method remains competitive with both the original and efficiency-gained PET methods.

We also compare our method with other token merging strategies, including bipartite soft matching (BSM) [2], average pooling, and max pooling. All methods except max pooling merge tokens by averaging matched pairs. Aver-

Method	$\Delta Acc (\%)$	Δ Time (ms)	Δ Mem (GB)
w/o merging	0	0	0
LTMP [3]	-0.38	+0.01	+5.84
ToMe [2]	-0.27	-0.56	-0.21
BSM [2] at 6th layer	-0.18	-0.61	-1.38
+ checkerboard	-0.17	-0.61	-1.38
+ key refinement (w/o gradient stopping)	-1.70	-0.61	-1.41
Ours	-0.05	-0.62	-1.41

Table 3. Performance on other token reduction strategies. All models are ViT-B/16 consist of 12 transformer layers with Bi-AdaptFormer [16]. Each model starts with 197 tokens. LTMP reduces a variable number of tokens, while ToMe merges 8 tokens at each of the 12 layers. For BSM at 6*th* layer and the methods below, 98 tokens are merged at the 6*th* layer only. All methods are applied during both training and inference time.

Method	AdaptFormer [5]	FPET-AdaptFormer	LoRA [13]	FPET-LoRA
Training time (s/it)	5.62	4.52 (-19.57%)	6.24	5.03 (-19.39%)

Table 4. Average training time	across VTAB-1K [38] datasets.
--------------------------------	-------------------------------

age pooling relies on deterministic matching, BSM uses a heuristic and non-differentiable approach, while our method employs a learnable bipartite matching technique. Compared to others, our method achieves higher accuracy with negligible FLOPs overhead. At layer 0, where tokens are minimally processed and adapters are not yet applied, the regional prior is competitive to similarity-based methods. From layer 1 onward, our method consistently outperforms naive pooling, demonstrating more optimal token reduction. Notably, BSM shows degraded accuracy particularly in early layers, underscoring robustness and practicality of our method. Detailed numerical results corresponding to Fig. 3 are provided in the supplementary materials.

5.5. Further analysis

In this section, we further highlight the superiority of FPET. As shown in Tab. 4, FPET significantly improves training efficiency, reducing training time by 19% compared to the original implementation. The latency introduced by our module is also negligible, accounting for only 0.56% of the total pipeline latency. We evaluate FPET on other backbones, including DeiT-S and ViT-L, as reported in Tab. 5. Further results covering additional backbones, fewshot learning on the FGVC dataset, cross-modal retrieval and other analysis are provided in the supplementary materials, further demonstrating the robustness and versatility of our approach. Visualizations of the merged tokens are also included for qualitative analysis.



Figure 3. Trade-off between FLOPs (right y-axis, bar plot) and accuracy (left y-axis, line plot). All models are ViT-B/16 consisting of 12 transformer layers with AdaptFormer [5]. BSM refers to bipartite soft matching [2]. All methods reduce 196 tokens to 98 tokens at different layers.

Model	Iodel Method		Time (ms)	FLOPs (G)	Mem (GB)		
ViT-L	LoRA [13]	76.0	8.67	61.8	19.6		
	FPET-LoRA	76.0	6.60 (-23.9%)	46.7 (-24.4%)	15.5 (-21.0%)		
DeiT-S	AdaptFormer[5]	70.3	1.13	4.8	4.4		
	FPET-AdaptFormer	70.3	0.94(-20.7%)	3.7(-23.6%)	3.8(-12.5%)		

Table 5. Model performance on other backbones.

6. Conclusion

In this paper, we extend the concept of parameter efficiency in parameter-efficient tuning (PET) by exploring both inference latency and training computational efficiencyto enhance the applicability of PET. We propose Faster Parameter-Efficient Tuning (FPET), a novel framework that formulates PET as a token redundancy reduction problem. Our approach formulates token reduction in a parameterefficient and fully differentiable manner, enabling optimal token reduction for PET. Our FPET significantly improves inference speed and training efficiency while attaining comparable performance to the previous PET methods, demonstrating the effectiveness of the proposed token redundancy reduction module. As our FPET can be seamlessly integrated with existing PET techniques, we hope our study will foster research and provide a foundation of practical PET for real-world applications.

7. Acknowledgement

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2025-00515741) and the Yonsei Signature Research Cluster Program of 2024 (2024-22-0161).

References

- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 2, 5
- [2] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. In *ICLR*, 2023. 2, 3, 4, 7, 8
- [3] Maxim Bonnaerens and Joni Dambre. Learned thresholds token merging and pruning for vision transformers. In *TMLR*, 2023. 2, 3, 7, 8
- [4] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001. 3
- [5] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. In *NeurIPS*, pages 16664–16678, 2022. 1, 2, 3, 5, 6, 7, 8
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2
- [7] Ruizi Han and Jinglei Tang. Straightforward layer-wise pruning for more efficient visual adaptation. In *ECCV*, 2024.
 1, 2, 6, 7
- [8] Shwai He, Liang Ding, Daize Dong, Miao Zhang, and Dacheng Tao. Sparseadapter: An easy approach for improving the parameter-efficiency of adapters. In *EMNLP*, 2023. 1, 2
- [9] Xuehai He, Chunyuan Li, Pengchuan Zhang, Jianwei Yang, and Xin Eric Wang. Parameter-efficient model adaptation for vision transformers. In AAAI, pages 817–825, 2023. 1, 2
- [10] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations* and Remote Sensing, pages 2217–2226, 2019. 6
- [11] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017. 6
- [12] 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 *ICML*, pages 2790–2799, 2019. 1, 2
- [13] 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. In *ICLR*, 2022. 1, 2, 3, 5, 6, 7, 8
- [14] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, pages 709–727, 2022. 1, 2, 5, 6, 7

- [15] Shibo Jie and Zhi-Hong Deng. Fact: Factor-tuning for lightweight adaptation on vision transformer. In AAAI, pages 1060–1068, 2023. 1, 2, 6, 7
- [16] Shibo Jie, Haoqing Wang, and Zhi-Hong Deng. Revisiting the parameter efficiency of adapters from the perspective of precision redundancy. In *ICCV*, pages 17217–17226, 2023. 1, 2, 5, 6, 7, 8
- [17] Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. In *NeurIPS*, pages 1022–1035, 2021. 1, 2
- [18] Zhenglun Kong, Peiyan Dong, Xiaolong Ma, Xin Meng, Wei Niu, Mengshu Sun, Xuan Shen, Geng Yuan, Bin Ren, Hao Tang, et al. Spvit: Enabling faster vision transformers via latency-aware soft token pruning. In *ECCV*, pages 620–640, 2022. 2
- [19] Hyeongjun Kwon, Taeyong Song, Somi Jeong, Jin Kim, Jinhyun Jang, and Kwanghoon Sohn. Probabilistic prompt learning for dense prediction. In *CVPR*, pages 6768–6777, 2023. 2
- [20] Yann LeCun, Fu Jie Huang, and Leon Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *CVPR*, pages II–104, 2004. 6
- [21] Jonathan Li, Will Aitken, Rohan Bhambhoria, and Xiaodan Zhu. Prefix propagation: Parameter-efficient tuning for long sequences. In ACL, 2023. 1, 2
- [22] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. In *NeurIPS*, pages 109–123, 2022. 1, 2, 5, 6, 7
- [23] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. In *ICLR*, 2022. 2
- [24] Dongyang Liu, Meina Kan, Shiguang Shan, and Xilin CHEN. A simple romance between multi-exit vision transformer and token reduction. In *ICLR*, 2024. 2
- [25] Weiyang Liu, Zeju Qiu, Yao Feng, Yuliang Xiu, Yuxuan Xue, Longhui Yu, Haiwen Feng, Zhen Liu, Juyeon Heo, Songyou Peng, Yandong Wen, Michael J. Black, Adrian Weller, and Bernhard Schölkopf. Parameter-efficient orthogonal finetuning via butterfly factorization. In *ICLR*, 2024. 2
- [26] Yajing Liu, Yuning Lu, Hao Liu, Yaozu An, Zhuoran Xu, Zhuokun Yao, Baofeng Zhang, Zhiwei Xiong, and Chenguang Gui. Hierarchical prompt learning for multi-task learning. In CVPR, pages 10888–10898, 2023. 2
- [27] S. Lloyd. Least squares quantization in pcm. *IEEE Transac*tions on Information Theory, 28(2):129–137, 1982. 3
- [28] Gen Luo, Minglang Huang, Yiyi Zhou, Xiaoshuai Sun, Guangnan Jiang, Zhiyu Wang, and Rongrong Ji. Towards efficient visual adaption via structural re-parameterization. arXiv preprint arXiv:2302.08106, 2023. 2, 5, 6, 7
- [29] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning. In *NeurIPS*, pages 26462–26477, 2022. 1, 2

- [30] Jungin Park, Jiyoung Lee, and Kwanghoon Sohn. Dualpath adaptation from image to video transformers. In *CVPR*, pages 2203–2213, 2023.
- [31] Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapterfusion: Nondestructive task composition for transfer learning. In ACL, pages 487–503, 2021. 1, 2
- [32] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, pages 13937–13949, 2021. 2
- [33] Kihyuk Sohn, Huiwen Chang, José Lezama, Luisa Polania, Han Zhang, Yuan Hao, Irfan Essa, and Lu Jiang. Visual prompt tuning for generative transfer learning. In *CVPR*, pages 19840–19851, 2023. 2
- [34] 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 *CVPR*, pages 7725–7735, 2023. 1, 2
- [35] Hongxu Yin, Arash Vahdat, Jose Alvarez, Arun Mallya, Jan Kautz, and Pavlo Molchanov. Adavit: Adaptive tokens for efficient vision transformer. In *CVPR*, pages 12299–12308, 2021. 2
- [36] Hongxu Yin, Arash Vahdat, Jose M. Alvarez, Arun Mallya, Jan Kautz, and Pavlo Molchanov. A-vit: Adaptive tokens for efficient vision transformer. In *CVPR*, pages 10809–10818, 2022. 2
- [37] Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In ACL, 2022. 1, 2, 6, 7
- [38] Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruyssen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, et al. The visual task adaptation benchmark. In arXiv preprint, arXiv:1910.04867, 2019. 1, 2, 5, 7, 8
- [39] Taolin Zhang, Jiawang Bai, Zhihe Lu, Dongze Lian, Genping Wang, Xinchao Wang, and Shu-Tao Xia. Parameter-efficient and memory-efficient tuning for vision transformer: a disentangled approach. In ECCV, 2024. 1, 2, 6, 7
- [40] Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. Neural prompt search. In arXiv preprint, arXiv:2206.04673, 2022. 1, 2, 5, 6, 7
- [41] Zhen-Ru Zhang, Chuanqi Tan, Haiyang Xu, Chengyu Wang, Jun Huang, and Songfang Huang. Towards adaptive prefix tuning for parameter-efficient language model fine-tuning. In ACL, pages 1239–1248, 2023. 1, 2
- [42] Hengyuan Zhao, Hao Luo, Yuyang Zhao, Pichao Wang, Fan Wang, and Mike Zheng Shou. Revisit parameter-efficient transfer learning: A two-stage paradigm. In *arXiv preprint*, 2023. arXiv:2303.07910. 1, 5
- [43] Jiawen Zhu, Simiao Lai, Xin Chen, Dong Wang, and Huchuan Lu. Visual prompt multi-modal tracking. In CVPR, pages 9516–9526, 2023. 2