

---

# Parameter and Computation Efficient Transfer Learning for Vision-Language Pre-trained Models

---

Qiong Wu<sup>1,2</sup>, Wei Yu<sup>1,2</sup>, Yiyi Zhou<sup>1,2</sup>, Shubin Huang<sup>1</sup>, Xiaoshuai Sun<sup>1,2</sup>, Rongrong Ji<sup>1,2</sup>

<sup>1</sup> Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University, 361005, P.R. China.

<sup>2</sup> Institute of Artificial Intelligence, Xiamen University, 361005, P.R. China.

{qiong, weiyu}@stu.xmu.edu.cn, zhuyiyi@xmu.edu.cn,  
shubinhuang@stu.xmu.edu.cn, {xssun, rrji}@xmu.edu.cn

## Abstract

With ever increasing parameters and computation, vision-language pre-trained (VLP) models exhibit prohibitive expenditure in downstream task adaption. Recent endeavors mainly focus on parameter efficient transfer learning (PETL) for VLP models by only updating a small number of parameters. However, excessive computational overhead still plagues the application of VLPs. In this paper, we aim at *parameter and computation efficient transfer learning* (PCETL) for VLP models. In particular, PCETL not only needs to limit the number of trainable parameters in VLP models, but also to reduce the computational redundancy during inference, thus enabling a more efficient transfer. To approach this target, we propose a novel *dynamic architecture skipping* (DAS) approach towards effective PCETL. Instead of directly optimizing the intrinsic architectures of VLP models, DAS first observes the significances of their modules to downstream tasks via a reinforcement learning (RL) based process, and then skips the redundant ones with lightweight networks, *i.e.*, adapters, according to the obtained rewards. In this case, the VLP model can well maintain the scale of trainable parameters while speeding up its inference on downstream tasks. To validate DAS, we apply it to two representative VLP models, namely ViLT and METER, and conduct extensive experiments on a bunch of VL tasks. The experimental results not only show the great advantages of DAS in reducing computational complexity, *e.g.*  $-11.97\%$  FLOPs of METER on VQA2.0, but also confirm its competitiveness against existing PETL methods in terms of parameter scale and performance. Our source code is given in our appendix.

## 1 Introduction

Inspired by the great success in natural language processing (NLP) [7, 15, 26, 30], large-scale pre-training on massive image-text pairs also becomes the *de-facto* standard in vision-language research [4, 20, 29, 53]. To accommodate the large-scale pre-training corpora, vision-language pre-trained (VLP) models [4, 8, 20, 24, 29, 45, 59] often adopt Transformer-based networks with sheer sizes of parameters and computations. In this case, directly transferring these VLP models to downstream tasks is excessively expensive in terms of memory footprint and computation overhead.

To reduce the costs of pre-training models, recent advances resort to *parameter efficient transfer learning* (PETL) for affordable downstream task adaptations [10, 22, 31, 48, 57, 59, 60]. In particular, the PETL methods aim to save the memory usage for downstream tasks by only updating or inserting a small number of trainable parameters rather than fully tuning the entire model. For instance, prompt-tuning methods [1, 5, 22, 30, 31, 40, 42, 59, 60] expand the input sequence with hand-craft or learnable tokens to bridge the gap between pre-training and downstream tasks. Practitioners

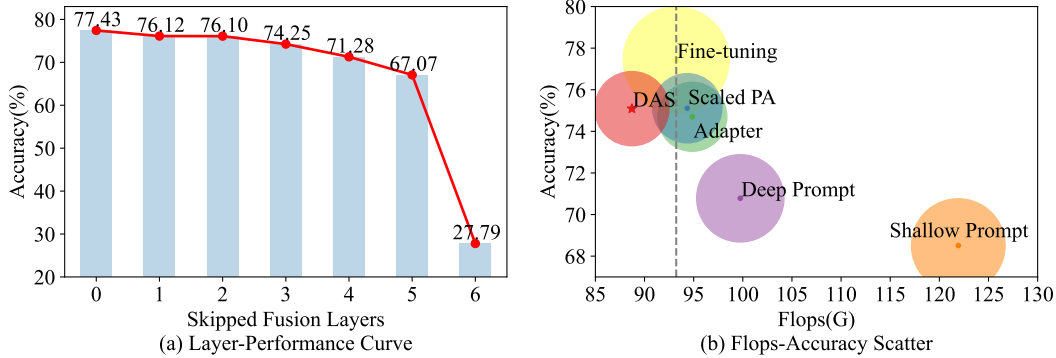


Figure 1: (a) The performance of METER [8] is barely affected when skipping a certain number of its Transformer layers. (b) The comparison on VQA2.0 between the conventional PETL methods [13, 17, 22, 30, 48] and the proposed *Dynamic Architecture Skipping* (DAS) for METER. The circle size represents the memory footprint. DAS is the only method faster than the original VLP model.

also insert lightweight neural networks called *Adapter* [10, 16, 36, 37, 48, 57] into the pre-trained models, thereby projecting the hidden features onto the semantic space of downstream tasks. More recently, these PETL methods have been successfully introduced to VLP models [10, 27] for either prompt-based image classification [22, 42, 59, 60] or conventional VL tasks like *visual question answering* [47, 48]. Despite the great successes, PETL methods still cannot reduce the computation complexity of VLP models, which is of more significance for applications.

In this paper, we study a novel problem called *parameter and computation efficient transfer learning* (PCETL). To achieve more efficient downstream task adaptations, PCETL is not only expected to maintain the scale of trainable parameters similar to PETL, but more importantly, also needs to reduce the computation complexity of pre-training models, thereby speeding up their inference on downstream tasks. In existing works, the efficiency of the network itself is largely attributed to its manually [6, 18, 35, 44] or automatically structural designs [51, 54]. Although the computation complexity can be further reduced by the compression methods, such as *pruning* [39, 3], *quantification* [9, 25, 58] or *distillation* [2, 21], these approaches usually require retraining after optimizing the network architecture, which is not applicable to the VLP models that are well pre-trained on massive data. On one hand, the large-scale pre-training data still needs a certain model capacity to learn these prior knowledge, thus it is hard to obtain a good trade-off between performance and computation overhead for the pre-training objectives. On the other hand, devising a small and effective model for each downstream task is still laborious and expensive, which also contradicts the target of PETL [16, 30], since fully fine-tune is often required.

In this case, we argue that the key to PCETL is to explore the parameter and computation redundancy in existing VLP models. It is generally assumed that the model scale is proportional to the complexity of the task [1, 14, 31]. To robustly serve a variety of downstream tasks, VLP models are pre-trained by multiple pre-training objectives based on tens of millions of image-text pairs [11, 41, 42, 46]. In this case, the excessive parameters are suitable for pre-training, but prone to redundant for a downstream task. As shown in Fig. 1-(a), the performance of METER [8] on VQA is barely affected when skipping a certain number of its Transformer layers. This empirical result also suggests that exploring a short-cut pathway in existing VLP models is a feasible way for PCETL.

To this end, we propose a novel *Dynamic Architecture Skipping* (DAS) approach towards efficient transfer learning for VLP models. By observing the module redundancy of VLP models, DAS can realize the optimal subnetwork routing of VLP models for a downstream task, thereby reducing the computation during inference. In practice, DAS regards this process as a  $k$ -armed bandit problem, and evaluates the importance of each VL layer/block to the downstream task via numerous subnetwork samplings and quick validations. Thus, the obtained rewards can be used to reflect the redundancy of VL modules and determine which layers to be skipped. Meanwhile, to achieve parameter efficiency, we also adopt lightweight networks, *i.e.* Adapter [16, 48], to serve the hidden feature adaptations and the short-cut connections of DAS for VLP models.

To validate DAS, we apply it to two representative VLP models, namely METER [8] and ViLT [24], on three VL benchmarks including VQA2.0 [11], NLVR<sup>2</sup> [46] and Flickr30K [41]. The experimental

results not only show the competitive performance of DAS against the fully finetune and PETL methods [13, 17, 22, 48], but also witness its great advantage in reducing the computation complexity of VLP models. For instance, DAS can help METER achieve 96.60% performance of full tuning on the VQA2.0 benchmark with only 1.65% trainable parameters, while decreasing 11.97% FLOPs. For the practical deployment of a specific VL task, DAS can reduce up to 93.75% parameters of the VLP models<sup>1</sup>. These results well confirm our assumption about the redundancy of VLP models on downstream tasks, and also validated the design of the proposed DAS.

Overall, our contributions can be summarized as three-fold:

- We raise a new problem called *parameter and computation efficient transfer learning* (PCETL) for VLP models, which not only requires to keep the scale of training parameters but also needs to reduce the computation complexity of VLP models on downstream tasks.
- We propose a novel *Dynamic Architecture Skipping* (DAS) for PCETL, which can explore the optimal short-cut pathway in VLP models with the combination of parallel adapters.
- On two VLP models and three benchmark datasets, the proposed DAS not only reduces the computation overhead by a large extent, *e.g.*,  $-11.97\%$  FLOPs of METER on VQA2.0, but also is on par with existing PETL methods in terms of parameter and performance.

## 2 Related Work

### 2.1 Vision-Language Pre-training

In recent years, the advancement in natural language processing (NLP) [26, 30] also sparks the prevalence of large-scale pre-training in vision-language (VL) research [4, 8, 20, 24, 29, 45, 59]. In particular, VL pre-training also accomplishes self-supervised learning on massive image-text pairs based on the generative prediction tasks, *e.g.* *masked language modeling* (MLM) and *masked image modeling* (MIM). Furthermore, *Image Text Matching* (ITM) [8, 24] is also applied to align two modalities. In terms of network architecture, most VLP models are equipped with two modality encoders to extract the visual and text features, *e.g.* BERT [7] and Faster-RCNN [43], respectively, based on which a stack of Transformer-based layers are deployed for cross-modal fusions [4, 8, 19, 20, 24, 28, 32, 33, 45, 50, 52]. For instance, ViL-BERT [33] and LXMERT [50] contain two independent Transformer-based branches [52] for region and text feature extractions, and another Transformer block is used for cross-modal interaction. To simplify the framework, Visual-BERT [28], VL-BERT [45] and UNITER [4] abandon additional branches and directly feed features into a single Transformer network. Then Pixel-BERT [20], CLIP-ViL [32], and METER [8] break the limitation of object detection backbones by directly applying grid features for multi-modal pre-training. To further simplify the model complexity, ViLT [24] directly feeds the word embeddings and image patches into the Transformer blocks. Additionally, CLIP [42] applies cross-modal contrastive learning for vision-language alignment with a shallow fusion layer for prediction. Overall, these VLP models often require more network branches due to the increase of modalities, resulting in more parameter and computation overheads. In this paper, we present the first attempt to evaluate their network redundancy on downstream tasks.

### 2.2 Parameter Efficient Transfer Learning

Parameter Efficient Transfer Learning (PETL) [10, 12, 13, 16, 17, 36–38, 47–49, 56, 57] aims to approach the fully-tuned performance on downstream by updating a small number of parameters. One of the main methodology in PETL is prompt-tuning [1, 5, 22, 30, 31, 40, 42, 59, 60], which is originally designed for large pre-trained language models such as GPT-3 [1]. Concretely, the hand-craft prompts [40, 42] expand the original input sequence with natural language and regard all problems as a generation task. To better fit downstream tasks, soft prompt tuning methods [22, 30, 60] replace the handcraft prompts with a sequence of trainable embeddings. In addition to prompt-tuning, adapter-based [10, 16, 36, 37, 48, 57] methods insert lightweight feed-forward networks into VLP models, and these methods transfer VLP models by projecting hidden features onto the downstream distributions [16, 48]. Furthermore, LoRA [17] is proposed to transfer VLP models without additional calculation overhead in the inference stage by updating the low-rank parts of the original parameters.

<sup>1</sup>When the model is only deployed for a task, the skipped layers can be also removed during deployment.

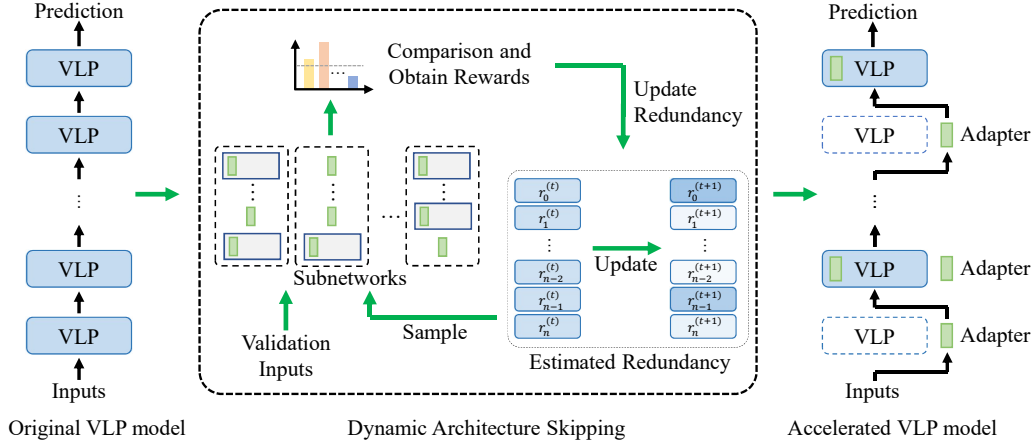


Figure 2: Illustration of *Dynamic Architecture Skipping* (DAS). DAS regards the network skipping as a  $k$ -armed bandit problem, and evaluates the redundancy of each VL layer/block via numerous subnetwork samplings. The accumulated rewards are used to determine which layers can be skipped, and adapters are also used for feature adaptations and short-cut connections.

Besides, Side-tuning [56] runs in parallel with the pre-trained models to adapt downstream tasks while overcoming the constraint from the concrete structure. In addition, LST [47] stacks the outputs of the pre-trained modules in a parallel path. Without feedback to the VLP model, LST alleviates the memory requirement in the transfer while increasing the computation overhead. Compared to fine-tuning the entire model, PETL methods significantly improve the efficiency in transferring VLP models to downstream tasks. However, all of the above methods take the original VLP model as the upper bound of inference efficiency. In this paper, the proposed DAS method is the first to reduce the computation of VLP models while maintaining competitive performance.

### 3 Preliminary

We first revisit the principle of PETL methods for VLP models. Given a vision-language pre-trained (VLP) model, denoted as  $G(\cdot)$ , the target of PETL is to achieve the parameter-efficient adaption on the downstream task, which can be summarized as

$$\operatorname{argmin}_{\sigma} \mathcal{L}(G(I, T | [\theta, \sigma])), \quad (1)$$

where  $\theta = \{\theta_1, \theta_2, \dots, \theta_n\}$  represent the parameters of  $n$  layers in the VLP model, and  $\theta$  is usually frozen in PETL.  $(I, T)$  denotes the image-text pair, and  $\sigma$  is a small number of updated parameters. Since all VLP layers are reserved on downstream tasks, PETL methods can only reduce the parameter expenditure but not the computation of VLP models. Moreover, most PETL methods often incur non-negligible latency during inference [17, 34].

According to the observation in Fig 1-(a), there exists obvious redundancy in the VLP models. To this end, the objective of the proposed task of *parameter and computation efficient transfer learning* (PCETL) can be defined by

$$\operatorname{argmin}_{\sigma, \mathbf{K}} \mathcal{L}(G(I, T | [\theta_{\mathbf{K}}, \sigma])), \quad (2)$$

where  $\theta_{\mathbf{K}} = \{\theta_{k_1}, \theta_{k_2}, \dots, \theta_{k_m}\} \in \theta$  is the parameters of VLP modules except the skipped ones. Via skipping the redundant layers in VLP models and the combination of PETL methods, VLP models can accelerate the inference speed and maintain the scale of updated parameters.

## 4 Dynamic Architecture Skipping

### 4.1 Redundancy Estimation

In this paper, we propose a novel transfer learning approach called *Dynamic Architecture Skipping* (DAS) towards the parameter and computation efficient adaption of VLP models.

---

**Algorithm 1** Dynamic Architecture Skipping

---

**Require:** The training and validation sets,  $D_t, D_v$ . VLP modules  $\{\theta_i\}$ . Adapters  $\{\sigma_i\}$ .

**Ensure:** The optimal network structure  $G_s$ .

Initialize the degree of redundancy of all modules,  $\mathbf{r}^{(0)} = [\mathbf{r}_1^{(0)}, \mathbf{r}_2^{(0)}, \dots, \mathbf{r}_n^{(0)}]$   
**for**  $t$  in  $T$  **do**  
    Calculate the score according to Eq. 3 and sample a subnetwork.  
    Update the parameter of adapters for adaptation and substitute  $\sigma$  by  $\mathcal{L}_{train}$ .  
    **if**  $t \bmod interval == 0$  **then**  
        Obtain a validation batch  $d_v \leftarrow D_v$ .  
        Evaluate the performance of  $c$  subnetworks and get their reward  $r_i = e^{-\mathcal{L}_{val}}$   
        Update the degree of redundancy  $\mathbf{r}^{(t)}$  according to Eq. 5.  
    **end if**  
**end for**  
Obtain the optimal structure  $\theta_K$  based on  $\mathbf{r}$ .

---

DAS first observes the model redundancy to downstream tasks before skipping the layers of VLP models. In practice, we regard this process as a  $k$ -armed bandit problem, as illustrated in Fig. 2. Firstly, we define the degree of redundancy as  $\mathbf{r} \in \mathbb{R}^n$ , where  $n$  is the number of VLP modules. To correctly estimate the redundancy, we equally initialize  $\mathbf{r}_i = 0$  and update it momentarily.

In each training step, we skip  $m$  modules according to uniform distribution based on  $\mathbf{r}$ , and train the sampled architectures on the downstream data. For  $t$ -th step, the action policy  $\pi_i^{(t)}$  for the  $i$ -th VLP module follows the distribution:

$$\pi_i^{(t)} \sim U(0, \rho(\mathbf{r}_i^{(t)})), \quad (3)$$

where  $U(a, b)$  is the uniform distribution between  $a$  and  $b$ . And  $\rho$  represent the Sigmoid function. We randomly pick a probability from the  $\pi_i^{(t)}$  of each module as the score  $s_i^{(t)}$ . According to the score  $s_i^{(t)}$ , the sampled subnetwork can be defined by

$$G_s = g_0 \circ g_1 \circ \dots \circ g_n, \quad (4)$$

where  $g_i = \begin{cases} \theta_i, & i \in \{j | s_j^{(t)} < s_m^{(t)}\}, \\ \sigma_i, & i \in \{j | s_j^{(t)} \geq s_m^{(t)}\}. \end{cases}$

Here,  $g_i \circ g_{i+1}$  represents the compositional function  $g_i(g_{i+1}(\cdot))$ .  $\theta_i$  denotes the original VL module, and  $\sigma_i$  is the lightweight module like adapter for short-cut connection. And  $s_m^{(t)}$  are the  $m$  largest values in the picked scores. Here, the module with a larger  $\mathbf{r}_i^{(t)}$  is more likely to be skipped during training. Meanwhile, Eq. 4 also help  $\sigma_i$  learn pre-trained knowledge from  $\theta_i$  in a distillation way [55].

Then, DAS observes the redundancy of VLP modules in a reinforcement learning manner, as shown in 2. DAS samples  $c$  candidate network structures and calculates their rewards according to their loss values during validation, *i.e.* reward  $v = e^{-loss}$ . Based on the rewards, the degree of redundancy  $\mathbf{r}$  can be updated by

$$\mathbf{r}_i^{(t+1)} = \mathbf{r}_i^{(t)} + (v_h - \frac{1}{c} \sum_{j=1}^c v_j). \quad (5)$$

Here,  $v_h$  denotes the reward of the sampled subnetwork, where the  $i$ -th module is skipped. When its validation loss is larger than the mean value, it suggests that this skipped module is more redundant. Eq. 5 is conducted at short training intervals to makes sure that most subnetworks can be sufficiently validated via numerous samplings, and the theorem of large numbers can guarantee the optimality of search results. The detailed search procedure of DAS is illustrated in Algorithm. 1. Finally, according to the degree of redundancy  $\mathbf{r}$ , we can select top- $m$  layers to be skipped, thereby reducing the computation complexity of VLP models.

## 4.2 Model Adapting

To reduce the updated parameter scale during adaptation, DAS also introduces lightweight adapters [16, 48] to serve the hidden feature transformations as well as the short-cut connections in

VLP models. Typically, an adapter is constructed by two linear projection layers and an activation function in between:

$$\text{adapter}(\mathbf{x}) = \text{ReLU}(\mathbf{x}\mathbf{W}_{in})\mathbf{W}_{out}. \quad (6)$$

Here,  $\mathbf{W}_{in} \in \mathbb{R}^{d \times h}$  and  $\mathbf{W}_{out} \in \mathbb{R}^{h \times d}$  are two trainable matrices, where  $h \ll d$ . For the  $i$ -th VLP module, the adaptation can be defined by

$$\mathbf{x}_i = \mathbf{x}_{i-1} + \text{VLP}(\mathbf{x}_{i-1}) + \text{adapter}(\mathbf{x}_{i-1}), \quad (7)$$

where  $\mathbf{x}_i$  is the output of the  $i$ -th component. In this manner, DAS can freeze most parameters in the VLP models during adaption, similar to existing PETL methods [48].

Notably, directly removing the redundant modules will make the subsequent layers to receive the hidden features with drastic changes. Meanwhile, we do not expect the fully tuning of the whole model. In this case, we apply the adapter to serve the short-cut connection of the skipped layers:

$$\mathbf{x}_i = \mathbf{x}_{i-1} + \text{adapter}_r(\mathbf{x}_{i-1}). \quad (8)$$

In this way, DAS can not only bridge the gap between feature transformations, but also retain parameter efficiency. Based on the estimated redundancy, DAS skips the redundant modules and finds out the optimal pathway for the downstream task with the helps of adapters, as shown in Fig. 2.

## 5 Experiment

### 5.1 Datasets and Experimental Setup

**Visual Question Answering.** We conduct experiments on VQA2.0 [11]. Instead of answering the question in open-ended natural language, it is converted into a classification task with 3, 129 classes. Following the previous setting [8, 24], the PETL methods and DAS are trained on the train and validation sets of VQA2.0, and we report the *test-dev* results from the online evaluation <sup>2</sup>.

**Natural Language for Visual Reasoning.** The NLVR<sup>2</sup> [46] is a dataset for classifying triplets of two images and a question into two classes. Because its form is different from the setup of VLP models, which has two images in one VL example, we feed these triplet examples to the model following the default setting of ViLT [24] and METER [8]. Under this setting, the paired images and the question are input to the network, respectively. And the classifier predicts the results according to the concatenation of two representations.

**Retrieval Task.** For cross-modal retrieval, we measure the performance on Flickr30K [41] re-splitted by Karpathy *et al.* [23]. We initialize the predictor for similarity measurement from the pre-trained ITM head. During the training, we randomly sample 15 instances as negative samples.

### 5.2 Implementation details

We validate DAS on two deep-fusion based VLP models, which are ViLT [24] and METER [8]. In terms of ViLT, we update the parameters of the additional components, classifier, class token and modal-type embeddings, while the rest are frozen. Following the most conventional setting [13, 48], the width of hidden states in adapters is set to 96. And the hidden dimension of the adapter used for the skip connection is set to 192 to retain a certain capacity. The VLP model is first warmed up for one epoch. In this epoch, the subnetwork is randomly sampled according to the skipped number  $m$ . Then the search runs 2 epochs and the redundancy observation is executed at 10-th step per interval. Finally, the optimal architecture will be trained for another 10 epochs. In terms of METER, we split its fusion layer into two modules, *i.e.* the vision and language ones, which are skipped separately. The hidden dimension of the adapter used as the skip connection is set to 192 for encoders and 288 for fusion layers. The rest settings are the same as ViLT. We conduct all experiments with a single NVIDIA Tesla A100 GPU and the settings not mentioned are the same as ViLT [24] and METER [8].

### 5.3 Experimental results

#### 5.3.1 Quantitative analysis

**Comparison with PETL methods.** We first compare DAS with a bunch of PETL methods [13, 17, 22, 30, 48] on the VLP models, of which results are given in Tab. 1. Here, **the suffix of DAS**

<sup>2</sup><https://eval.ai/web/challenges/challenge-page/830/overview>

Table 1: Comparison between DAS and the PETL methods for METER and ViLT on VQA, NLVR<sup>2</sup> and Flickr30K. The best performance is **bold** and the second best is underlined.

METER									
Method	Updated Parameter	VQA		NLVR <sup>2</sup>		Flickr30K		Avg.	
		test-dev	Additional FLOPs	test-P	Additional FLOPs	IR/TR R@1	Additional FLOPs	Per.	Additional FLOPs
Full Tuning	323.31M	77.43	0.00	83.05	0.00	82.22/94.30	0.00	84.25	0.00
Classifier Only	-	69.93	0.00	73.23	0.00	78.80/89.00	0.00	77.74	0.00
Shallow Prompt [30]	0.30M	68.51	28.71G	65.69	26.84G	74.20/88.60	28.71G	74.25	28.71G
Deep Prompt [22]	1.84M	70.78	6.53G	72.64	5.59G	78.84/89.40	6.53G	77.92	6.53G
LoRA [17]	0.29M	74.00	0.00	78.82	0.00	79.86/92.60	0.00	81.32	0.00
Adapter [48]	5.34M	74.70	1.64G	79.93	1.38G	80.38/91.90	1.64G	81.73	1.64G
Scaled PA [13]	3.59M	<b>75.11</b>	1.12G	80.38	0.66G	80.40/ <b>93.20</b>	1.12G	<b>82.27</b>	1.12G
<b>DAS<sub>4</sub>-Fusion</b>	5.34M	74.80	<b>-11.16G</b>	80.11	<b>-5.13G</b>	80.12/91.80	<b>-11.16G</b>	81.71	<b>-9.15G</b>
<b>DAS<sub>4</sub>-Global</b>	6.23M	<u>75.09</u>	<u>-4.51G</u>	<b>80.69</b>	<u>-3.67G</u>	<b>80.42</b> /91.40	<u>-6.06G</u>	81.90	<u>-4.74G</u>

ViLT									
Method	Updated Parameter	VQA		NLVR <sup>2</sup>		Flickr30K		Avg.	
		test-dev	Additional FLOPs	test-P	Additional FLOPs	IR/TR R@1	Additional FLOPs	Per.	Additional FLOPs
Full Tuning	115.43M	71.26	0.00	76.13	0.00	64.40/83.50	0.00	73.82	0.00
Classifier Only	-	65.75	0.00	66.08	0.00	57.42/78.00	0.00	66.81	0.00
Shallow Prompt [30]	0.15M	66.47	+19.53G	66.47	+19.53G	55.92/74.80	+19.53G	65.92	+19.53G
Deep Prompt [22]	1.84M	69.30	+5.14G	73.34	+5.14G	58.64/79.50	+5.14G	70.20	+5.14G
LoRA [17]	0.15M	68.44	0.00	72.77	0.00	57.44/77.70	0.00	69.09	0.00
Scaled PA [13]	1.80M	70.40	+0.44G	75.13	+0.44G	61.88/79.00	+0.44G	71.60	0.44G
Adapter [48]	3.56M	<b>70.85</b>	+0.86G	<b>75.51</b>	+0.86G	<b>62.68/81.40</b>	+0.86G	<b>72.61</b>	0.86G
<b>DAS<sub>4</sub>-Fusion</b>	5.34M	70.66	-0.80G	75.13	-0.80G	60.66/80.80	-1.03G	71.41	-1.03G

Table 2: Ablation study of different numbers of skipped layers by DAS. “Fusion” denotes the skipping is only for the fusion branch, while “Global” refers to the entire METER.

METER							
Method	Skipped Number	VQA		NLVR <sup>2</sup>		Avg.	
		test-dev	Additional FLOPs	test-P	Additional FLOPs	Per.	Additional FLOPs
Baseline	0	75.28	1.68G	81.28	0.99G	78.28	1.34G
DAS-Fusion	2	74.92	-9.06G	80.07	-2.66G	77.50	-5.86G
	4	74.80	-11.16G	80.11	-4.14G	77.46	-7.65G
	6	74.67	-17.58G	78.16	-9.97G	76.42	-13.78G
	8	73.70	-24.00G	79.30	-11.45G	76.50	-17.72G
DAS-Global	2	75.24	-3.96G	81.37	-2.19G	78.31	-3.08G
	4	75.13	-4.51G	81.34	-3.67G	78.24	-4.09G
	6	75.02	-5.06G	80.04	-4.22G	77.53	-4.64G
	8	74.95	-5.61G	79.61	-8.34G	77.28	-6.97G

ViLT							
Method	Skipped Number	VQA		NLVR <sup>2</sup>		Avg.	
		test-dev	Additional FLOPs	test-P	Additional FLOPs	Per.	Additional FLOPs
Baseline	0	70.13	0.73G	76.26	0.73G	73.20	0.73G
DAS	1	69.28	-1.03G	74.89	-1.03G	72.09	-1.03G
	2	67.64	-2.79	73.00	-2.79G	70.32	-2.79G




Figure 3: The comparison between DAS and random skipping for METER on the VQA2.0.

denotes the number of skipped layers, and “*fusion*” and “*global*” refer to the range of network skipping, *i.e.*, only the fusion branch or the complete model. From Tab. 1, the first observation is that existing PETL methods can largely reduce the amount of updated parameters for VLP models. For instance, the prompt-based methods only require about 1M parameters for two VLP models, nearly 300 times less than full tuning. Meanwhile, their performance gap to fully tuning is also marginal, *e.g.*, Scaled PA [13]. However, we can also see that all these PETL methods cannot reduce the computation of VLP models, and some of them incurs obvious increases in FLOPs, *e.g.* +28.71G by Shallow Prompt [30] on METER. The most computation efficient one is LoRA [17], which applies the re-parameterization technique to merge the additional modules into the VLP model, taking no extra computations. However, its performance obviously lags behind other PETL methods and our DAS. In stark contrast, DAS is the only method that can reduce the computation overhead on the downstream VL tasks, *e.g.*, -11.16G FLOPs by DAS<sub>4</sub>-Fusion on VQA. More importantly, its updated parameter scale is only slightly larger than Adapter and Scaled PA, while the overall performance is still competitive. These results well confirm the effectiveness of DAS towards PCETL.

**Ablation of the number of skipped layers.** In Tab. 2, we report the results of skipping different numbers of layers by DAS. In terms of METER, we can first observe that skipping a few layers has limited impact on its performance, *e.g.*, skipping up to 8 fusion layers only has about 2.2% performance drops, strongly suggesting the redundancy of this VLP model. However, DAS can only reduce about one layer for ViLT without notable performance degradations. To explain, METER is a deep VLP model with two independent encoders to extract the features of image and text, while ViLT

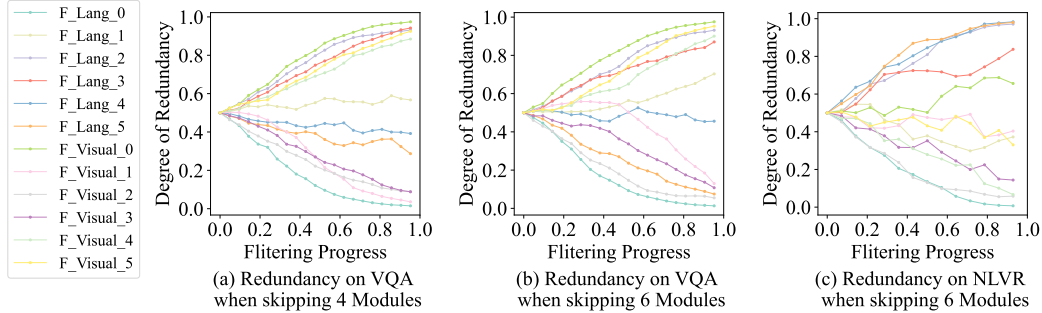


Figure 4: The change of the redundancies of METER’s fusion layers. The horizontal axis shows the progress of training and the vertical axis represents the degree of redundancy based on Eq. 5. F\_Lang and F\_Visual denote the language and visual modules in the fusion branch.

processes the multi-modal information from image pixels and word embeddings. In this case, ViLT is more compact than METER for downstream tasks, and this is also reflected in their parameter scales and performance. In terms of METER, we can also see the difference in optimizing the fusion branch and the entire model, *i.e.* DAS-Fusion and DAS-Global. With the same number of skipped layers, DAS-Fusion can reduce more FLOPs since the length of multi-modal inputs is often larger than the single-modal one. Meanwhile, when evaluating the entire model, DAS often tends to reduce the layers in the language encoders<sup>3</sup>, which also suggests that natural language understanding is often less difficult in the VL tasks. Overall, these results well confirm our motivation about the redundancy of VLP models in downstream VL tasks, especially the ones with independent encoders like METER.

**Reliability of DAS.** Considering that DAS is an RL-based search approach, we also examine its stability and reliability via comparing to random skipping, of which results are given in Tab. 2. It can be seen that DAS is consistently better than random skipping without regard to the number of skipping layers, well confirming its effectiveness. In particular, when the number of skipped layers increases, the performance deviation of random skipping will become much more obvious, *e.g.*  $\pm 1.33$  for skipping 8 layers. Instead, DAS is still more stable and superior than the random solution, which also suggests the importance of correct redundancy estimations. Overall, these results further confirm the effectiveness and reliability of the proposed DAS.

### 5.3.2 Qualitative analysis

To obtain deep insight into DAS, we also visualize its changes of redundancy estimations during search, as depicted in Fig. 4. From these plots, we can observe several patterns of DAS across tasks. The first is that the most redundant layers can quickly emerge during DAS, especially when the number of skipped layers is smaller, *e.g.* Fig. 4.a. Meanwhile, for some uncertain layers, their redundancies can be gradually determined after a short period of oscillation. We can also see that the numbers of the skipped language and visual layers are similar on VQA and NLVR. However, the preference about the skipping layers is still different on the two tasks, see Fig. 4.b and Fig. 4.c.

In Fig. 5, we visualize the subnetworks of METER searched by DAS, which are also the final results of Fig. 4. As discussed above, the preferences of the hidden layers of METER are still different for two tasks. In terms of VQA, DAS tends to discard the language modules at the middle level while skipping the visual ones on the top and bottom of the network. In contrast, the language layers skipped on NLVR are all the top ones. These search results can somewhat reflect the properties of these two tasks. For instance, VQA is a reasoning task, thus it focuses more on the high-level semantics of the input text, while NLVR needs a detailed comparison between two images and one sentence, so the last few language layers may be more redundant to this task.

Overall, Fig. 4 and Fig. 5 not only confirm the feasibility of DAS towards PCETL, but also yields some interesting findings about the downstream adaption of VLP models.

## 6 Limitation

Currently, DAS has two main limitations. First, it still needs to set the number of layers to skip. In our future work, we will introduce the computation or hardware constraints to DAS for the more

<sup>3</sup>The detailed results are given in our Appendix.



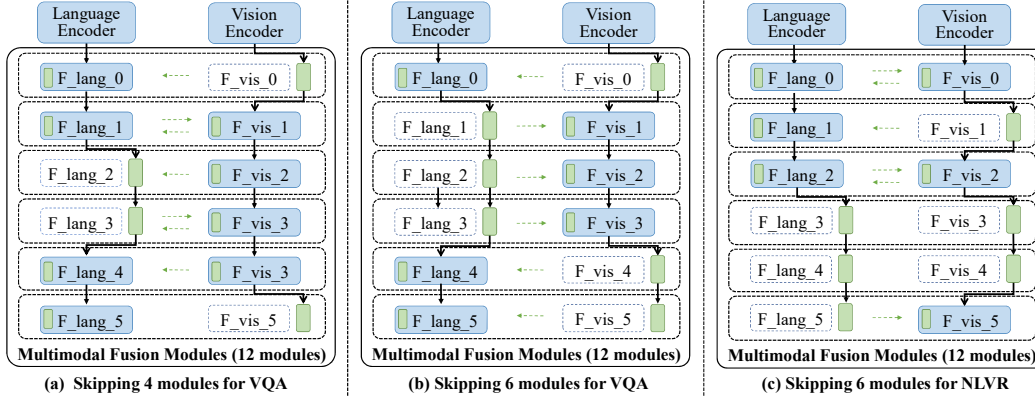


Figure 5: The optimal subnetworks of METER searched by DAS. Here, the green modules are trainable adapters while the blue ones are frozen during adaption.

automatic network skipping. Second, DAS only regards the complete Transformer layers in VLP models as the skipping candidates, limiting its potential in pathway routing. In the future, we will extend the search space with more detailed components, such as MHA and FFN.

## 7 Conclusion

In this paper, we propose a new problem for vision-language pre-trained (VLP) models termed *parameter and computation efficient transfer learning* (PCETL). Existing transfer learning solutions for VLP models can only save the parameter expenditure during downstream task adaption, *e.g.*, the PETL ones, while the excessive computation is still a unconquered problem. To this end, we propose a novel approach called *Dynamic Architecture Skipping* (DAS) towards effective PCETL. DAS can observe the redundancies of VLP modules to downstream tasks via a reinforcement learning based process, and then skip the redundant ones to speed up inference. Meanwhile, DAS also adopts lightweight adapters to serve the hidden feature adaptations and the short-cut connections, thereby reducing the scale of trainable parameters. On two VLP models and three VL tasks, DAS not only shows a great superiority in reducing computation, but is also on par with the PETL methods in terms of parameter overhead and performance.

## References

- [1] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [2] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [3] Jianda Chen, Shangyu Chen, and Sinno Jialin Pan. Storage efficient and dynamic flexible runtime channel pruning via deep reinforcement learning. *Advances in Neural Information Processing Systems (NeurIPS)*, pages 14747–14758, 2020.
- [4] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: universal image-text representation learning. In *European Conference on Computer Vision (ECCV)*, pages 104–120, 2020.
- [5] Leyang Cui, Yu Wu, Jian Liu, Sen Yang, and Yue Zhang. Template-based named entity recognition using BART. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1835–1845, 2021.
- [6] Zihang Dai, Hanxiao Liu, Quoc V. Le, and Mingxing Tan. Coatnet: Marrying convolution and attention for all data sizes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 4171–4186, 2019.
- [8] Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, et al. An empirical study of training end-to-end vision-and-language transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18166–18176, 2022.
- [9] Steven K Esser, Jeffrey L McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. *arXiv preprint arXiv:1902.08153*, 2019.
- [10] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *Computing Research Repository (CoRR)*, 2021.
- [11] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6325–6334, 2017.
- [12] Demi Guo, Alexander M. Rush, and Yoon Kim. Parameter-efficient transfer learning with diff pruning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4884–4896, 2021.
- [13] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *International Conference on Learning Representations (ICLR)*, 2022.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision (ECCV)*, 2016.
- [15] Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. Deberta: Decoding-enhanced bert with disentangled attention. *Computing Research Repository (CoRR)*, 2020.
- [16] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning (ICML)*, pages 2790–2799, 2019.
- [17] 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 *International Conference on Learning Representations (ICLR)*, 2022.
- [18] Gao Huang, Shichen Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Condensenet: An efficient densenet using learned group convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [19] Zhicheng Huang, Zhaoyang Zeng, Yupan Huang, Bei Liu, Dongmei Fu, and Jianlong Fu. Seeing out of the box: End-to-end pre-training for vision-language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12976–12985, 2021.
- [20] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *Computing Research Repository (CoRR)*, 2020.
- [21] Ashraf Islam, Chun-Fu Richard Chen, Rameswar Panda, Leonid Karlinsky, Rogerio Feris, and Richard J Radke. Dynamic distillation network for cross-domain few-shot recognition with unlabeled data. *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3584–3595, 2021.
- [22] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge J. Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision (ECCV)*, pages 709–727, 2022.
- [23] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI / PAMI)*, pages 664–676, 2017.
- [24] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning (ICML)*, pages 5583–5594, 2021.
- [25] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018.

- [26] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3045–3059, 2021.
- [27] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9694–9705, 2021.
- [28] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *Computing Research Repository (CoRR)*, 2019.
- [29] Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. UNIMO: towards unified-modal understanding and generation via cross-modal contrastive learning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2592–2607, 2021.
- [30] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4582–4597, 2021.
- [31] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. GPT understands, too. *Computing Research Repository (CoRR)*, 2021.
- [32] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *Computing Research Repository (CoRR)*, 2019.
- [33] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 13–23, 2019.
- [34] Gen Luo, Minglang Huang, Yiyi Zhou, Xiaoshuai Sun, Guannan Jiang, Zhiyu Wang, and Rongrong Ji. Towards efficient visual adaption via structural re-parameterization. *Computing Research Repository (CoRR)*, 2023.
- [35] Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Yan Wang, Liujuan Cao, Yongjian Wu, Feiyue Huang, and Rongrong Ji. Towards lightweight transformer via group-wise transformation for vision-and-language tasks. *IEEE Transactions on Image Processing (TIP)*, 2022.
- [36] Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1022–1035, 2021.
- [37] Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 565–576, 2021.
- [38] Yuning Mao, Lambert Mathias, Rui Hou, Amjad Almahairi, Hao Ma, Jiawei Han, Scott Yih, and Madian Khabza. Unipelt: A unified framework for parameter-efficient language model tuning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 6253–6264, 2022.
- [39] Fanxu Meng, Hao Cheng, Ke Li, Huixiang Luo, Xiaowei Guo, Guangming Lu, and Xing Sun. Pruning filter in filter. *Advances in Neural Information Processing Systems (NeurIPS)*, pages 17629–17640, 2020.
- [40] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. Language models as knowledge bases? In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2463–2473, 2019.
- [41] Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *International Journal of Computer Vision (IJCV)*, pages 74–93, 2017.
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, pages 8748–8763, 2021.
- [43] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 91–99, 2015.

- [44] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [45] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations (ICLR)*, 2020.
- [46] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 6418–6428, 2019.
- [47] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. LST: ladder side-tuning for parameter and memory efficient transfer learning. *Computing Research Repository (CoRR)*, 2022.
- [48] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. VL-ADAPTER: parameter-efficient transfer learning for vision-and-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5217–5227, 2022.
- [49] Yi-Lin Sung, Varun Nair, and Colin Raffel. Training neural networks with fixed sparse masks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 24193–24205, 2021.
- [50] Hao Tan and Mohit Bansal. LXMERT: learning cross-modality encoder representations from transformers. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5099–5110, 2019.
- [51] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5998–6008, 2017.
- [53] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. SimVLM: Simple visual language model pretraining with weak supervision. In *International Conference on Learning Representations (ICLR)*, 2022.
- [54] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [55] Canwen Xu, Wangchunshu Zhou, Tao Ge, Furu Wei, and Ming Zhou. Bert-of-theseus: Compressing BERT by progressive module replacing. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7859–7869, 2020.
- [56] Jeffrey O. Zhang, Alexander Sax, Amir Zamir, Leonidas J. Guibas, and Jitendra Malik. Side-tuning: A baseline for network adaptation via additive side networks. In *European Conference on Computer Vision (ECCV)*, pages 698–714, 2020.
- [57] Renrui Zhang, Rongyao Fang, Wei Zhang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free clip-adapter for better vision-language modeling. *Computing Research Repository (CoRR)*, 2021.
- [58] Yunshan Zhong, Mingbao Lin, Mengzhao Chen, Ke Li, Yunhang Shen, Fei Chao, Yongjian Wu, and Rongrong Ji. Fine-grained data distribution alignment for post-training quantization. In *European Conference on Computer Vision (ECCV)*, pages 70–86, 2022.
- [59] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16795–16804, 2022.
- [60] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision (IJCV)*, pages 2337–2348, 2022.