

# An Empirical Study of Parameter Efficient Fine-tuning on Vision-Language Pre-train Model

Yuxin Tian

College of Computer Science  
Sichuan University  
Chengdu, China  
cs.yuxintian@outlook.com

Mouxin Yang

College of Computer Science  
Sichuan University  
Chengdu, China  
yangmouxing@gmail.com

Yunfan Li

College of Computer Science  
Sichuan University  
Chengdu, China  
yunfanli.gm@gmail.com

Dayiheng Liu

College of Computer Science  
Sichuan University  
Chengdu, China  
losinuris@gmail.com

Xingzhang Ren

School of Software and Microelectronics  
Peking University  
Beijing, China  
xzhren@pku.edu.cn

Xi Peng

College of Computer Science  
Sichuan University  
Chengdu, China  
pengx.gm@gmail.com

Jiancheng Lv\*

College of Computer Science  
Sichuan University  
Chengdu, China  
lvjiancheng@scu.edu.cn

**Abstract**—Recent studies applied Parameter Efficient Fine-Tuning techniques (PEFTs) to efficiently narrow the performance gap between pre-training and downstream. There are two important factors for various PEFTs, namely, the accessible data size and fine-tunable parameter size. A natural expectation for PEFTs is that the performance of various PEFTs is positively related to the data size and fine-tunable parameter size. However, according to the evaluation of five PEFTs on two downstream vision-language (VL) tasks, we find that such an intuition holds only if the downstream data and task are not consistent with pre-training. For downstream fine-tuning consistent with pre-training, data size no longer affects the performance, while the influence of fine-tunable parameter size is not monotonous. We believe such an observation could guide the choice of training strategy for various PEFTs.

**Index Terms**—Parameter Efficient Fine-tuning, Vision-Language Pre-training

## I. INTRODUCTION

Vision-language pre-training (VLP) has emerged as a fundamental paradigm to boost the performance of downstream VL tasks. Most existing works boost the performance of the pre-trained model by designing novel pre-training tasks [1], increasing the size of both pre-training dataset [2] and the model parameters [3], [4]. Although VLP has shown promising zero-shot performance on the downstream tasks, fine-tuning still plays an indispensable role in narrowing the gap between pre-training and downstream domains.

Taking the fine-tuning cost into consideration, numerous Parameter Efficient Fine-Tuning (PEFT) methods have been proposed to adapt the VLP models to downstream tasks in an efficient manner. More specifically, PEFT methods only fine-tune a few parameters while freezing the most pre-trained parameters. In this study, we mainly focus on the PEFT methods that insert additional parameters into different positions of the VLP models, which we refer to as the exogenous PEFT.

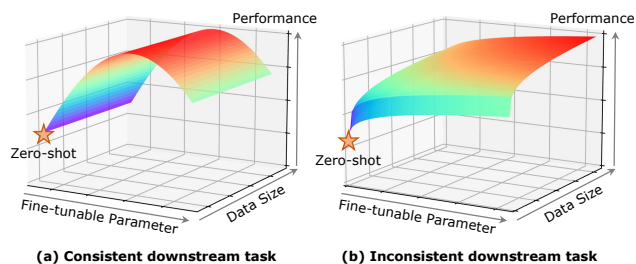


Fig. 1. The performance is only affected by the size of fine-tunable parameters when the downstream task and data are consistent with pre-training. Otherwise, the performance is positively related to data and parameter size.

For clarity, we divide the exogenous PEFT methods into *embedding composition* ones (e.g., prompt-tuning [5] and prefix-tuning [6]) and *layer composition* ones (e.g., Adapter [7], [8] and LoRA [9]), according to the inserted position. Almost all existing PEFT methods seek to achieve competitive performance compared to the full fine-tuning counterparts while embracing high training efficiency. Intuitively, one may expect that the performance of various PEFTs is positively related to the accessible data size and fine-tunable parameter size.

To evaluate such an expectation, we conduct an empirical analysis with five PEFTs on two downstream VL tasks. Specifically, we first offer a novel unified view of the investigated prompt-tuning, prefix-tuning, LoRA, serial adapter-tuning, and parallel adapter-tuning. Then, considering the differences between the pre-training and downstream fine-tuning, the accessible downstream data size, and the size of the fine-tunable parameters, we conduct a series of experimental evaluations on two widely-used VL datasets, *i.e.*, MSCOCO Caption and VQA<sub>v2</sub>. It is worth noting that the image caption task is usually adopted as a pre-training task in the VLP,

\* Corresponding Author

but not with VQA. Therefore, the image caption and the corresponding data in the downstream tasks could be regarded as consistent with the VLP, while the VQA one is inconsistent.

The contributions of this study mainly lie in the new empirical observations. As illustrated by Fig. 1, if the downstream task and dataset are not consistent with pre-training, the data size and the fine-tunable parameter size are positively related to the performance. If consistent, the data size no longer affects the performance of various PEFTs while the influence of fine-tunable parameter size is not monotonous. We believe that such observations would guide the training strategy design of various PEFTs. Furthermore, our experimental results also reveal an additional phenomenon: Considering the training efficiency and performance, layer composition (*e.g.*, LoRA) could be a better choice for the downstream adaptation of the VLP model.

## II. RELATED WORK

### A. PEFT for VLP model

A growing body of research has been devoted to finding parameter-efficient alternatives to adapt large-scale VLP models to downstream tasks and to reduce the cost of various aspects such as memory and storage. The representative works for the first line are prompt-tuning [5] and prefix-tuning [6]. The early prompt-based works employ prompt-tuning [5] to manage the few-shot visual tasks [10]. The most recent researches expand the application of prefix-tuning into the VLP [11] model and achieve comparable performance with full finetuning at low cost. The other line of work is based on the adapter [7] and LoRA [9]. Most recent works [12] resort to the adapter layer to extend the multimodal abilities of the generative large-scale language model. Besides, LoRA [9] are also used to adapt text-only LLM for multimodal tasks [13]. Although all of these methods have demonstrated their effectiveness, there is currently no analysis to investigate how the data size and fine-tunable parameter size affect various PEFTs.

### B. Empirical PEFT analysis

Recent studies have conducted numerous experiments to examine the possible factors that influence the performance and robustness of PEFT. Chen et.al [14] reveal that full fine-tuning cannot be entirely replaced by PEFT approaches currently in NLP since it cannot attain superior performance to full fine-tuning when given adequate fine-tuning budget and data size. Sung et.al [15] finds that the vanilla adapter [7] achieved the best VL task performance among its variants. Regarding robustness, Chen et.al [16] deem that neither full fine-tuning nor PEFT approaches consistently provide robustness for data corruption. Differently, this study investigates an inevitable problem *Are the data size and fine-tunable parameter size positively related to the VL performance of various PEFTs?*

## III. A UNIFIED VIEW FOR PEFT

In this section, we briefly review the investigated PEFTs from a unified perspective view. The exogenous parameter fine-tuning efficiently trains the VLP model by incorporating

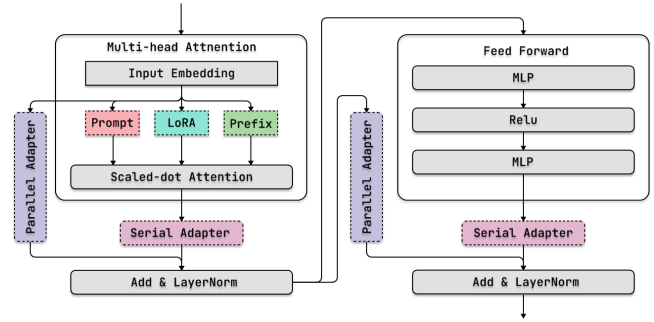


Fig. 2. Unified view of evaluated PEFT methods within a transformer block.

additional trainable parameters into the input, *e.g.*, **prompt-tuning** and **prefix-tuning**, or intermediate layers of the model, *e.g.*, **LoRA**, **serial adapter-tuning**, and **parallel adapter-tuning**. To be specific, given a pre-trained VLP model  $F$  with parameters  $\Theta$ , PEFT resorts to fine-tuning the additional parameters  $\Phi$  to adapt the pre-trained model  $F$  from pre-training task to downstream task, *e.g.*, VQA. In formal, with the input embedding  $\mathbf{x} \in \mathbb{R}^{n \times d}$  for a transformer block and a ground truth  $\mathbf{y}$ , the objective of such PEFT could be unified as follow:

$$\begin{aligned} \arg \min_{\Phi} \mathcal{L}(F(\mathbf{x}; \Theta, \Phi), \mathbf{y}) \\ = \mathcal{L}(f_n \circ \dots \circ f_1(\mathbf{x}; \Theta_1, \Phi_1), \mathbf{y}), \end{aligned} \quad (1)$$

where  $\mathcal{L}$  and  $f_i$  denotes the loss function of downstream task and the  $i$ -th transformer block, respectively. According to the position of the extra parameters introduced by PEFT, we categorize the tested prompt-tuning and prefix-tuning as *embedding composition* and serial adapter-tuning, parallel adapter-tuning, and LoRA as *layer composition*. We will elaborate on them in the following.

### A. Full finetuning

The most straightforward way to adapt the pre-trained VLP models to downstream tasks is full finetuning. Yet, directly updating the full set of  $\Theta$  costs vast computational resources, particularly as model size continues to increase.

### B. Embedding composition

Prior study [5] proposed ‘‘continuous’’ prompt which concatenates trainable parameters  $\Phi$  to the embedding of the transformer layers and hence we entitle them as embedding composition. For example, prompt-tuning [5] prepends the model input embedding  $\mathbf{x} \in \mathbb{R}^{n \times d}$  with a learnable embedding  $\Phi := \mathbf{p} \in \mathbb{R}^{p \times d}$  optimized directly through gradient descent, where  $n/p$  denotes the length of the input/prompt embedding and  $d$  is the dimension of them. According to Eq.1, the objective of prompt-tuning could be rewritten following:

$$\arg \min_{\mathbf{p}} \mathcal{L}(F([\mathbf{p} : \mathbf{x}]), \mathbf{y}), \quad (2)$$

where  $[\cdot : \cdot]$  indicates the concatenation operation. Instead of adding parameters to the input embedding of the model,

prefix-tuning [6] prepends trainable tokens, *i.e.*,  $\Phi := \{\mathbf{p}_i\}^n$ , to the hidden states of all the transformer blocks. Formally, the corresponding objective is as follows:

$$\arg \min_{\{\mathbf{p}_i\}_1^n} \mathcal{L}(f_n^p \circ f_{n-1}^p \circ \dots \circ f_1^p(\mathbf{x}), \mathbf{y}), \quad (3)$$

where  $f_i^p(\mathbf{x}) := f_i([\mathbf{p}_i : \mathbf{x}])$ . Usually, one could use an additional MLP encoder  $P$  to encode  $\{\mathbf{p}_i\}^n$  for better training stability.

### C. Layer composition

Differing from embedding composition, layer composition inserts an extra sub-network with few trainable parameters into transformer blocks. For example, [7] firstly proposes a vanilla adapter layer that consists of two fully connected layers. Following effort [8] extends it and proposes a parallel variant instead of inserting the adapter layer sequentially. Formally, inserting two serial/parallel adapter layers  $\{A_{i,1}^{\{s,p\}}, A_{i,2}^{\{s,p\}}\}$  to the  $i$ -th transformer block as additional parameters  $\Phi := \{A_{i,1}^{\{s,p\}}, A_{i,2}^{\{s,p\}}\}^n$ , one could have the following objective of adapter-tuning:

$$\arg \min_{\{A_{i,1}, A_{i,2}\}^n} \mathcal{L}(f_n^a \circ f_{n-1}^a \circ \dots \circ f_1^a(\mathbf{x}), \mathbf{y}), \quad (4)$$

where  $f_i^a$  could be the  $i$ -th transformer block with either sequential adapter layer  $f_i^{sa}$  or parallel adapter layer  $f_i^{pa}$ :

$$f_i^{sa}(\mathbf{x}) := AL \circ A_{i,2}^s \circ FFN \circ AL \circ A_{i,1}^s \circ MHA(\mathbf{x}, \mathbf{x}), \quad (5)$$

$$f_i^{pa}(\mathbf{x}) := AL(A_{i,2}^p(h_{med}) + FFN(h_{med})) \quad (6)$$

$$h_{med} = AL(A_{i,1}^p(\mathbf{x}) + MHA(\mathbf{x}, \mathbf{x})) \quad (7)$$

$$A_i^{\{s,p\}}(\mathbf{x}) := ReLU(\mathbf{x}\mathbf{W}_{down})\mathbf{W}_{up} + \mathbf{x} \quad (8)$$

where  $AL$ ,  $FFN$ , and  $MHA$  denote the skip-connection and layernorm, feed-forward, and multi-head attention within a transformer block, respectively. Further, the matrices, *i.e.*,  $\mathbf{W}_{down} \in \mathbb{R}^{d_h \times l}$  and  $\mathbf{W}_{up} \in \mathbb{R}^{l \times d_h}$ , are the trainable parameters of a adapter layer, where  $d_h$  denotes the dimension of hidden state.

In addition to the adapter-tuning, LoRA [9] inserts reparametrized fully connected layers into the self-attention layers, which is inspired by the intrinsic dimensionality [17]. In detail, LoRA utilizes the low-rank decomposition matrices [17] to reparameterize the additional MLP layers. In formal, given the number of head  $k$ , the hidden state of  $i$ -th attention head  $\mathbf{h}_i$ , and the trainable parameters, *i.e.*,  $\Phi := \{\Phi_{i,1}, \Phi_{i,2}\}^n$  which are plugged into the  $MHA$  layers, one could have the following objective for LoRA:

$$\arg \min_{\{\Phi_{i,1}, \Phi_{i,2}\}^n} \mathcal{L}(f_n^l \circ f_{n-1}^l \circ \dots \circ f_1^l(\mathbf{x}), \mathbf{y}). \quad (9)$$

The transformer block with LoRA layer could be as follows:

$$f_i^l(\mathbf{x}) := AL \circ FFN \circ AL \circ MHA_i^l(\mathbf{x}, \mathbf{x}), \quad (10)$$

$$MHA_i^l(\mathbf{g}, \mathbf{x}) = [\mathbf{h}_i^1 : \mathbf{h}_i^2 : \dots : \mathbf{h}_i^k] \mathbf{W}_o, \quad (11)$$

$$\mathbf{h}_i^j = Attn(\mathbf{x}(\mathbf{W}_q^j + \Phi_{i,1}), \quad (12)$$

$$\mathbf{g}(\mathbf{W}_k^j + \Phi_{i,2}), \mathbf{g}\mathbf{W}_v^j),$$

$$\Phi_{i,m} = \mathbf{B}\mathbf{A}, \quad (13)$$

where the parameters of multi-head attention are

$$\mathbf{W}_o \in \mathbb{R}^{d \times d}, \mathbf{W}_q^j, \mathbf{W}_k^j, \mathbf{W}_v^j \in \mathbb{R}^{d \times \frac{d}{k}}, \quad (14)$$

and the low-rank trainable matrices for the LoRA layer are

$$\mathbf{B} \in \mathbb{R}^{d \times r}, \mathbf{A} \in \mathbb{R}^{r \times \frac{d}{k}}, r \ll \frac{d}{k}. \quad (15)$$

$MHA$  functions as self-attention, such that  $\mathbf{g} = \mathbf{x}$ . Otherwise, it operates as cross-attention.

## IV. EXPERIMENTAL SETUP

### A. Base model

We adopt the recently-proposed VLP model, namely, mPLUG [18], as our base model, which employs the discriminative-generative pre-training tasks [3], [4]. In brief, mPLUG adopts the pre-trained CLIP-ViT [2] (ViT-B/16) as visual encoder and two  $BERT_{base}$  [19] models for textual-visual feature fusing and text decoding (See supplementary material for more details).

### B. Dataset and task

We investigate two widely-used and distinct VL downstream tasks in this study, *i.e.*, visual question answering on VQAv2 [20] and image captioning on MSCOCO Caption [21] (See supplementary material for dataset statistics). The VQA is a multi-modal image understanding task that requires the VLP model to answer the textual questions referring to the corresponding image. Following [3], [4], we treat it as an answer generation without constraints for better generality. Image captioning is a multi-modal text generation task that asks for a VLP model to generate an accurate and fluent caption for a given image.

### C. Implementation details

Both for captioning and VQA, the visual encoder (CLIP-ViT-B/16) takes the resized  $256 \times 256$  image as input. The text encoder takes a caption prompt  $\{a \text{ picture of}\}$  and a question  $\{Question: \{\#question\} Answer: \}$  as the input for captioning and VQA, respectively. Then the caption and answer would be generated by the following text decoder. During training, we used the AdamW optimizer with a weight decay of 0.05. The learning rate is warmed up to the highest learning rate in the first epoch, and decayed to the lowest learning rate following a cosine schedule. All the experiments are optimized by cross-entropy loss and conducted on four V100-32G-SMX2 GPUs with a total batch size of 256 (See supplementary material for details). We report the best performance metric amongst the whole training. To be specific, the CIDEr [22] on Karpathy-test split is used for image captioning, whilst the overall accuracy on VQAv2 validation split is reported <sup>1</sup> by following [3], [18].

<sup>1</sup>The submission on the test set is limited: <https://eval.ai/web/challenges/challenge-page/830/phases>

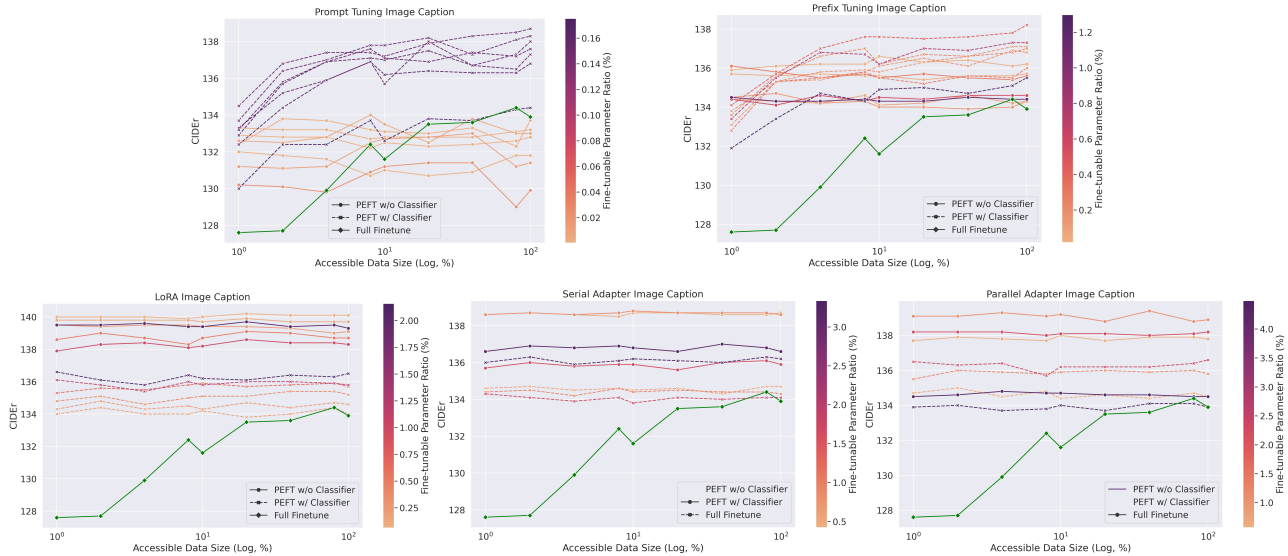


Fig. 3. **Layer composition PEFTs achieve better performance than embedding composition on MSCOCO Caption.** Both layer and embedding composition PEFTs could achieve comparable performance with full fine-tuning. The performance of the tested PEFTs is regardless of the accessible data size. Increasing the size of fine-tunable parameters by simultaneously fine-tuning the final classifier only improves the performance of prompt-tuning, and hurts that of the layer composition.

#### D. Evaluation protocol

We conduct experiments on two VL downstream tasks with five PEFTs, *i.e.*, prompt-tuning, prefix-tuning, LoRA, serial adapter-tuning, and parallel adapter-tuning, to investigate *how the accessible data size and the fine-tunable parameter size affects the performance of them*. Firstly, the available data for different tasks varies, so we employ a random sampling to select diverse proportional Image-Caption and QA pairs from MSCOCO Caption and VQAv2 datasets for training. Second, the available hardware constraints on the available fine-tunable parameters of various PEFT. To explore *how the size of the fine-tunable parameters influence the performance of different PEFTs*, we modulate the length of the prompt, the length of the prefix, the rank of LoRA, as well as the hidden size of both the serial adapter and the parallel adapter. (See supplementary material for the details).

### V. RESULTS AND ANALYSIS

#### A. PEFT with various data sizes

In this section, we investigate how the accessible data size influences the performance of various PEFTs. According to Fig. 3, one could find that the performance of the tested PEFTs is regardless of different accessible data sizes of MSCOCO Caption, which is different from our intuition. However, from Fig. 4, the performance of all the tested PEFTs steadily increases when the accessible training data size of VQAv2 increases. This could be imputed to that captioning on MSCOCO Caption is consistent with pre-training while VQAv2 is not. To be specific, vision-language pre-training usually utilizes a large-scale of image-text pairs [3], [4], [18] and the downstream data may inevitably be used for pre-training, *e.g.*, MSCOCO Caption.

#### B. PEFT with various parameter sizes

In this section, we investigate how the fine-tunable parameter size affects the performance of various PEFTs. From Fig. 4 and 3, the fine-tunable parameter size does affect the performance of various PEFTs. To be specific, the fine-tunable parameter size is positively related to the performance when one adopts the VL downstream task and data that are different from pre-training, *e.g.*, VQA. Such a phenomenon is consistent with our intuition. If we increase the fine-tunable parameter size by simultaneously fine-tuning the final classifier with various PEFTs, the performance on VQAv2 can be further improved and they could even achieve superior performance than full fine-tuning.

However, when the VL downstream task and data are consistent with the pre-training, *i.e.*, captioning on MSCOCO Caption, there exists an optimal fine-tunable parameter choice, which is different from our intuition. Additionally, increasing the fine-tunable parameter size by simultaneously fine-tuning the final classifier could not always improve the performance on MSCOCO Caption. We deem such a phenomenon also attributes to the consistency between downstream fine-tuning and pre-training. When the downstream fine-tuning is consistent with pre-training, too many fine-tunable parameters could lead to over-fitting. (See supplementary material for further discussion.)

#### C. Further analysis

Besides, our experimental results also provide a comparison of various PEFTs. Fig. 3 indicates that the tested PEFTs always outperform the full fine-tuning except for the prompt-tuning on MSCOCO Caption. Such phenomenon may be attributed to the fact that prompt-tuning possesses less than 0.16% fine-tunable parameters which is less than others. On the other






Fig. 4. **Layer composition PEFTs achieve better composition on VQA2.** Empirically, Layer composition PEFTs could achieve comparable performance with full fine-tuning, while embedding composition PEFTs cannot. The performance of the tested PEFTs is positively correlated to the accessible training data and fine-tunable parameters. Additionally, simultaneously fine-tuning the final classifier of the model could further boost the performance and even achieve superior performance than full fine-tuning.

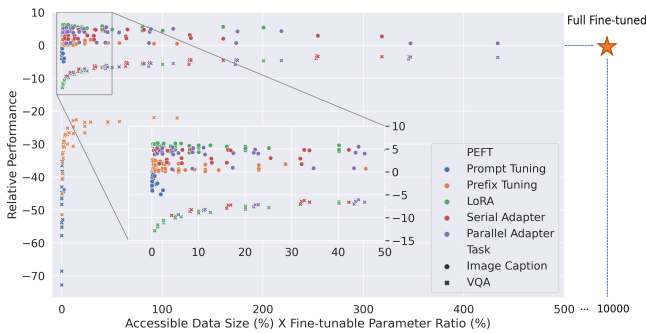


Fig. 5. The comparison of various PEFTs.

hand, the embedding composition PEFTs lag far from full fine-tuning, while the layer embedding composition PEFTs achieve comparable performance to the full fine-tuning, illustrated by Fig. 4. Eventually, by considering the product of data size and fine-tunable parameter size as computation cost, one could easily find which PEFT is better for the tested VLP model. According to Fig. 5, the layer composition offers superior training efficiency and performance on the two downstream VL tasks.

## VI. CONCLUSION

In this paper, we conduct a comprehensive study on five PEFTs on two VL downstream datasets to investigate: *how accessible data size and fine-tunable parameter size affect the performance of various PEFTs.* Referring to our experiments, we find that if the downstream data and task are not consistent with pre-training, increasing the fine-tunable parameter size or accessible data size benefits the performance of PEFT as expected. Nevertheless, if consistent, the data size does

not affect the performance and the fine-tunable parameter holds an optimal size choice. Such a phenomenon could guide the training strategy design of various PEFTs. We leave the exploration of the influence of the base model, more downstream tasks, and data to future work.

## ACKNOWLEDGMENT

## REFERENCES

## REFERENCES

- [1] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang, "OFA: Unifying Architectures, Tasks, and Modalities Through a Simple Sequence-to-Sequence Learning Framework," in *Proc. of ICML, 2022*, pp. 23318–23340.
- [2] 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 *Proc. of ICML, 2021*, pp. 8748–8763.
- [3] 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 *Proc. of NeurIPS, 2021*, pp. 9694–9705.
- [4] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *Proc. of ICML, 2022*, pp. 12888–12900.
- [5] Brian Lester, Rami Al-Rfou, and Noah Constant, "The power of scale for parameter-efficient prompt tuning," *arXiv preprint arXiv:2104.08691*, 2021.
- [6] Xiang Lisa Li and Percy Liang, "Prefix-Tuning: Optimizing Continuous Prompts for Generation," in *Proc. of ACL, 2021*, pp. 4582–4597.
- [7] 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 *Proc. of ICML, 2019*, pp. 2790–2799.
- [8] Junxian He, Chungting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig, "Towards a Unified View of Parameter-Efficient Transfer Learning," in *Proc. of ICLR, 2022*.

- [9] 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 *Proc. of ICLR*, 2022.
- [10] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu, “Learning to Prompt for Vision-Language Models,” *Int. J. Comput. Vis.*, pp. 2337–2348, 2022.
- [11] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan, “Maple: Multi-modal prompt learning,” in *Proc. of CVPR*, 2023, pp. 19113–19122.
- [12] Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao, “LLaMA-Adapter: Efficient Fine-tuning of Language Models with Zero-init Attention,” *CoRR*, vol. abs/2303.16199, 2023.
- [13] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qian Qi, Ji Zhang, and Fei Huang, “mPLUG-Owl: Modularization Empowers Large Language Models with Multimodality,” *CoRR*, vol. abs/2304.14178, 2023.
- [14] Guanzheng Chen, Fangyu Liu, Zaiqiao Meng, and Shangsong Liang, “Revisiting Parameter-Efficient Tuning: Are We Really There Yet?,” in *Proc. of EMNLP*, 2022, pp. 2612–2626.
- [15] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal, “VL-ADAPTER: Parameter-Efficient Transfer Learning for Vision-and-Language Tasks,” in *Proc. of CVPR*, 2022, pp. 5217–5227.
- [16] Shuo Chen, Jindong Gu, Zhen Han, Yunpu Ma, Philip H. S. Torr, and Volker Tresp, “Benchmarking Robustness of Adaptation Methods on Pre-trained Vision-Language Models,” *CoRR*, vol. abs/2306.02080, 2023.
- [17] Armen Aghajanyan, Sonal Gupta, and Luke Zettlemoyer, “Intrinsic Dimensionality Explains the Effectiveness of Language Model Fine-Tuning,” in *Proc. of ACL*, 2021, pp. 7319–7328.
- [18] Chenliang Li, Haiyang Xu, Junfeng Tian, Wei Wang, Ming Yan, Bin Bi, Jiabo Ye, He Chen, Guohai Xu, Zheng Cao, Ji Zhang, Songfang Huang, Fei Huang, Jingren Zhou, and Luo Si, “mPLUG: Effective and Efficient Vision-Language Learning by Cross-modal Skip-connections,” in *Proc. of EMNLP*, 2022, pp. 7241–7259.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *Proc. of NAACL*, 2019, pp. 4171–4186.
- [20] Yash Goyal, Tejas Khot, Aishwarya Agrawal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh, “Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering,” *Int. J. Comput. Vis.*, pp. 398–414, 2019.
- [21] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick, “Microsoft COCO Captions: Data Collection and Evaluation Server,” *CoRR*, vol. abs/1504.00325, 2015.
- [22] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh, “CIDEr: Consensus-based image description evaluation,” in *Proc. of CVPR*, 2015, pp. 4566–4575.
- [23] Bin Bi, Chenliang Li, Chen Wu, Ming Yan, Wei Wang, Songfang Huang, Fei Huang, and Luo Si, “PALM: Pre-training an Autoencoding&Autoregressive Language Model for Context-conditioned Generation,” in *Proc. of EMNLP*, 2020, pp. 8681–8691.

## APPENDIX A EXPERIMENTAL SETUP

### A. Data Details

We investigate two widely-used and distinct VL downstreaming tasks in this study, *i.e.*, visual question answering on VQAv2 [20] and image captioning on MSCOCO Caption [21]. The dataset statistics are shown in Table I.

### B. Base model

We utilize the recent VLP model, namely, mPLUG [18], as our base model, which follows the discriminative-generative pre-train tasks [3], [4]. In general, mPLUG resorts to three discriminative tasks, *i.e.*, Image-Text Contrastive Learning [3],

Image-Text Matching [3], and Masked Language Modeling [19], for multi-modal representation alignment and understanding, and one multi-modal text generation task, *i.e.*, PrefixLM [23] for multi-modal text-generation, respectively. In this study, all the experiments are conducted with mPLUG-base model, which adopts one pre-trained with CLIP-ViT [2] (ViT-B/16) as visual encoder and two *BERT*<sub>base</sub> [19] models for textual and visual feature fusing and text decoding.

### C. Implementation details

Both for captioning and VQA, the visual encoder (CLIP-ViT-B/16) takes the resized  $256 \times 256$  image as input. The text encoder takes a caption prompt  $\{a\ picture\ of\}$  and a question  $\{Question:\ \{\#question\}\ Answer:\ \}$  as the input for captioning and VQA, respectively. Then the caption and answer would be generated by the following text decoder. For the image captioning on MSCOCO Caption, we adopt iteration-based training instead of epoch-based training<sup>2</sup> since such data is widely used [3], [4], [18] and has been learned in the pre-training. In contrast, we used epoch-based training for VQA on VQAv2 for it is not introduced in the pre-training. During training, we adopt the AdamW optimizer with a weight decay of 0.05. The learning rate is warmed up to the highest learning rate in the first epoch, and decayed to the lowest learning rate following a cosine schedule. All the experiments are optimized by cross-entropy loss and conducted on four V100-32G-SMX2 GPUs with a total batch size of 256.

## APPENDIX B

### PEFT EASILY OVER-FITTING ON MSCOCO CAPTION

We also provide the validation curve of the five PEFTs on MSCOCO Caption, illustrated from Fig 6 to Fig. 10. One could also that since the captioning on MSCOCO Caption is consistent with the pre-training, scaling up the fine-tunable parameter size would not monotonously improve the performance of PEFTs. We argue that if the downstream task is consistent with pre-training, too many fine-tunable parameters would lead to over-fitting.

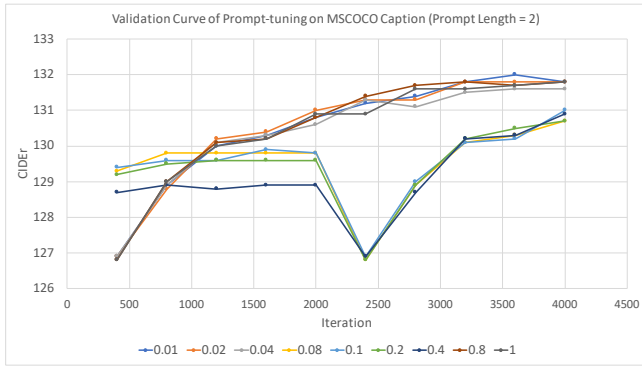
TABLE I  
DATASET STATISTICS

Split	VQAv2		MSCOCO Caption	
	Images	QA Pairs	Images	Captions
Train	113.2K	605.1K	113.2K	566.8K
Val	5.0K	26.7K	5.0K	5.0K
Test	5.0K	26.3K	5.0K	5.0K


<sup>2</sup>Epoch denotes how many times the model sees the complete dataset.

TABLE II  
IMPLEMENTATION DETAILS FOR THE PARAMETER SIZE.

Dataset	Method	Data Ratio	Highest LR	LR Scheduling	Epochs/Steps	Modulated Hyper-paramter
MSCOCO	Prompt-tuning	1%, 2%, 4%, 8%, 10%, 20%,40%, 80%, 100%	3e-4	Warmup from 0 to highest LR at the first 400 steps and cosine decay to 0	4000 Steps	The length of prompt: 2,4,8,16,32,64,128
	Prefix-tuning	1%, 2%, 4%, 8%, 10%, 20%,40%, 80%, 100%	3e-4	Warmup from 0 to highest LR at the first 400 steps and cosine decay to 0	4000 Steps	The length of prefix: 2,4,8,16,32,64,128
	LoRA	1%, 2%, 4%, 8%, 10%, 20%,40%, 80%, 100%	5e-6	Warmup from 0 to highest LR at the first 400 steps and cosine decay to 0	4000 Steps	The rank of LoRA: 2,4,8,16,32
	Serial Adapter	1%, 2%, 4%, 8%, 10%, 20%,40%, 80%, 100%	1e-6	Warmup from 0 to highest LR at the first 400 steps and cosine decay to 0	4000 Steps	The hidden size of serial adapter: 16,32,64,128
VQAV2	Parallel Adapter	1%, 2%, 4%, 8%, 10%, 20%,40%, 80%, 100%	1e-6	Warmup from 0 to highest LR at the first 400 steps and cosine decay to 0	4000 Steps	The hidden size of parallel adapter: 16,32,64,128
	Prompt-tuning	10%, 20%,40%, 80%, 100%	3e-4	Warmup from 1e-6 to highest LR at the first 800*data_ratio steps and cosine decay to 1e-6	10 Epochs	The length of prompt: 2,4,8,16,32,64,128
	Prefix-tuning	10%, 20%,40%, 80%, 100%	3e-4	Warmup from 1e-6 to highest LR at the first 800*data_ratio steps and cosine decay to 1e-6	10 Epochs	The length of prefix: 2,4,8,16,32,64,128
	LoRA	10%, 20%,40%, 80%, 100%	3e-4	Warmup from 1e-6 to highest LR at the first 800*data_ratio steps and cosine decay to 1e-6	10 Epochs	The rank of LoRA: 2,4,8,16,32
VQAV2	Serial Adapter	10%, 20%,40%, 80%, 100%	3e-4	Warmup from 1e-6 to highest LR at the first 800*data_ratio steps and cosine decay to 1e-6	10 Epochs	The hidden size of serial adapter 16,32,64,128,256,512
	Parallel Adapter	10%, 20%,40%, 80%, 100%	3e-4	Warmup from 1e-6 to highest LR at the first 800*data_ratio steps and cosine decay to 1e-6	10 Epochs	The hidden size of parallel adapter 16,32,64,128,256,512

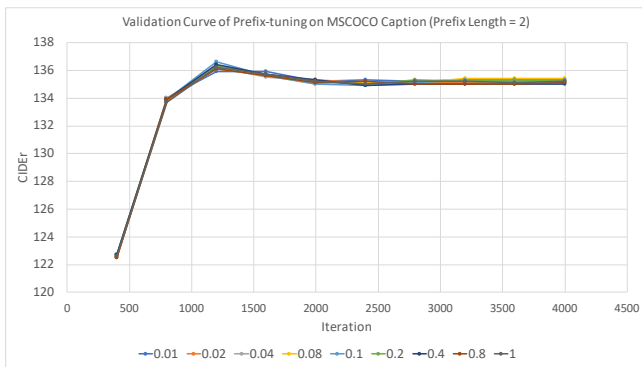


(a) Prompt-length=2.




(b) Prompt-length=128.

Fig. 6. Validation curve of prompt-tuning on MSCOCO Caption.




(a) Prefix-length=2.




(b) Prefix-length=128.

Fig. 7. Validation curve of prefix-tuning on MSCOCO Caption.






(a) Rank=2.




(b) Rank=128.

Fig. 8. Validation curve of LoRA on MSCOCO Caption.

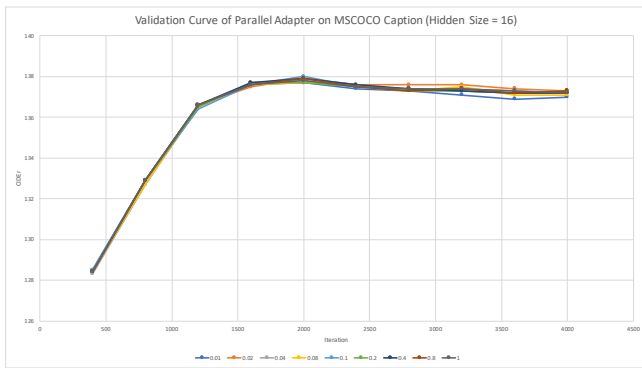


(a) Hidden-size=16.




(b) Hidden-size=128.

Fig. 9. Validation curve of serial adapter on MSCOCO Caption.



(a) Hidden-size=16.



(b) Hidden-size=128.

Fig. 10. Validation curve of parallel adapter on MSCOCO Caption.