

Pruning All-Rounder: Rethinking and Improving Inference Efficiency for Large Vision Language Models

Wei Suo¹, Ji Ma¹, Mengyang Sun¹, Lin Yuanbo Wu², Peng Wang^{1,‡}, Yanning Zhang¹

¹Northwestern Polytechnical University, China.

²Swansea University, United Kingdom.

{suowei1994, maji, sunmenmian}@mail.nwpu.edu.cn {peng.wang, ynzhang}@nwpu.edu.cn

Abstract

Although Large Vision-Language Models (LVLMs) have achieved impressive results, their high computational costs pose a significant barrier to wide application. To enhance inference efficiency, most existing approaches can be categorized as parameter-dependent or token-dependent strategies to reduce computational demands. However, parameter-dependent methods require retraining LVLMs to recover performance while token-dependent strategies struggle to consistently select the most relevant tokens. In this paper, we systematically analyze the above challenges and provide a series of valuable insights for inference acceleration. Based on these findings, we propose a novel framework, the Pruning All-Rounder (PAR). Different from previous works, PAR develops a meta-router to adaptively organize pruning flows across both tokens and layers. With a self-supervised learning manner, our method achieves a superior balance between performance and efficiency. Notably, PAR is highly flexible, offering multiple pruning versions to address a range of acceleration scenarios. The code for this work is publicly available at <https://github.com/ASGO-MM/Pruning-All-Rounder>.

1. Introduction

Substantial progress has been made in the development of Large Vision-Language Models (LVLMs) in recent years [3, 9, 25, 26, 30, 40–42]. Despite being able to handle various Vision-Language (VL) tasks, the applications of these models become restricted due to the increasing computational costs. In fact, the significant computational burden of LVLMs originates from billions of model parameters and lengthy input sequences. Thus, there are two paradigms to address such a problem: parameter-dependent and token-dependent compression strategies.

[‡]Corresponding authors.

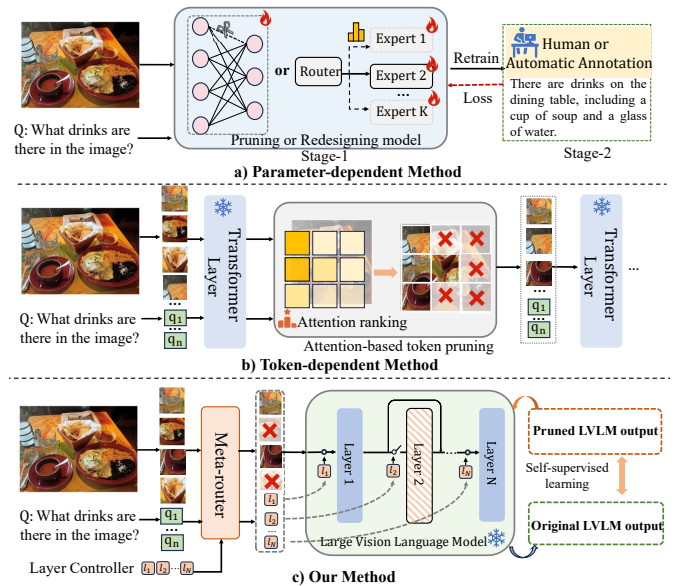


Figure 1. Comparison of different model compression methods. (a) The parameter-dependent strategy requires an additional fine-tuning process to bridge the performance gap. (b) The token-dependent pruning method accelerates inference by reducing the input length. (c) Our method focuses on jointly compressing the inputs and parameters in a self-supervised manner.

As shown in Fig.1(a), the parameter-dependent acceleration approaches mostly follow an inherent two-stage paradigm to enhance the inference efficiency. In practice, they first require localizing irrelevant weights [11, 15, 38] or redundant structures [8, 45, 47] within LVLMs. Then, these pruned or newly designed models are re-trained to minimize performance degradation. On the other hand, as shown in Fig. 1(b), the token-dependent pruning technologies compress inference costs by reducing the model’s input length [2, 5, 7, 35, 37]. Considering that not all visual tokens are related to a given question, this approach usually applies attention scores as a criterion to discard [7] or merge

non-significant tokens [6].

Although essential progress has been made, they face the following limitations: 1) Parameter-dependent acceleration methods mostly involve an additional procedure to fine-tune the pruned model [8, 47] or retrain a new VL model from scratch [11, 15]. This complex pipeline requires substantial annotations and huge computational resources to update the parameters of LVLMs. 2) For a long visual input, even if each token only has two options (*i.e.*, keep or discard), the solution space is still vast due to combination explosion [46]. How to select the most relevant tokens is still a significant challenge. 3) Despite parameters and inputs jointly leading to oversized computational costs, previous efforts treat them as two independent research lines and ignore their relationship.

In this paper, we systematically analyze the above limitations (a detailed discussion can be found in Sec. 3) and introduce new insights as follows: 1) We extend the layer redundancy study [31] to the multimodal domain and find that the scale of LVLMs can be effectively reduced by directly skipping some redundant layers. This insight suggests a promising parameter compression strategy without complex procedures to update LVLMs’ weights after pruning. 2) With regard to token pruning, extensive results show that these rule-based token selection methods (*e.g.*, attention score ranking [7]) cannot consistently localize the most relevant tokens. Inspired by [20], we add noise to rule-based token selection results and empirically show that some better token combinations can be captured by expanding selection space with noise. 3) We experimentally demonstrate that tokens and layers are correlated. Therefore, building a joint optimization model is crucial for integrating the pruning process.

Based on these insights, we argue that there is still a substantial opportunity to improve model efficiency, particularly for models already deployed. Consequently, we introduce a new framework called **Pruning All-Rounder**, abbreviated as PAR. As illustrated in Fig.1 (c), rather than relying on a single strategy to accelerate inference, our PAR focuses on jointly compressing redundant layers and tokens with self-supervised learning. Specifically, we first develop a simple yet effective transformer-based block as a meta-router to adaptively organize the pruning workflow. Additionally, considering that the traditional rule-based token selection method [7] can only offer a limited reference, we employ a noise perturbation mechanism to expand the search space. Finally, without depending on human or automated labeling, the meta-router can be optimized by contrasting changes in model responses. Benefiting from the above designs, PAR not only achieves a better trade-off between performance and computational costs but also offers several potential advantages: **1) Convenience.** Our method effectively compresses the computational scales without up-

dating LVLMs’ parameters after pruning or relying on extensive human-labeled annotations, making it conveniently applicable to already deployed models. **2) Model Agnostic.** Without bells and whistles, the PAR can be regarded as a plug-and-play module and easily integrated into other frameworks. **3) Flexibility.** Different domains have distinct pruning preferences. For example, discarding tokens would be risky in medical scenarios [1]. As an all-rounder, our method provides users with multiple selectable versions. Our contributions are as follows:

1) We systematically discuss the redundancy of LVLMs at both the layer and token levels. By conducting qualitative and quantitative analyses, our work draws a series of meaningful insights for inference acceleration.

2) We propose a new framework Pruning All-Rounder to reduce the computational costs of LVLMs in a self-supervised manner. Different from the previous works, the PAR can jointly compress the parameters and tokens without re-adjusting the model’s weights after pruning.

3) The proposed PAR introduces multiple alternative layouts to reduce the FLOPs while maintaining the original capacity of models. Extensive experiments show the effectiveness of our method across nine different benchmarks.

2. Related work

2.1. Large Vision-Language Models

In recent years, Large Vision-Language Models (LVLMs) have attracted significant attention due to their powerful reasoning capabilities. Unlike previous specialized vision-language models, these LVLMs inherit the capabilities of language models and can follow different human instructions. Typically, these models [3, 9, 25, 26] consist of three main modules: Visual Encoder, Large Language Model (LLM), and Multimodal Connector. However, with the continuous expansion of model scale, they are constrained by high computational costs. In this paper, we aim to accelerate the reasoning process of LVLMs to facilitate more user-friendly models. Meanwhile, since LLMs constitute the majority of parameters in LVLMs [10, 38, 44], our work focuses on compressing the LLM component to enhance overall efficiency.

2.2. Parameter Pruning

Many researchers have attempted to reduce the computational costs of LVLMs by decreasing model scales, which has led to the development of parameter pruning methods [10, 11, 14, 38, 44]. Existing parameter pruning approaches for VLMs can be broadly categorized into two types: task-specific pruning [10, 44] and task-agnostic pruning [11, 14, 38]. Given that prevalent LVLMs can handle a diverse array of vision-language tasks, task-agnostic pruning appears more practical than task-specific




Figure 2. Preliminary experiments of LVLMs: (a) The layer-level redundancy analysis. We observed that some redundant layers can be directly skipped without significantly affecting inference performance. (b) The token-level redundancy analysis. Some better token combinations can be captured by introducing noise. (c) The changes in cosine similarity across layers and tokens. The results show that layers and tokens are correlated.

approaches. However, these methods [11, 14, 38] typically involve complex weight search algorithms and require an additional retraining stage after pruning, which limits their practicality, particularly for already deployed LVLMs. Meanwhile, several recent efforts have focused on redesigning model architectures, such as Mixture-of-Experts (MoE) [47] and Mixture-of-Depth (MoD) [29], to reduce computational burden. Unlike these approaches, we demonstrate the significant layer redundancy present in current LVLMs, and remove redundant layers without model retraining while preserving the original capabilities.

2.3. Token Pruning

One well-known drawback of the Transformer [43] architecture is its quadratic computational complexity with respect to input length. Therefore, lengthy visual inputs represent another key factor impacting model efficiency. To accelerate LVLMs from the input perspective, researchers have proposed pruning [2, 7, 19, 24, 50] or merging [5, 37, 51] unnecessary tokens. For example, FastV [7] is the first work to identify redundant tokens through attention score ranking. VTW [24] analyzes attention scores between different token types (*i.e.*, instruction, visual, and text tokens) and discovers that visual tokens in deeper layers are rarely attended to, thus simply pruning all visual tokens in these layers. Using a slightly different attention-based approach, LLaVA-PruMerge [37] and HiRED [2] visualize attention scores between CLS tokens and spatial visual tokens. Their results show that many attention values are near-zero, revealing the redundancy of numerous visual tokens. More recently, Fit-Prune [50] prunes tokens by minimizing attention distribution divergence between pre-pruning and post-pruning states. Different from the above methods, our work aims to explore an easier-to-deploy approach that jointly compresses redundant layers and tokens without updating the LVLMs’ parameters after pruning.

3. Preliminary

3.1. Large Vision-Language Models

Given an image I , LVLMs can perform multiple Vision-Language tasks by following the user’s textual instruction Q . To accomplish the corresponding reasoning task, image input I would first be encoded into token sequence $[v_1, v_2, \dots, v_s]$ by a visual encoder (*e.g.*, CLIP-ViT [33]) and a multi-modal connector (*e.g.*, MLP [26] or Q-former [9]), where s denotes the sequence length of image tokens. Meanwhile, the textual input Q is converted as corresponding language embedding $[q_1, q_2, \dots, q_c]$, where c is the length of text input. Finally, the response of t -th generative step can be obtained in an auto-regressive manner:

$$p(y_t) = \prod_{t=1}^L p_{\theta}(y_t | v_1, \dots, v_s, q_1, \dots, q_c, y_{<t}). \quad (1)$$

Here, we omit the systematic prompt tokens for simplicity. In general, the above generation process can also be formalized as $y = f(x, \theta)$, where y , x , and θ are outputs, inputs and parameters of LVLMs, respectively. The f denotes the mapping from inputs to outputs. It can be seen that the high computational costs of LVLMs mainly arise from calculations with θ and x . Therefore, two research paradigms are presented for model acceleration: parameter-dependent and token-dependent strategies.

However, several issues remain to be answered before accelerating the LVLMs: 1) Existing parameter-dependent acceleration methods need to fine-tune pruned LVLMs to bridge the performance gap [11, 14], which limits their feasibility of compressing already deployed models. 2) Previous works have proved that not all visual tokens are related to inference [7], but how to select the most relevant tokens is still an open question. 3) The θ and x jointly lead to vast computational costs, while their relationship has yet to be thoroughly evaluated.

Next, we aim to answer the above questions with the popular VL model LLaVA-1.5-7B [25]. Meanwhile, to ensure that the conclusions are generalizable, related experiments are established by combining five benchmarks: ScienceQA [28], AOKVQA [36], POPE [23], SEED-Bench-Image [21] and MM-Bench [27].

3.2. Layer Redundancy

As mentioned above, these parameter-dependent compression strategies involve an additional process that updates the parameters of pruned LVLMs to minimize performance degeneration [15, 48, 49], which is unfriendly for practical applications. In this section, we aim to explore a more *easy-to-deploy* approach for parameter compression. Since existing LVLMs are composed of a stack of transformer blocks, we focus on these basic computational units (*i.e.*, layers).

Though the field of NLP has proved LLMs can directly short-cut some redundant layers [16, 31], the related study is still limited in our community. Inspired by [48, 49], we conduct related experiments to quantitatively analyze the layer-wise redundancy of LVLMs. In particular, as shown in Fig. 2(a), given the above five benchmarks, we first compute the averaged cosine similarity of the hidden states between two adjacent layers. The results show that these adjacent layers exhibit high similarity, meanwhile, the last 16 layers have higher redundancy (*i.e.*, the blue line). Correspondingly, the reasoning performance (*i.e.*, the red line) shows a limited degradation when skipping these latter layers. More importantly, ***the layer-level redundancy phenomenon offers a promising parameter compression strategy due to avoiding additional training process.***

3.3. Token Redundancy

In this section, we focus on lengthy visual input which is the other important reason for high computational costs. Considering that not all visual tokens are relevant to a given textual input, existing token selection approaches mostly apply attention score ranking to select the token subset [2, 7, 37]. For example, FastV [7] uses attention scores from the second layer to select tokens. However, we argue that such a rule-based method is incapable of consistently selecting the most relevant tokens in the vast solution space.

To validate this hypothesis, we utilize the attention ranking in FastV [7] as a reference and add the Gaussian noise $\mathcal{N} \sim (0, 10)$ with 5 different seeds to explore whether better token combinations exist. As shown in Fig. 2 (b), we compute accuracies with three token pruning strategies (*i.e.*, random selection, attention ranking, and perturbed attention ranking) under different token preserving ratios (*i.e.*, 10%, 20% and 30%). From the red box (attention ranking) and the green box (random selection), we find that attention-based token selection methods outperform random dropping by a large margin. However, by adding Gaussian

noise [52] to the attention values (the blue box), we see that the performance can be enhanced further. The above results indicate that the rule-based method can only provide a limited reference. Meanwhile, as a plausible approach, the noise perturbation has the potential to capture better token combinations by expanding the search space. Based on the above experiments, we conclude that ***although using attention ranking provides an important pruning reference, some better token combinations can be obtained by introducing the noise perturbation.***

3.4. Joint Compression Discussion

To further enhance inference efficiency, a natural idea is to integrate the above two pruning paradigms. In this section, we empirically investigate this strategy by computing the layer-wise cosine similarity across different token combinations. As we already demonstrated that deep layers exhibit higher redundancy compared to the shallow layers in Sec. 3.2, the related experiments are conducted using these deep layers (*i.e.*, layers after 16th). Specifically, as shown in Fig. 2(c), we first calculate the cosine scores from the 17th to the 20th layer where all tokens are preserved (*i.e.*, “All tokens”). It can be observed that the 17th and 20th layers can be prioritized for removal if layer-wise redundancy is solely considered. However, when 50% tokens are dropped using attention ranking, we observe a dramatic change in the cosine similarity of these layers. Furthermore, we add 5 randomly sampled Gaussian noise $\mathcal{N} \sim (0, 10)$ to the attention scores and remove tokens based on perturbed scores. We find the cosine similarity of these layers is not fixed and the token combinations will greatly influence the ranking of layer-wise importance. We conclude that ***the layer-dependent and token-dependent strategies are correlated, and thus a joint optimization model is crucial to integrate them.***

3.5. Summary

Below, we summarize our investigations against layer redundancy and token redundancy.

(1) For redundant parameters, we observe that directly skipping some deep layers does not dramatically affect the reasoning ability of LVLMs. This finding extend the layer redundancy study in NLP [31] to the multimodal domain and provides an *easy-to-deploy* parameter compression paradigm without additional processes to update the parameters of pruned LVLMs.

(2) Regarding the lengthy visual input, we reveal that the existing attention-based token selection methods cannot consistently select the most important tokens and lead to sub-optimal results. Fortunately, some better token combinations can be captured by introducing the Gaussian noise.

(3) We find that the token dropping and layer skipping are correlated. Therefore, an ideal acceleration method

should jointly model them to achieve a better trade-off between performance and computational costs.

4. Method

Based on these insights, we propose a new framework Pruning All-Rounder (PAR), which can adaptively drop redundant tokens and short-cut unnecessary layers in a self-supervised manner. As shown in Fig. 3, the PAR focuses on leveraging a lightweight transformer block as the meta-router to compress the high computational costs of LVLMs. More importantly, our approach can be individually optimized or implemented to tackle different scenarios. In other words, existing works such as rule-based token pruning [7] or early exiting mechanism [12] can also be regarded as specific versions of our method. Next, we will introduce our approach in detail.

4.1. Pruning All-Rounder

For a well-trained LVLM, the core objective of our work is to build a pruning flow by a meta-router \mathcal{O}_ϕ with parameters ϕ . Meanwhile, considering the complex relationship between tokens and layers, this \mathcal{O}_ϕ can adaptively model them and provide more flexible pruning layouts.

Specifically, given a lengthy visual input sequence $V = \{v_i\}_{i=1}^S$ and textual instruction $Q = \{q_i\}_{i=1}^C$, our work would only prune visual token V due to $S \gg C$ [7, 24]. Moreover, we build a group of learnable embedding $L = \{l_i\}_{i=1}^N$ as a layer-wise controller to manage the layer-skipping operation, where N is the number of candidate layers. Because Sec. 3.2 has proved that the last 16 layers exhibit higher redundancy, we set the $N = 16$ to represent the index of these layers. Then, the meta-router \mathcal{O}_ϕ is utilized to model the pruning process by feeding the V , Q and L . To ensure lightweight, our \mathcal{O}_ϕ is a basic transformer block [43] with $\sim 40M$ trainable parameters. The above computations can be formulated as:

$$Z_{out} = \mathcal{O}_\phi(\text{concat}[V; Q; L]), \quad (2)$$

where concat represent concatenate operation. Since our work focuses on dropping redundant visual tokens and layers, the $Z_{out} = [Z_{out}^V; Z_{out}^L]$ would only collect the corresponding outputs from V and L , while the outputs from Q are omitted. Taking advantage of the self-attention mechanism [43], each element of Z_{out} can adaptively aggregate the information from other inputs. Finally, we split the Z_{out} to Z_{out}^V and Z_{out}^L , which would be mapped into the corresponding action space (*i.e.*, keep or discard). We take Z_{out}^V as an example, the ranked scores P_{act}^V are obtained by:

$$Z_{act}^V = \sigma(\text{MLP}(Z_{out}^V)), P_{act}^V = \text{Rank}(Z_{act}^V), \quad (3)$$

where Rank and σ refer to the ranking operation and sigmoid active function. MLP denotes a Multi-Layer Perceptron. For Z_{out}^L , we also use the Eq. 3 to obtain P_{act}^L . During

inference, we directly select top- K layers and top- M tokens to discard. Meanwhile, the pruning flow can also be customized by only ranking either Z_{act}^V or Z_{act}^L to tackle different scenarios. For example, users can only rank Z_{act}^V to achieve token pruning.

4.2. Model Optimization

It can be noted that Eq.3 involves a non-differentiable operation, meanwhile, the meta-router \mathcal{O}_ϕ is difficult to optimize without explicit labels. To address this challenge, we reformulate the training procedure of \mathcal{O}_ϕ as a preference optimization task and use the DPO technique [34] to train our meta-router. Furthermore, the above training process can be implemented in a self-supervised manner, allowing it to work without any human or automated labeling.

Preference Data Construction. The core of DPO technology is to guide the model to generate expected responses by constructing high-quality positive and negative samples. In this paper, we extend this idea into the model compression domain with self-supervised learning.

In particular, given an image-question pair (V, Q) in an unlabeled multi-modal dataset, we remove K layers and M tokens to construct contrastive pruning actions. Here, the K is randomly selected from the last 16 layers based on the conclusions in Sec 3.2. As for visual tokens, to explore better combinations, we apply Gaussian noise to perturb the rule-based scores ranking [7] as mentioned in Sec 3.3 and discard M tokens with the lower scores. By performing the above operation 5 times for each sample, we can build a pruning action set \mathcal{S} . Moreover, as a meta-router, this preference set \mathcal{S} can also be specifically used for layer or token pruning when the K or M is set to 0. Then, we divide the \mathcal{S} into positive action \mathcal{S}^+ and negative action \mathcal{S}^- by comparing them with the original outputs \mathcal{S}_o (*i.e.*, outputs of the full LVLM). Considering that the outputs of the pruned model will exhibit significant changes compared to the original ones, we use KL-divergence [18] to measure whether the reasoning is correctly executed. Finally, following [22], we use balanced positive and negative sets to construct the preference dataset. In our experiments, we observed that satisfactory results have been achieved using only 5,000 unlabelled samples.

Preference Optimization Process. Compared to previous reinforcement learning [4, 39], the DPO uses a policy model and a reference model to avoid the complex optimization process. Benefiting from its simplicity and stability, we apply this technology to optimize our \mathcal{O}_ϕ . Meanwhile, previous works [17, 32] have demonstrated that better or comparable performance than the original DPO can be achieved by removing the reference model. Therefore,

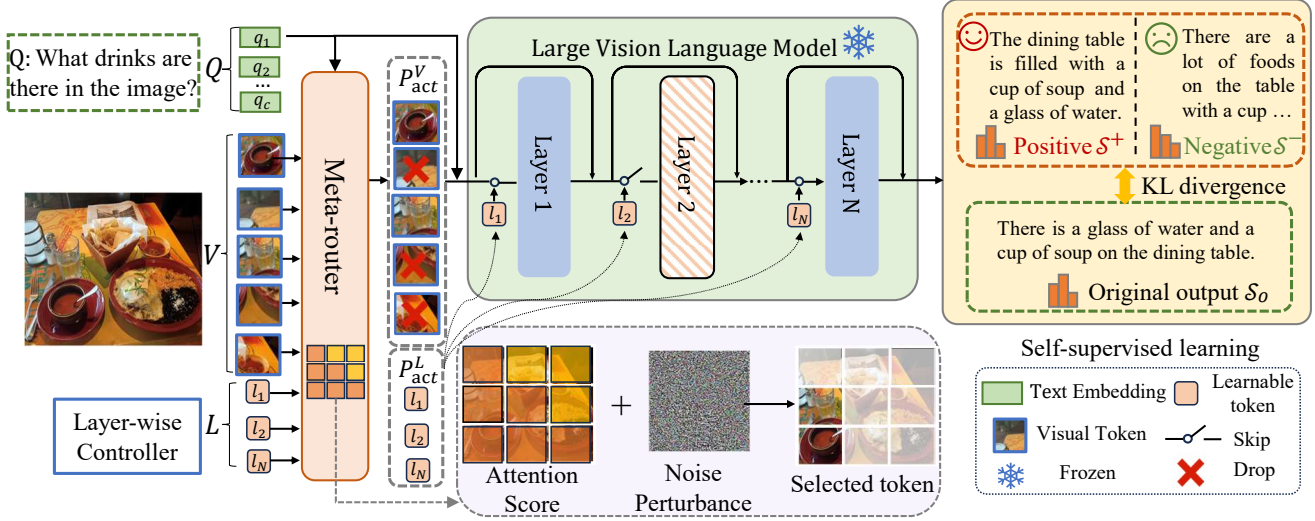


Figure 3. Overview of our Pruning All-Rounder (PAR). Given an image and corresponding question, our method focuses on applying a meta-router to adaptively prune redundant tokens and layers. Without human or automatic labeling, our PAR can be optimized by calculating the KL-divergence scores. Finally, as an all-rounder, the meta-router also provides multiple selectable pruning layouts to tackle different scenarios.

given the preference dataset $\mathcal{D} = \{\mathcal{S}^+, \mathcal{S}^-\}$, our optimization objective is formulated as follows:

$$\max_{\mathcal{O}_\phi} \mathbb{E}_{(x, \mathcal{S}^+, \mathcal{S}^-) \sim \mathcal{D}} [\log \sigma(\beta \log \mathcal{O}_\phi(\mathcal{S}^+ | x)) - \beta \log \mathcal{O}_\phi(\mathcal{S}^- | x)], \quad (4)$$

where $x = (V, Q, L)$ represents the input sequence. We follow [17, 32] and set the β as 1. With the above optimization objective, the PAR can learn to select the suitable model acceleration strategy based on different inputs. Note that we only optimize the lightweight \mathcal{O}_ϕ , while the parameters of LVLm would be frozen. More importantly, the \mathcal{O}_ϕ is only optimized to determine which visual tokens (or model layers) to prune **without changing the features of visual tokens or performing re-alignment between vision and language**. By keeping the parameters of LVLms frozen, we can enhance the efficiency of LVLms with minimal training overhead (*e.g.*, approximately 10 minutes of training time on a single NVIDIA 3090 GPU).

5. Experiment

5.1. Experimental Settings

Benchmarks and Backbones. We conduct extensive experiments on eight multimodal benchmarks, including AOKVQA [36], Science-QA [28], MME [13], POPE [23], MMBench [27], MMBench-CN [27], LLaVA-Bench-in-the-Wild [25] and SEED-Bench-Image [21]. Moreover, we apply PAR into two prevalent LVLms, including LLaVA-v1.5-7B [26] and Qwen-VL-Chat-9B [3]. During inference,

we adopt the same settings as in FastV [7] (such as the systematic prompt) for fair comparisons.

Implementation Details. Following FastV [7] and LLa-Volta [6], we report the theoretical TFLOPs and the corresponding reduction ratio. The PAR is trained on four NVIDIA 3090 GPUs and the batch size is 4. For fair comparisons, we adopt the same settings as in FastV [7] such as systematic prompt and decoding strategy. In our experiments, 5k data samples (less than 1%) are randomly sampled from LLaVA-80k [25] to construct preference data. Note that we set the pruning actions $K = 4$ and $M = 400$. During testing, the K and M can also be assigned by users without retraining a new \mathcal{O}_ϕ .

5.2. Quantitative Evaluation

Main Results. In Table 1, we compare the performance of our PAR with previous acceleration methods on two prevalent LVLms (*i.e.*, LLaVA-v1.5-7B [26] and Qwen-VL-Chat-9B [3]). For LLaVA-v1.5 [26], we quote related results from [6]. Our meta-router is trained under the setting of removing 400 tokens and 4 layers, and this result is marked with †. In Table 1, we use T and L to represent the number of retained tokens and layers for clarity. From the table, it can be observed that our PAR obtains better performance across nine benchmarks. Specifically, the results show that our PAR can outperform the state-of-the-art method FastV [7] by 3.3% while achieving more compression ratio of approximately 6.8%. In addition, when we further compress the model by only keeping

Table 1. For simplicity, the T and L represent the number of tokens and layers that are **preserved**. † denotes our meta-router is trained on this setting and other results are directly tested. * indicates that the original paper did not provide the results, and they are reproduced under the same settings. The results are sorted in descending order of TFLOPs. Note that the best results within the same TFLOPs range are highlighted in **bold**.

Method	TFLOPs (↓)	AOKVQA	SQA	MME	POPE	MMB	MMB ^{CN}	LLaVA ^W	SEED ^I
LLaVA-1.5-7B [26] ($T=576, L=32$)	11.05 (100%)	77.8	70.8	1510	86.1	65.3	59.4	65.5	66.7
+ ShortGPT [31]*	8.29 (75.0%)	74.2	64.6	964	69.7	50.4	37.5	52.9	54.3
+ RoE-LLaVA [47]	8.24 (74.6%)	-	68.7	-	-	64.6	-	-	57.8
+ γ -MoD [29]	7.98 (71.9%)	-	64.7	1342.1	86.0	59.3	-	-	-
+ PAR ($T=256, L=28$)	6.89 (62.4%)	77.7	70.4	1351	80.6	62.6	54.5	64.8	65.2
+ PAR ($T=176, L=28$)†	6.19 (56.0%)	77.7	70.0	1300	76.1	62.2	53.8	63.8	62.7
+ Random [6]	5.93 (53.7%)	72.7	69.3	1142	55.8	39.7	33.3	47.6	52.2
+ LLaVolta [6]	5.93 (53.7%)	74.9	69.4	1150	70.1	56.4	46.5	55.6	55.7
+ FastV [7]	5.93 (53.7%)	75.5	69.4	1298	65.6	60.1	53.0	54.8	56.3
+ PAR ($T=144, L=24$)	5.18 (46.9%)	77.6	69.8	1292	75.9	61.1	53.1	63.6	60.2
+ PAR ($T=128, L=24$)	5.07 (45.9%)	76.4	69.6	1286	73.6	60.2	50.7	62.5	59.4
Qwen-VL-Chat-9B [3] ($T=256, L=32$)	10.01 (100%)	75.6	68.2	1487	86.5	60.6	56.7	73.5	65.4
+ ShortGPT [31]*	7.52 (75.1%)	63.5	52.6	1398	81.1	46.6	39.2	61.3	56.0
+ Random [6]*	7.21 (72.0%)	70.1	64.9	1138	80.2	44.3	37.6	62.5	57.7
+ LLaVolta [6]*	7.21 (72.0%)	71.6	65.3	1336	80.8	51.1	45.8	64.0	59.8
+ FastV [7]*	7.21 (72.0%)	72.2	65.9	1405	81.4	53.5	49.1	64.8	60.1
+ PAR ($T=102, L=30$)†	7.34 (73.3%)	74.7	66.8	1464	82.5	57.3	55.1	69.3	61.7
+ PAR ($T=72, L=28$)	6.61 (66.0%)	73.5	66.3	1425	81.9	54.3	50.3	68.5	60.2

Table 2. Performance Comparison of different model settings. * denotes reproduced results, and LLaVA-1.5-7B [26] is used as the backbone.

	Method	TFLOPs (↓)	POPE	SEED ^I
<i>Token-level</i>				
1	Random Dropping	5.93 (53.7%)	55.8	52.2
2	LLaVolta [6]	5.93 (53.7%)	70.1	55.7
3	FastV [7]	5.93 (53.7%)	65.6	56.3
4	<i>joint-then-token</i>	5.93 (53.7%)	71.3	57.9
5	<i>token-router</i>	5.93 (53.7%)	72.7	59.0
<i>Layer-level</i>				
6	Random Dropping	8.29 (75.0%)	74.2	52.6
7	ShortGPT [31]*	8.29 (75.0%)	79.3	58.0
8	<i>joint-then-layer</i>	8.29 (75.0%)	82.1	60.5
9	<i>layer-router</i>	8.29 (75.0%)	82.9	61.1

128 visual tokens and 24 layers, we observe that our method exceeds the SOTA [7] by 7.8% lower TFLOPs. More importantly, compared to those training-based methods (e.g., RoE-LLaVA [47]), our PAR still obtains the SOTA performance without updating the parameters of pruned LVLMs. Finally, we expand our approach to the Qwen-VL-Chat [3]. The results show that the PAR can be regarded as a plug-and-play module and conveniently transferred into other VL models. As a preliminary exploration, we also report different versions across various combinations of tokens and

Table 3. **Ablation study.** We ablate key components to demonstrate the effectiveness of our method. LLaVA-1.5-7B [26] is used as the backbone, and “Full” denotes the original model. The “Layer”, “Token” and “Joint” indicate layer skipping, token dropping and joint pruning, respectively. DPO denotes the DPO strategy, while “Noise” is the Gaussian Noise Perturbation.

	TFLOPs	Layer	Token	DPO	Noise	POPE	SEED ^I
<i>Full ($T=576, L=32$)</i>							
1	11.05 (100%)					86.1	66.7
<i>Layer ($T=576, L=28$)</i>							
2	8.29 (75.0%)	✓				81.4	61.6
3	8.29 (75.0%)	✓		✓		85.8	65.1
<i>Token ($T=176, L=32$)</i>							
4	6.96 (63.0%)		✓			72.8	59.8
5	6.96 (63.0%)		✓	✓		74.9	61.2
6	6.96 (63.0%)		✓	✓	✓	76.5	63.0
<i>Joint ($T=176, L=28$)</i>							
7	6.19 (56.0%)	✓	✓			64.4	54.5
8	6.19 (56.0%)	✓	✓	✓	✓	76.1	62.7

layers. Notably, **one-training is enough**, these results can be directly obtained without retraining the meta-router. The flexibility allows the model to quickly adapt to different resource constraints and performance requirements.

User: What is the fire hydrant in the image surrounded by?

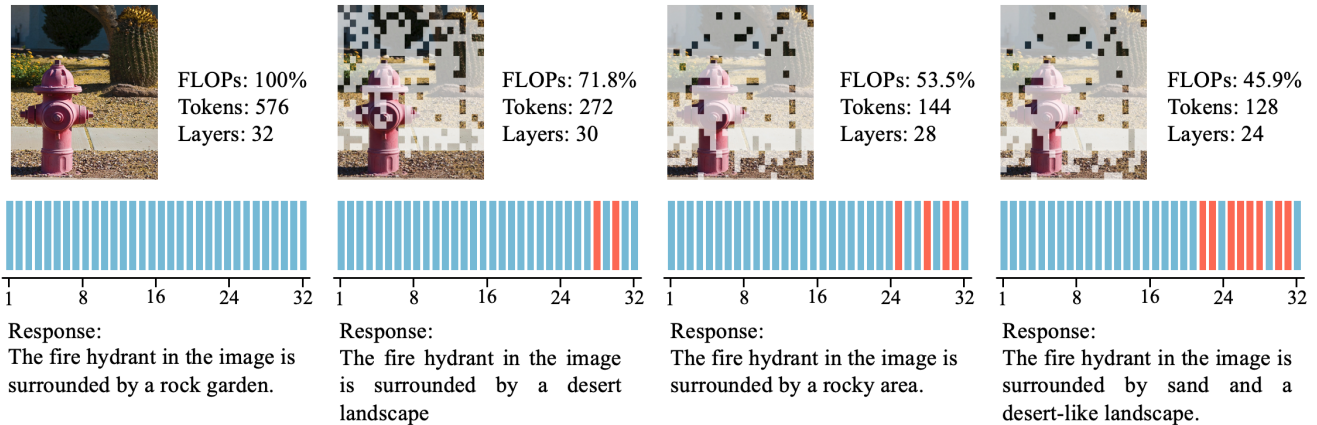


Figure 4. Qualitative results of our PAR method using LLaVA-1.5-7B [26]. The tokens and layers are gradually ablated to compress the model’s FLOPs.

Token or Layer Pruning. In Table 2, we further conduct experiments on the LLaVA-v1.5-7B [26] to verify the effectiveness of our method with specific settings. Specifically, we first train a meta-router (*i.e.*, $T=176$, $L=28$ in Table 1). During inference, we only perform the token or layer compression branch, referred to as “joint-then-token” and “joint-then-layer”, respectively. It can be observed that the PAR still achieves outstanding performance when only performing single-branch compression. Additionally, in the “token-router” and “layer-router”, we retrain two specialized routers to execute layer pruning or token removal, respectively. We observe that our PAR exceeds these rule-based strategies [7, 31] by a big margin. The above results show that the proposed PAR has great flexibility to tackle different acceleration scenarios.

5.3. Ablation Study

As shown in Table 3, we conduct several ablation studies on the POPE [36] and SEED [21] to systematically discuss our method. In the first row of Table 3, we report the original results with LLaVA-v1.5-7B [26]. Moreover, we construct three groups of experiments to verify the effectiveness of different modules. Specifically, we first focus on the layer-level redundancy in 2-3 rows of this Table. In row 2, we skip the first four layers with the highest cosine similarity. Then, the DPO strategy is introduced to train a specific layer-router (*i.e.*, the row 3). It can be observed that our method achieves better results compared to simply skipping specific layers. In rows 4-6, we explore the token-level compression method. In row 4, we report corresponding results by dropping visual tokens using FastV [7]. Then the DPO and Gaussian Noise are incorporated to construct different token selection manners in rows 5-6. These results indicate that our method provides a more effective

strategy for token pruning. Finally, we directly combine the rule-based layer skipping [31] and token pruning [7] in row 7. The result exhibits a significant decline due to the mutual influence between the token and layer. Conversely, our PAR in row 8 can improve the simple combining method by 11.7 and 8.2 points on POPE and SEED.

5.4. Qualitative Results

As shown in Fig.4, we visualize the model acceleration process on LLaVA-v1.5-7B [26] to further analyze and verify the proposed PAR. In this figure, we progressively ablate visual tokens and layers to reduce the FLOPs ratio (from left to right). From the figure, we find that our PAR effectively mitigates computation costs while maintaining reasoning ability. Notably, when PAR removes $\sim 80\%$ (*i.e.*, only remaining 128 tokens) of the visual tokens and skips 8 layers, resulting in only 45.9% FLOPs ratio, the model is still able to follow the instruction and provide a reasonable response.

6. Conclusion

In this paper, we focus on reducing the computational cost of Large Vision-Language Models (LVLMs) across parameters and inputs perspectives. By building comprehensive comparisons, we empirically demonstrate that LVLMs suffer from serious redundancy across layers and tokens. Based on a series of meaningful insights, we propose a new Pruning All-Rounder (PAR) framework, which can jointly compress the model’s parameters and inputs in a self-supervised learning manner. As a preliminary exploration, we expect our work to serve as a strong baseline and inspire our community to explore more effective acceleration strategies.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. U23B2013).

References

- [1] Ahmed Shihab Albahri, Ali M Duhaim, Mohammed A Fadhel, Alhamzah Alnoor, Noor S Baqer, Laith Alzubaidi, Osamah Shihab Albahri, Abdullah Hussein Alamoodi, Jinshuai Bai, Asma Salhi, et al. A systematic review of trustworthy and explainable artificial intelligence in healthcare: Assessment of quality, bias risk, and data fusion. *Information Fusion*, 96:156–191, 2023. 2
- [2] Kazi Hasan Ibn Arif, JinYi Yoon, Dimitrios S. Nikolopoulos, Hans Vandierendonck, Deepu John, and Bo Ji. Hired: Attention-guided token dropping for efficient inference of high-resolution vision-language models in resource-constrained environments, 2024. 1, 3, 4
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond, 2023. 1, 2, 6, 7
- [4] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022. 5
- [5] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster, 2023. 1, 3
- [6] Jieneng Chen, Luoxin Ye, Ju He, Zhao-Yang Wang, Daniel Khashabi, and Alan Yuille. Llavolta: Efficient multi-modal models via stage-wise visual context compression, 2024. 2, 6, 7
- [7] Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models, 2024. 1, 2, 3, 4, 5, 6, 7, 8
- [8] Xiaodong Chen, Yuxuan Hu, Jing Zhang, Yanling Wang, Cuiping Li, and Hong Chen. Streamlining redundant layers to compress large language models, 2024. 1, 2
- [9] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. 1, 2, 3
- [10] Zhiyuan Fang, Jianfeng Wang, Xiaowei Hu, Lijuan Wang, Yezhou Yang, and Zicheng Liu. Compressing visual-linguistic model via knowledge distillation, 2021. 2
- [11] Matteo Farina, Massimiliano Mancini, Elia Cunegatti, Gaowen Liu, Giovanni Iacca, and Elisa Ricci. Multi-flow: Shifting towards task-agnostic vision-language pruning, 2024. 1, 2, 3
- [12] Zhengcong Fei, Xu Yan, Shuhui Wang, and Qi Tian. Deecap: Dynamic early exiting for efficient image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12216–12226, 2022. 5
- [13] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. Mme: A comprehensive evaluation benchmark for multimodal large language models, 2024. 6
- [14] Zhe Gan, Yen-Chun Chen, Linjie Li, Tianlong Chen, Yu Cheng, Shuohang Wang, Jingjing Liu, Lijuan Wang, and Zicheng Liu. Playing lottery tickets with vision and language, 2021. 2, 3
- [15] Zhe Gan, Yen-Chun Chen, Linjie Li, Tianlong Chen, Yu Cheng, Shuohang Wang, Jingjing Liu, Lijuan Wang, and Zicheng Liu. Playing lottery tickets with vision and language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 652–660, 2022. 1, 2, 4
- [16] Andrey Gromov, Kushal Tirumala, Hassan Shapourian, Paolo Glorioso, and Daniel A. Roberts. The unreasonable ineffectiveness of the deeper layers, 2024. 4
- [17] Jiwoo Hong, Noah Lee, and James Thorne. Reference-free monolithic preference optimization with odds ratio. *arXiv preprint arXiv:2403.07691*, 2024. 5, 6
- [18] Shuyi Ji, Zizhao Zhang, Shihui Ying, Liejun Wang, Xibin Zhao, and Yue Gao. Kullback–leibler divergence metric learning. *IEEE transactions on cybernetics*, 52(4):2047–2058, 2020. 5
- [19] Yutao Jiang, Qiong Wu, Wenhao Lin, Wei Yu, and Yiyi Zhou. What kind of visual tokens do we need? training-free visual token pruning for multi-modal large language models from the perspective of graph. *arXiv preprint arXiv:2501.02268*, 2025. 3
- [20] Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13872–13882, 2024. 2
- [21] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension, 2023. 4, 6, 8
- [22] Lei Li, Zhihui Xie, Mukai Li, Shunian Chen, Peiyi Wang, Liang Chen, Yazheng Yang, Benyou Wang, and Lingpeng Kong. Silk: Preference distillation for large visual language models. *arXiv preprint arXiv:2312.10665*, 2023. 5
- [23] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023. 4, 6
- [24] Zhihang Lin, Mingbao Lin, Luxi Lin, and Rongrong Ji. Boosting multimodal large language models with visual tokens withdrawal for rapid inference, 2024. 3, 5
- [25] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023. 1, 2, 4, 6
- [26] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2024. 1, 2, 3, 6, 7, 8

- [27] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. Mmbench: Is your multi-modal model an all-around player?, 2024. 4, 6
- [28] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafford, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521, 2022. 4, 6
- [29] Yaxin Luo, Gen Luo, Jiayi Ji, Yiyi Zhou, Xiaoshuai Sun, Zhiqiang Shen, and Rongrong Ji. γ -mod: Exploring mixture-of-depth adaptation for multimodal large language models. *arXiv preprint arXiv:2410.13859*, 2024. 3, 7
- [30] Ji Ma, Wei Suo, Peng Wang, and Yanning Zhang. C3 l: content correlated vision-language instruction tuning data generation via contrastive learning. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 1155–1163, 2024. 1
- [31] Xin Men, Mingyu Xu, Qingyu Zhang, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han, and Weipeng Chen. Shortgpt: Layers in large language models are more redundant than you expect, 2024. 2, 4, 7, 8
- [32] Yu Meng, Mengzhou Xia, and Danqi Chen. Simpo: Simple preference optimization with a reference-free reward. *arXiv preprint arXiv:2405.14734*, 2024. 5, 6
- [33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3
- [34] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2024. 5
- [35] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification, 2021. 1
- [36] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge, 2022. 4, 6, 8
- [37] Yuzhang Shang, Mu Cai, Bingxin Xu, Yong Jae Lee, and Yan Yan. Llava-prumerge: Adaptive token reduction for efficient large multimodal models, 2024. 1, 3, 4
- [38] Dachuan Shi, Chaofan Tao, Ying Jin, Zhendong Yang, Chun Yuan, and Jiaqi Wang. Upop: Unified and progressive pruning for compressing vision-language transformers, 2023. 1, 2, 3
- [39] Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*, 2023. 5
- [40] Wei Suo, Mengyang Sun, Weisong Liu, Yiqi Gao, Peng Wang, Yanning Zhang, and Qi Wu. S3c: Semi-supervised vqa natural language explanation via self-critical learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2646–2656, 2023. 1
- [41] Wei Suo, Lanqing Lai, Mengyang Sun, Hanwang Zhang, Peng Wang, and Yanning Zhang. Rethinking and improving visual prompt selection for in-context learning segmentation. In *European Conference on Computer Vision*, pages 18–35. Springer, 2024.
- [42] Wei Suo, Lijun Zhang, Mengyang Sun, Lin Yuanbo Wu, Peng Wang, and Yanning Zhang. Octopus: Alleviating hallucination via dynamic contrastive decoding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 29904–29914, 2025. 1
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023. 3, 5
- [44] Jianfeng Wang, Xiaowei Hu, Pengchuan Zhang, Xiujun Li, Lijuan Wang, Lei Zhang, Jianfeng Gao, and Zicheng Liu. Minivlm: A smaller and faster vision-language model, 2021. 2
- [45] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E. Gonzalez. Skipnet: Learning dynamic routing in convolutional networks, 2018. 1
- [46] Cheng-En Wu, Jinhong Lin, Yu Hen Hu, and Pedro Morgado. Patch ranking: Efficient clip by learning to rank local patches. *arXiv preprint arXiv:2409.14607*, 2024. 2
- [47] Qiong Wu, Zhaoxi Ke, Yiyi Zhou, Gen Luo, Xiaoshuai Sun, and Rongrong Ji. Routing experts: Learning to route dynamic experts in multi-modal large language models, 2024. 1, 2, 3, 7
- [48] Qiong Wu, Zhaoxi Ke, Yiyi Zhou, Gen Luo, Xiaoshuai Sun, and Rongrong Ji. Routing experts: Learning to route dynamic experts in multi-modal large language models, 2024. 4
- [49] Qiong Wu, Wei Yu, Yiyi Zhou, Shubin Huang, Xiaoshuai Sun, and Rongrong Ji. Parameter and computation efficient transfer learning for vision-language pre-trained models. *Advances in Neural Information Processing Systems*, 36, 2024. 4
- [50] Weihao Ye, Qiong Wu, Wenhao Lin, and Yiyi Zhou. Fit and prune: Fast and training-free visual token pruning for multi-modal large language models. *arXiv preprint arXiv:2409.10197*, 2024. 3
- [51] Hongxu Yin, Arash Vahdat, Jose Alvarez, Arun Mallya, Jan Kautz, and Pavlo Molchanov. Advait: Adaptive tokens for efficient vision transformer, 2022. 3
- [52] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017. 4