

# ST3: Accelerating Multimodal Large Language Model by Spatial-Temporal Visual Token Trimming

Jiedong Zhuang<sup>1,2</sup>, Lu Lu<sup>2</sup>, Ming Dai<sup>1</sup>, Rui Hu<sup>1</sup>, Jian Chen<sup>2</sup>, Qiang Liu<sup>2</sup>, Haoji Hu<sup>1\*</sup>

<sup>1</sup>Zhejiang University

<sup>2</sup>Alibaba Cloud Computing

zhuangjiedong@zju.edu.cn, ll200214@alibaba-inc.com, mingdai997@outlook.com, rickhu@zju.edu.cn, j.chen@alibaba-inc.com, yf.yifeng@alibaba-inc.com, haoji\_hu@zju.edu.cn

## Abstract

Multimodal large language models (MLLMs) enhance their perceptual capabilities by integrating visual and textual information. However, processing the massive number of visual tokens incurs a significant computational cost. Existing analysis of the MLLM attention mechanisms remains shallow, leading to coarse-grain token pruning strategies that fail to effectively balance speed and accuracy. In this paper, we conduct a comprehensive investigation of MLLM attention mechanisms with LLaVA. We find that numerous visual tokens and partial attention computations are redundant during the decoding process. Based on this insight, we propose Spatial-Temporal Visual Token Trimming ( $ST^3$ ), a framework designed to accelerate MLLM inference without retraining.  $ST^3$  consists of two primary components: 1) Progressive Visual Token Pruning (PVTP), which eliminates inattentive visual tokens across layers, and 2) Visual Token Annealing (VTA), which dynamically reduces the number of visual tokens in each layer as the generated tokens grow. Together, these techniques deliver around  $2\times$  faster inference with only about 30% KV cache memory compared to the original LLaVA, while maintaining consistent performance across various datasets. Crucially,  $ST^3$  can be seamlessly integrated into existing pre-trained MLLMs, providing a plug-and-play solution for efficient inference.

## Introduction

The field of generative AI has recently witnessed explosive advancements, primarily driven by the development of large language models (LLMs) (Achiam et al. 2023; Touvron et al. 2023). Building upon this progress, Multimodal Large Language Models (MLLMs) (Liu et al. 2024a,b) have emerged as a powerful approach, integrating vision encoders like CLIP (Radford et al. 2021) to extract visual features. This integration enhances the reasoning capabilities of LLMs, enabling them to excel in complex tasks such as visual question answering (VQA) (Lu et al. 2022; Kembhavi et al. 2016), visual reasoning (Fu et al. 2024; Yue et al. 2024; Liu et al. 2023), visual grounding (Dai et al. 2024; Zhuang et al. 2025) and video captioning (Chai et al. 2024). However, the substantial computational demands of MLLMs pose a significant challenge, with their billions of parameters demanding




Figure 1: Comparison of various models on dataset ScienceQA\_Img (Lu et al. 2022) with the circle size representing their FLOPs. Our method achieves the highest accuracy in 13B parameter models while maintaining the minimum FLOPs and decoding latency, even outperforming the smaller model LLaVA-1.5-7B.

exponential increases in computational resources as the input sequence length grows.

Moreover, converting high-resolution images into visual tokens further exacerbates the inference cost (Liu et al. 2024b; Dong et al. 2024b; Luo et al. 2024), creating a significant obstacle to the practical applications of MLLMs in real-time scenarios. Previous methods (Cha et al. 2024; Li et al. 2022; Luo et al. 2024; Wu et al. 2024) attempt to address this by *introducing learnable modules* before feeding the tokens into the model, aiming to *reduce the length of the visual sequence*. While these methods can effectively reduce the inference latency, they typically require fine-tuning or retraining the entire MLLM. Given the current trend of rapidly increasing model sizes, the requirement of retraining or fine-tuning is suboptimal, since these processes are extremely resource-intensive and time-consuming for models with billions of parameters. On the other hand, recent works (Chen et al. 2024b; Lin et al. 2024) have focused on reducing the number of visual tokens in the inference process without retraining. These studies analyzed the atten-

\*Corresponding author.

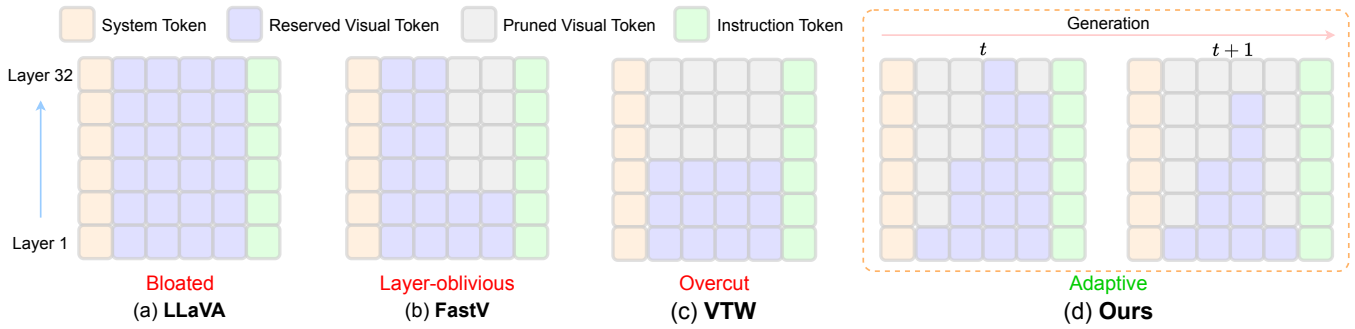


Figure 2: Illustration of our method compared with existing visual token pruning methods. (a) LLaVA (Liu et al. 2024a) keeps all visual tokens. (b) FastV (Chen et al. 2024b) prunes a fixed number of tokens in deeper layers. *This layer-oblivious paradigm overlooks the variability of attention patterns across layers.* (c) VTW (Lin et al. 2024) prunes all visual tokens in the latter half layers, *leading to a permanent loss of visual information in deeper layers.* Additionally, these three methods maintain the same quantity of visual tokens throughout the entire generation process, *requiring a substantial KV cache memory budget.* (d) Our method prunes inattentive visual tokens progressively as the layer goes deeper, while dynamically reducing tokens in the generation process. It maximizes the inference efficiency by exploiting the limit of the model’s dependence on visual tokens.

tion scores within the decoder and observed that the “Attention Sink” phenomenon (Xiao et al. 2023) is also present in MLLMs: a majority of the visual tokens receive a low level of attention, and these *unattended tokens can be directly pruned with negligible impact* on the model’s overall performance. However, these methods (Chen et al. 2024b; Lin et al. 2024) overlook two important aspects: 1) the iterative nature of the text generation process, where each newly generated token affects the allocation of visual attention in subsequent generation steps; 2) the potential performance gain from leveraging the internal distribution of the collection of visual tokens.

Although some token pruning methods (Xiao et al. 2023; Zhang et al. 2024b; Li et al. 2024b) have been proposed in LLM, token pruning in MLLM remains relatively underexplored. Due to the inherent differences between the modalities of images and texts, it is challenging to directly apply existing LLM-based techniques to MLLM. To bridge this gap, we conduct an in-depth investigation of the attention mechanisms during MLLM inference, using LLaVA as a case study. Our analysis reveals several intriguing phenomena that have not been observed in previous empirical studies (Chen et al. 2024b; Lin et al. 2024). For example, Fig. 4d shows that the attention within the visual token changes dramatically as the layers deepen: in the shallow layers, the attention is distributed uniformly across the visual tokens, whereas in the deep layers, the model selectively discards most visual tokens, focusing the attention on a subset of tokens. Additionally, Fig. 4b shows that as the length of the generated text tokens increases, the attention received by visual tokens gradually diminishes, indicating a decreasing reliance on visual features in the later stages of the generation process. Furthermore, by examining the cosine similarity between the attention of the different layers on various datasets, we find that the attention patterns between adjacent layers are highly similar (Fig. 3 and Fig. 4c). We refer to this phenomenon as “lazy layer”, where the subsequent layer essentially mimics the attention pattern of the preceding layer.

By incorporating the insights gained from our empirical studies, we propose an effective MLLM token pruning framework for accelerating MLLM inference. We first introduce a Progressive Visual Token Pruning (**PVTP**) mechanism, which progressively prunes inattentive visual tokens as the model’s layers deepen. Inspired by the “lazy layer” phenomenon, **PVTP** leverages a *step-wise* mechanism to perform token pruning, ensuring the same visual tokens are reused across the layers between two pruning executions. Furthermore, we introduce a method called Visual Token Annealing (**VTA**) during the generation process, which dynamically adjusts the number of retained visual tokens in the current decoding process based on the length of the generated text sequence. By combining **PVTP** and **VTA**, our approach differs significantly from previous token pruning techniques, as illustrated in Fig. 2. Our method maintains the quality of generated outputs and requires significantly fewer FLOPs, while achieving a lower decoding latency compared to existing methods on a wide range of multimodal tasks.

In summary, the main contributions of this paper can be outlined as follows:

- We conduct a comprehensive analysis of the visual attention, systematically elucidating the basis for principled and effective visual token pruning in MLLMs.
- We propose a novel, retraining-free acceleration framework for MLLMs, which achieves a more than 50% FLOPs reduction for LLaVA and improves its decoding speed by 2×, without significant impact on accuracy.
- We validate the superiority of our method over existing methods and quantify the impact of our proposed techniques across a wide range of multimodal task datasets.

## Related Work

**Multimodal Large language models.** The explosive growth of Large Language Models (LLMs) (Achiam et al. 2023; Bai et al. 2023a; Chiang et al. 2023; Touvron et al. 2023) has promoted rapid advancements of Multimodal




Figure 3: Illustration of the similarity between the attention scores of all layers in LLaVA-1.5-7B. High similarity scores are distributed around the diagonal, indicating adjacent layers have more similar attention patterns.

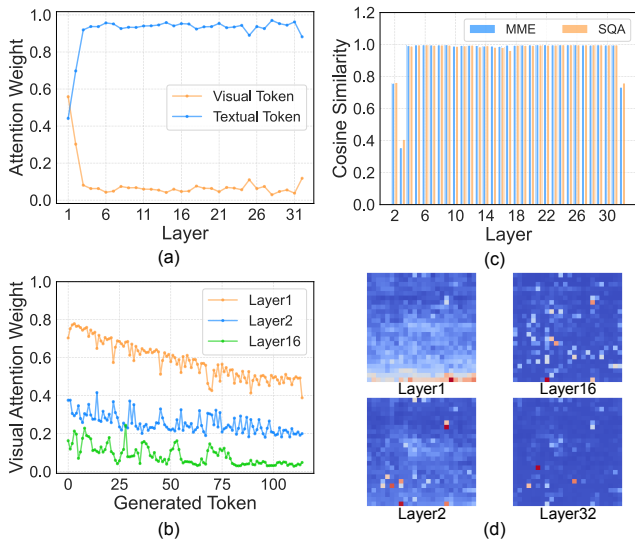


Figure 4: (a) Attention weight in various layers. Visual attention exhibits a persistently low magnitude after the layer3. (b) The weight of visual attention changes with the length of the generated text sequence. (c) Cosine similarity of the attention scores between each layer and its previous layer. A value close to 1 indicates that the attention score distributions in the two layers are nearly identical. (d) Visual attention. High attention tokens decrease as layers deepen.

Large Language Models (MLLMs), which combine the language understanding capabilities of LLMs with the visual processing abilities of vision models. These MLLMs have demonstrated remarkable success in tasks requiring the integration of textual and visual information, such as visual question answering, image captioning, and multimodal reasoning. The pioneering model CLIP (Radford et al. 2021) has taken a big stride in bridging the gap between the visual and textual modalities. Subsequent works (Alayrac et al. 2022; Li et al. 2022, 2023; Liu et al. 2024a,b) have built upon the CLIP model, using its visual encoder to extract visual tokens. They leverage the wealth of image-text paired datasets and the power of cross-modal alignment and optimization, resulting in a considerable enhancement of learning efficiency. This progression signifies an important advancement in the realm of MLLMs, expanding the range of applications by holistically embracing both textual

and visual modalities. Other impactful works include Gemini (Anil et al. 2023; Reid et al. 2024), MiniGPT4 (Zhu et al. 2023), InternVL (Chen et al. 2024c), CogVLM (Wang et al. 2023) and InternLM-XComposer (Zhang et al. 2023; Dong et al. 2024a; Zhang et al. 2024a; Dong et al. 2024b). While these works have been important drivers in the thriving development of the MLLMs field, our work focuses on exploring the redundant tokens and computations in MLLMs (e.g. LLaVA), with the goal of improving their inference efficiency.

**Vision Token Compression For MLLMs.** To reduce the unnecessary computational cost caused by the bloated visual tokens while effectively aligning the visual modality with the text modality, many methods (Alayrac et al. 2022; Bai et al. 2023b; Li et al. 2023; Luo et al. 2024; Cha et al. 2024; Li et al. 2024a) adopt a projector between the visual encoder and the LLM. Resampler (Bai et al. 2023b) and Q-former (Li et al. 2022, 2023) use cross-attention that transfers the visual information to a few learnable tokens. Honeybee (Cha et al. 2024) and MobileVLM (Chu et al. 2023, 2024) leverage convolution layers to aggregate local features and generate the compressed tokens. DeCo (Yao et al. 2024) employs a naive pooling technique to down-sample the visual features. SeTok (Wu et al. 2024) consolidates tokens carrying similar semantics. AcFormer (Liu et al. 2024c) employs attention focus of the visual encoder as anchors to drive the information aggregation. VoCo (Ye et al. 2024) compresses tokens via intrinsic token distillation. LLaVolta (Chen et al. 2024a) incorporates stage-wise visual context compression to progressively compress the visual tokens from heavily to lightly. These methods reduce the number of visual tokens, but they all require additional training of the MLLMs. To alleviate this overhead, several training-free methods have emerged. CrossGet (Shi et al. 2023) merges visual tokens through Complete-Graph Soft Matching. LLaVA-Prumerge (Shang et al. 2024) performs adaptive tokens merging before feeding them into the LLM, but ignores the necessity of modality interactions. FastV (Chen et al. 2024b) and VTW (Lin et al. 2024) leverage the attention scores within the LLM decoder to evaluate the importance of each token, and then prune the inattentive tokens. However, their analysis on attention is not comprehensive, lacking a dynamic, holistic observation of the entire generation process. Compared to these methods, we conduct a more systematic analysis of the attention in MLLMs and propose a training-free dynamic pruning technique.

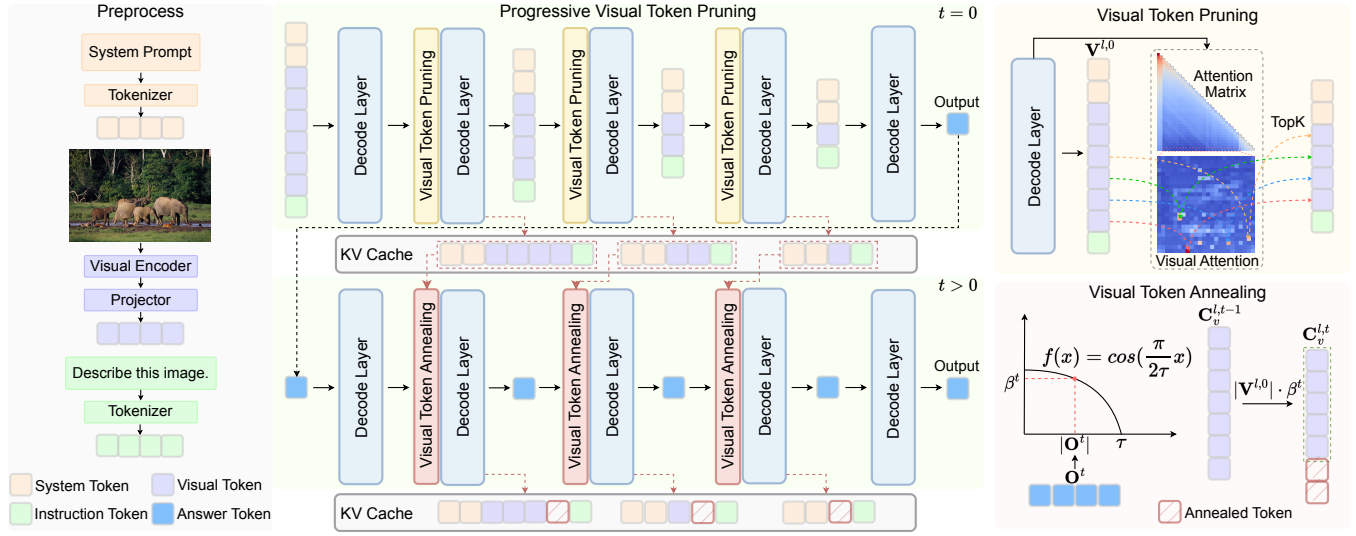


Figure 5: Overview of  $ST^3$  framework. **Preprocess** first converts the input from various modalities into tokens. All tokens are concatenated and fed into the LLM’s decoder layers. **Progressive Visual Token Pruning (PVT)** gradually prunes away non-critical visual tokens throughout the entire decoding forward process. Top-right illustrates the details of **Visual Token Pruning**. It extracts the visual token attention from the attention matrix of the previous decoding layer, and selectively retains the TopK most important tokens. The predicted token is then used as the input for the next generation step. **Visual Token Annealing (VTA)** employs a cosine function to control the decay of visual token KV cache (Pope et al. 2023) in each generation step.

## Methodology

### Preliminary

In this work, we use LLaVA (Liu et al. 2024a,b), a widely-adopted model with an exemplar MLLM architecture, to demonstrate our proposed visual token pruning strategy. LLaVA is composed of three pivotal components: 1) Visual Encoder, which employs the CLIP-ViT-L/14 (Radford et al. 2021; Dosovitskiy et al. 2020) as its backbone to convert an input image into visual embedding; 2) Visual Projector, which maps the visual embedding to the text embedding space through a fully connected layer; 3) LLM, which takes visual tokens and text tokens as input, and predicts the next text token in an autoregressive manner. After preprocessing different modalities input, all tokens are concatenated and the input of layers in LLM can then be represented as:

$$\mathbf{X}^{l,t} = [\mathbf{S}, \mathbf{V}^{l,t}, \mathbf{I}, \mathbf{O}^t] \quad (1)$$

where  $\mathbf{X}^{l,t}$  and  $\mathbf{V}^{l,t}$  denote the inputs and visual tokens of the  $l$ -th layer in  $t$ -th generation,  $\mathbf{S}$  and  $\mathbf{I}$  are system tokens and instruction tokens, and  $\mathbf{O}^t$  is the output token at  $t$ -th generation initialized as  $\emptyset$ . The pipeline of our method is illustrated in Fig. 5.

### Progressive Visual Token Pruning

Token pruning (Bolya et al. 2022; Norouzi et al. 2024; Wu et al. 2023) techniques are well-developed for ViTs, and these approaches are often able to effectively retain the most crucial tokens. However, these methods typically incur a significant computational overhead, and their effectiveness on MLLMs remains unclear. In this paper, we propose an approach based on attention score ranking to screen pivotal

tokens, which has been proven effective in FastV (Chen et al. 2024b). Specifically, attention scores are derived from the causal self-attention operation within a decoder layer in LLMs. The self-attention in  $l$ -th layer can be formulated as:

$$\text{Self-attn}(Q^{l,t}, K^{l,t}, V^{l,t}) = A^{l,t} V^{l,t} \quad (2)$$

where  $A^{l,t} = \text{Softmax}\left(\frac{Q^{l,t}(K^{l,t})^T + M}{\sqrt{d}}\right)$ ,  $Q^{l,t}$ ,  $K^{l,t}$  and  $V^{l,t}$  are query, key and value projected from  $X^{l,t}$ , and  $M$  is an upper triangular matrix where all non-zero elements are set to  $-\infty$  and the diagonal elements are set to 0.

Considering that the model relies solely on the last token to predict the next output token autoregressively (Lin et al. 2024), we take its last row as the attention score  $s^{l,t}$  and extract the attention scores on visual tokens  $\alpha^{l,t}$ . By sorting the  $\alpha^{l,t}$ , we retain visual tokens with higher importance:

$$\mathbf{V}^{l,t} = \mathbf{V}^{l,t}[\text{Top}(G, \alpha_i^{l-1,t})] \quad G = |\mathbf{V}^{l,t}| - N; t = 0 \quad (3)$$

where  $\alpha_i^{l-1,t}$  means attention score of  $i$ -th visual token,  $\text{Top}(G, *)$  denotes indices of highest  $G$  values,  $|\mathbf{V}^{l,t}|$  is the length of  $\mathbf{V}^{l,t}$ , and  $N$  is the number of pruned visual token.

The strategy we adopt is tailored to the model’s empirical behaviors and aligns well with our analysis of the attention mechanism in MLLMs (as in Fig. 4d): In the shallow layers, the attention on visual tokens is quite dense, and the model requires a larger number of visual tokens to gather information; as the depth increases, the visual attention becomes sparser, which means the model focusing on a smaller set of tokens. Therefore, to preserve the crucial shallow-level information, we avoid pruning tokens in the first three layers.

Moreover, we further conduct an in-depth analysis and statistical evaluation of the attention score  $s^{l,t}$ . Specifically,

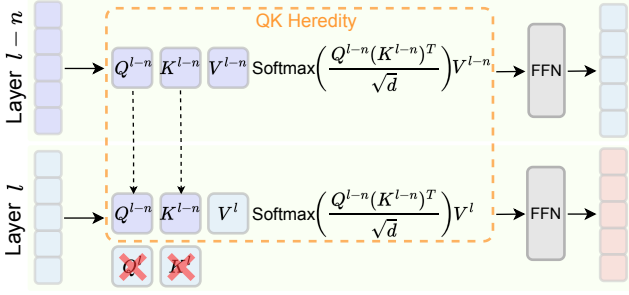


Figure 6: QK heredity. Lazy layers  $l$  inherits the  $Q^{l-n}$  and  $K^{l-n}$  from the shallower layer.

we calculate the cosine similarity  $c^{l,t}$  of the attention score distributions between each pair of adjacent layers. This can be expressed by the following formula:

$$c^{l,l-1,t} = \text{Cosine-similarity}(s^{l,t}, s^{l-1,t}) \quad l \in [2, 32] \quad (4)$$

Based on observation (Fig. 4c) of  $s^{l,t}$ , we find the distributions of attention scores between each layer  $l$  and its previous layer  $l-1$  are similar. Building upon the insights gained from this, we hypothesize that a similar pattern may exist across the multi layers as well. To validate this, we further expand Eq. 4:

$$c^{l,l',t} = \text{Cosine-similarity}(s^{l,t}, s^{l',t}) \quad l, l' \in [2, 32] \quad (5)$$

where  $c^{l,l',t}$  refers to cosine similarity of the attention score distributions between all layers. As shown in Fig. 3, the attention scores between adjacent layers indeed possess a shared pattern. In other words, we find “lazy” layers which are essentially mimicking their preceding layers. Based on these findings, we introduce two techniques to exploit this “lazy layer” phenomenon and reduce the computational overhead: 1) a strided pruning approach that skips over layers instead of pruning every layer, which reduces  $\text{Top}(G,*)$  computation, and 2) QK heredity at certain layers which reduces attention computation by inheriting the attention pattern from previous layers. Specifically, we formulate the attention operation in a “lazy layer”  $l$  as:

$$\text{Attention}(Q^l, K^l, V^l) = \text{Self-attn}(Q^{l-n}, K^{l-n}, V^l) \quad (6)$$

We let the “lazy layer” inherit the  $Q^{l-n}$ ,  $K^{l-n}$  and  $\text{Softmax}\left(\frac{Q^{l-n}(K^{l-n})^T}{\sqrt{d}}\right)$  results from the previous layer. Fig. 6 illustrates the principle of QK heredity.

### Visual Token Annealing

Based on another finding from our experiments: the proportion of visual attention decreases as the generated text tokens increase (as in Fig. 4b), we propose a strategy called Visual Token Annealing (VTA). VTA dynamically adjusts the quantity of KV cache corresponding to visual tokens retained at each layer based on generated text tokens over the whole autoregressive process. VTA can be formulated as:

$$C_v^{l,t} = C_v^{l,t-1} \cdot [|\mathbf{V}^{l,0}| \cdot \beta^t] \quad t > 0 \quad (7)$$

$$\beta^t = \begin{cases} \cos(|\mathbf{O}^t| \cdot \frac{\pi}{2\tau}) & |\mathbf{O}^t| < \tau \\ 0 & |\mathbf{O}^t| \geq \tau \end{cases} \quad (8)$$

where  $C_v$  means visual KV cache (sorted),  $|\mathbf{O}^t|$  denotes length of  $\mathbf{O}^t$ , and  $\tau$  is a hyper-parameter controlling the rate of decay in visual tokens while enforcing a truncation on the maximum generation length. The rationale for selecting the cosine function for VTA is based on the following two considerations when  $|\mathbf{O}^t| \in [0, \tau]$ : 1) the first derivative of the cosine function is negative, ensuring gradual shrinkage of visual tokens as the text tokens are generated; 2) its second derivative is also negative, reinforcing a larger compression of visual tokens in the later stages of the generation process.

## Experiments

### Comparison on Single-token Answer Datasets

We compare our method with existing methods on mainstream single-token answer datasets, with the results shown in Tab. 1. Among them, for the LLaVA-NeXT model, we use the LLaVA-1.6-vicuna version. Compared to the original LLaVA series models, our method reduces more than 50% FLOPs and achieves over 2-3 $\times$  improvements in inference speed, with almost no loss in accuracy. Additionally, compared to the current state-of-the-art methods, our method shows clear advantages in terms of inference speed and computational cost, while maintaining comparable or even better accuracy. The experimental results are consistent with our observations that MLLMs tend to focus their attention on a few tokens at deeper layers. Thus, as layers deepen, only fewer tokens need to be retained. This property allows us to drastically improve inference efficiency by pruning a large ratio of tokens without compromising accuracy.

### Ablation on Progressive Visual Token Pruning

We conduct ablation experiments on pruning parameters stride  $S$  and pruning ratio  $R$ . The proportion of conserved visual tokens in the last layer  $C$  can be calculated as:

$$C = 1 - \lfloor (L-3)/S \rfloor \cdot R - P \quad (9)$$

where  $P$  means pruning ratio in 4-th layer, which is set to 50% inspired by FastV (Chen et al. 2024b), and “ $\lfloor * \rfloor$ ” is the round-down operation. Tab. 3 shows the results when  $C = 1\%$  on LLaVA-1.5-7B and  $C = 5\%$  on LLaVA-1.5-13B. Lower FLOPs cannot achieve the lowest latency as a smaller pruning stride incurs a larger overhead. Thus, in practice, the selection of  $S$  and  $R$  is optimized with the primary objective of attaining the minimal latency with competitive accuracy. Incorporating QK heredity (Fig. 6) along with PVTP can further improve acceleration performance. Tab. 4 shows the effect of implementing QK heredity at different layers: in shallow layers, it leads to a decrease in accuracy, while the performance is robust when it is applied to later layers. We believe this is due to the errors introduced by QK heredity compounding as layers deepen.

### Comparison on Long-sentences Answer Datasets

For longer sentence generation, we employ the widely adopted KV cache (Pope et al. 2023) technique. After the

Methods	FLOPs↓	Latency↓	MME↑	AI2D↑	SQA↑	MMMU↑	MMB↑	POPE↑
LLaVA-1.5-7B	9.38T	70.80ms	1861.97	55.25	69.46	35.20	<b>64.00</b>	<b>86.99</b>
FastV‡ (Chen et al. 2024b)	9.38T ↓ 00.00%	70.80ms ↓ 00.00%	1862.12	55.34	68.77	35.00	63.83	85.13
FastV† (Chen et al. 2024b)	5.78T ↓ 38.38%	42.38ms ↓ 40.14%	1854.87	55.38	68.77	35.10	63.83	85.13
VTW (Lin et al. 2024)	5.38T ↓ 42.64%	46.24ms ↓ 34.69%	1849.51	<b>55.41</b>	<b>69.61</b>	<b>36.10</b>	63.92	86.89
ST <sup>3</sup> (Ours)	<b>4.37T ↓ 53.41%</b>	<b>40.12ms ↓ 43.33%</b>	<b>1866.71</b>	<b>55.41</b>	68.96	35.30	63.83	85.28
LLaVA-1.5-13B	17.81T	128.49ms	1817.95	59.26	72.78	35.00	<b>68.81</b>	87.09
FastV‡ (Chen et al. 2024b)	17.81T ↓ 00.00%	128.49ms ↓ 00.00%	1817.55	58.87	72.98	33.90	68.30	86.17
FastV† (Chen et al. 2024b)	10.73T ↓ 39.75%	73.34ms ↓ 42.92%	<b>1860.51</b>	58.84	72.88	34.70	68.56	86.67
VTW (Lin et al. 2024)	10.14T ↓ 43.07%	80.74ms ↓ 37.16%	1836.18	<b>59.42</b>	72.68	<b>35.10</b>	<b>68.81</b>	<b>87.17</b>
ST <sup>3</sup> (Ours)	<b>7.88T ↓ 55.76%</b>	<b>63.15ms ↓ 50.85%</b>	1830.51	58.81	<b>73.18</b>	34.60	68.21	86.81
LLaVA-NeXT-7B	30.95T	408.61ms	1850.56	65.25	<b>70.15</b>	35.40	<b>67.10</b>	87.62
FastV‡ (Chen et al. 2024b)	30.95T ↓ 00.00%	408.61ms ↓ 00.00%	1806.93	65.12	69.21	35.70	66.41	<b>87.78</b>
FastV† (Chen et al. 2024b)	17.83T ↓ 42.39%	207.34ms ↓ 49.26%	1804.95	64.64	69.21	35.10	66.32	87.47
VTW (Lin et al. 2024)	16.74T ↓ 45.91%	229.81ms ↓ 43.76%	<b>1852.56</b>	<b>65.41</b>	70.00	<b>35.80</b>	<b>67.10</b>	87.60
ST <sup>3</sup> (Ours)	<b>13.87T ↓ 55.19%</b>	<b>148.85ms ↓ 63.57%</b>	1822.42	64.51	69.06	35.10	66.67	87.39
LLaVA-NeXT-13B	57.66T	691.02ms	<b>1902.39</b>	70.27	73.57	<b>36.10</b>	<b>69.33</b>	<b>87.53</b>
FastV‡ (Chen et al. 2024b)	57.66T ↓ 00.00%	691.02ms ↓ 00.00%	1887.00	69.95	72.98	35.90	68.81	87.48
FastV† (Chen et al. 2024b)	32.19T ↓ 44.17%	346.71ms ↓ 49.83%	1874.82	69.98	73.03	35.90	68.47	87.30
VTW (Lin et al. 2024)	30.61T ↓ 46.91%	363.98ms ↓ 47.33%	1879.96	70.24	73.43	35.30	69.07	87.44
ST <sup>3</sup> (Ours)	<b>22.02T ↓ 61.81%</b>	<b>218.71ms ↓ 68.35%</b>	1886.23	<b>70.29</b>	<b>73.63</b>	35.30	69.07	87.44

Table 1: Comparison of existing training-free token pruning method on MLLMs with single-token answer datasets. SQA means the ScienceQA(Lu et al. 2022) image subset, MMMU represents the validation subset of MMMU(Yue et al. 2024), and MMB denotes the english subset of MMBench(Liu et al. 2023). Latency refers to the total time required for the forward propagation to pass through all decode layers. “‡” means attention mask implementation and “†” means token drop implementation in its code base. The best results are in **bold**.




Figure 7: *Left*: Changes in decode latency under distinct batch sizes. *Right*: Relationship between KV cache Memory, decode latency and the length of generated sequence, with imagesize=336\*336 and batchsize=16.

first token generation, subsequent token generations directly read the Key-Value from the KV cache for the attention computation. Therefore, the implementation of Visual Token Annealing (VTA) actually acts on the KV cache by evicting the Key-Value pairs corresponding to the visual token positions. Based on extensive experimentation, we find: (1) It is unnecessary to *retain the entire KV cache of visual tokens in the first generation*, and (2) KV cache of Visual tokens can be *gradually discarded as the generation unfolds*. Tab. 2 demonstrates that our method can achieve a comparable generation quality, while only requiring less than 50% or even 30% of the KV cache memory compared to LLaVA. Notably, VTW (Lin et al. 2024) fails on long-token generation datasets while it performs well on single-token datasets, showing our method’s robustness in comparison. In the left of Fig. 7, we show the results from testing the latency of generating one token for different methods with KV cache under

varying batchsizes. PVTP exhibits lower latency compared to existing methods across various settings. The right of Fig. 7 illustrates the relationship between the length of the generated sequence and the KV cache memory usage. It shows that VTA maintains a low memory usage throughout the entire generation pipeline. The kinks in the blue and orange line represent the points where all the cache corresponding to visual tokens is cleared at that node.

To understand why partial KV cache is enough, we note that: 1) as shown in Tab. 5, a large proportion of the Top 50% attention visual tokens overlaps across different steps, suggesting that the KV cache trimmed after the first generation step can be reused in the subsequent generation steps; 2) the distribution of visual token attention remains similar at each generation step (see Fig. 8). The overlap is defined by:

$$\text{Overlap}_n = \frac{|\text{TopK}(T_1) \cap \text{TopK}(T_n)|}{|\text{TopK}(T_1)|} \quad (10)$$

where  $T_n$  means the visual tokens in  $n$ -th generation step.  $|\cdot|$  denotes obtaining the number of elements in a set.

### Ablation on Visual Token Annealing

Tab. 6 provides a comparative analysis of different decay schemes. The Exponential attenuation function is formulated as  $\beta^t = e^{-|\mathbf{O}^t|/\sigma}$ , and the Linear attenuation function is  $\beta^t = 1 - \frac{|\mathbf{O}^t|}{\tau}$  when  $|\mathbf{O}^t| < \tau$  and 0 otherwise. The Cosine function demonstrates superior performance compared to the other two, which aligns well with our intuition because it has a lower annealing rate in the initial generation process and thus better preserves the integrity of visual information. In contrast, the Exponential function has a faster

Methods	Memory↓	Coco2017↑	Flickr30k↑	Nocaps↑
LLaVA-1.5-7B	4.94G	110.43	74.89	105.53
FastV‡	4.94G	<b>110.80</b>	74.70	105.36
VTW	2.47G	67.20	40.65	95.78
ST <sup>3</sup> (Ours)	<b>2.27G</b>	110.66	<b>74.91</b>	<b>105.62</b>
LLaVA-1.5-13B	7.70G	115.57	79.56	<b>109.31</b>
FastV‡	7.70G	115.91	79.68	108.85
VTW	3.75G	101.67	65.87	95.21
ST <sup>3</sup> (Ours)	<b>2.51G</b>	<b>116.33</b>	<b>79.86</b>	108.45
LLaVA-NeXT-7B	15.68G	<b>99.87</b>	<b>68.47</b>	<b>88.37</b>
FastV‡	15.68G	98.54	67.18	86.48
VTW	7.84G	82.50	57.81	43.76
ST <sup>3</sup> (Ours)	<b>7.21G</b>	98.57	66.93	86.29
LLaVA-NeXT-13B	24.48G	<b>101.93</b>	<b>66.70</b>	<b>88.18</b>
FastV‡	24.48G	101.29	66.13	87.86
VTW	12.24G	86.86	56.52	43.26
ST <sup>3</sup> (Ours)	<b>7.98G</b>	101.37	65.99	87.96

Table 2: Comparison of existing training-free token pruning method on MLLMs with image caption datasets. “Memory” means KV cache memory, which is calculated over the same sample after first token is generated (batchsize=16). “Coco2017” refers to validation subset of COCO2017 caption(Chen et al. 2015). “Flickr30k” and “Nocaps” are validation and test splits in the original datasets(Plummer et al. 2015; Agrawal et al. 2019). The best results are in **bold**.

decay rate in the initial stages and the decay rate decreases uniformly over time, conflicting with the model’s reliance on visual tokens in the early generation process. The results reveal that convex functions may be more suitable for VTA. A study of the parameter  $\tau$  is conducted in the cosine function. Notably, the accuracy plunges when  $\tau = 10$ , which can be attributed to the mismatch between the  $\tau$  and the actual length of the generated tokens: visual tokens are completely discarded once the generated tokens length exceeds 10, resulting in the loss of image information in the subsequent generation process. Therefore, we recommend setting  $\tau$  to be greater than the preset maximum generation length.

Step	1	2	3	4	5
Overlap	100.00%	90.97%	82.29%	84.72%	86.46%
Step	6	7	8	9	10
Overlap	88.89%	82.99%	86.81%	88.19%	88.54%

Table 5: Ratio of Top 50% attention visual token in each generation step that overlap with the first step.




Figure 8: Attention on visual tokens in different generation step. The distribution of visual attention is similar across the generation process.

LLaVA-1.5-7B ( $C = 1\%$ )					
$S$	$R$	FLOPs	Latency	MME	AI2D
1	1.75%	<b>4.03T</b>	43.88ms	1836.47	54.63
2	3.50%	4.08T	41.05ms	1862.15	54.53
4	7.00%	4.20T	40.74ms	1840.29	54.99
7	12.25%	4.37T	<b>40.12ms</b>	<b>1866.71</b>	<b>55.41</b>
14	24.50%	4.78T	40.24ms	1861.87	55.41
28	49.00%	5.59T	42.67ms	1858.10	55.31
LLaVA-1.5-13B ( $C = 5\%$ )					
$S$	$R$	FLOPs	Latency	MME	AI2D
1	1.25%	<b>7.64T</b>	68.28ms	1826.94	58.78
2	2.50%	7.72T	63.96ms	<b>1851.51</b>	58.91
3	3.75%	7.80T	64.54ms	1833.87	59.10
4	5.00%	7.88T	<b>63.15ms</b>	1830.51	58.81
6	7.50%	8.04T	63.34ms	1827.58	<b>59.07</b>
9	11.25%	8.29T	64.06ms	1832.09	58.71

Table 3: Results of the pruning stride “ $S$ ” and the pruning ratio “ $R$ ” per step. The best results are in **bold**.

Methods	FLOPs↓	Latency↓	MME↑	AI2D↑	SQA↑	MMMU↑
ST <sup>3</sup>	4.37T	40.12ms	1866.71	55.41	68.96	35.30
7→8~10	4.30T	38.66ms	1480.70	43.10	62.47	34.20
14→15~17	4.31T	38.89ms	1825.35	49.32	62.62	32.10
21→22~24	4.33T	39.15ms	1864.18	54.70	69.01	36.30
28→29~31	4.34T	38.88ms	1863.54	55.25	69.06	35.70

Table 4: Results of implementing QK heredity (Fig. 6) across multiple layers. “4→5~7” means reusing attention score of layer4 from layer5 to layer7 based on ST<sup>3</sup>.

Function	Coco2017	Flickr30k	Nocaps
Exp. ( $\sigma = 20$ )	110.29	73.65	103.55
Linear ( $\tau = 50$ )	110.50	74.33	104.59
Cosine ( $\tau = 50$ )	<b>110.66</b>	<b>74.91</b>	<b>105.62</b>
Cosine ( $\tau = 100$ )	110.58	74.47	104.96
Cosine ( $\tau = 75$ )	110.55	74.49	104.99
Cosine ( $\tau = 25$ )	110.49	74.27	104.64
Cosine ( $\tau = 10$ )	103.05	64.45	96.68

Table 6: Comparison of the attenuation function with varied parameters on LLaVA-1.5-7B. “Exp.” means exponential. The best results are in **bold**.

## Conclusion

This paper introduces Spatial-Temporal Visual Token Trimming (ST<sup>3</sup>), a novel framework for pruning visual tokens of Multimodal Large Language Models (MLLM) across both *layer* and *time* dimensions. ST<sup>3</sup> is based on two key observations from our analysis of MLLM attention mechanisms: numerous visual tokens and partial attention computations are redundant during decoding, and subsequent layers often mimic the attention patterns of preceding ones (“lazy layers”). Experimental results show that ST<sup>3</sup> delivers around 2× faster inference with only about 30% KV cache memory compared to the original LLaVA, while maintaining consistent performance across various datasets. We anticipate that this work provides valuable empirical insights and advances the field of MLLM token pruning.

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Agrawal, H.; Desai, K.; Wang, Y.; Chen, X.; Jain, R.; Johnson, M.; Batra, D.; Parikh, D.; Lee, S.; and Anderson, P. 2019. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8948–8957.
- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hason, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736.
- Anil, R.; Borgeaud, S.; Wu, Y.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023a. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023b. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Bolya, D.; Fu, C.-Y.; Dai, X.; Zhang, P.; Feichtenhofer, C.; and Hoffman, J. 2022. Token merging: Your vit but faster. *arXiv preprint arXiv:2210.09461*.
- Cha, J.; Kang, W.; Mun, J.; and Roh, B. 2024. Honeybee: Locality-enhanced projector for multimodal llm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13817–13827.
- Chai, W.; Song, E.; Du, Y.; Meng, C.; Madhavan, V.; Bar-Tal, O.; Hwang, J.-N.; Xie, S.; and Manning, C. D. 2024. AuroraCap: Efficient, Performant Video Detailed Captioning and a New Benchmark. *arXiv preprint arXiv:2410.03051*.
- Chen, J.; Ye, L.; He, J.; Wang, Z.-Y.; Khashabi, D.; and Yuille, A. 2024a. LLaVolta: Efficient Multi-modal Models via Stage-wise Visual Context Compression. *arXiv preprint arXiv:2406.20092*.
- Chen, L.; Zhao, H.; Liu, T.; Bai, S.; Lin, J.; Zhou, C.; and Chang, B. 2024b. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. *arXiv preprint arXiv:2403.06764*.
- Chen, X.; Fang, H.; Lin, T.-Y.; Vedantam, R.; Gupta, S.; Dollár, P.; and Zitnick, C. L. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024c. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24185–24198.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna>, 3(5).
- Chu, X.; Qiao, L.; Lin, X.; Xu, S.; Yang, Y.; Hu, Y.; Wei, F.; Zhang, X.; Zhang, B.; Wei, X.; et al. 2023. Mobilevlm: A fast, strong and open vision language assistant for mobile devices. *arXiv preprint arXiv:2312.16886*.
- Chu, X.; Qiao, L.; Zhang, X.; Xu, S.; Wei, F.; Yang, Y.; Sun, X.; Hu, Y.; Lin, X.; Zhang, B.; et al. 2024. Mobilevlm v2: Faster and stronger baseline for vision language model. *arXiv preprint arXiv:2402.03766*.
- Dai, M.; Yang, L.; Xu, Y.; Feng, Z.; and Yang, W. 2024. SimVG: A Simple Framework for Visual Grounding with Decoupled Multi-modal Fusion. *arXiv preprint arXiv:2409.17531*.
- Dong, X.; Zhang, P.; Zang, Y.; Cao, Y.; Wang, B.; Ouyang, L.; Wei, X.; Zhang, S.; Duan, H.; Cao, M.; et al. 2024a. Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model. *arXiv preprint arXiv:2401.16420*.
- Dong, X.; Zhang, P.; Zang, Y.; Cao, Y.; Wang, B.; Ouyang, L.; Zhang, S.; Duan, H.; Zhang, W.; Li, Y.; et al. 2024b. Internlm-xcomposer2-4khd: A pioneering large vision-language model handling resolutions from 336 pixels to 4k hd. *arXiv preprint arXiv:2404.06512*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fu, C.; Chen, P.; Shen, Y.; Qin, Y.; Zhang, M.; Lin, X.; Yang, J.; Zheng, X.; Li, K.; Sun, X.; Wu, Y.; and Ji, R. 2024. MME: A Comprehensive Evaluation Benchmark for Multimodal Large Language Models. *arXiv:2306.13394*.
- Kembhavi, A.; Salvato, M.; Kolve, E.; Seo, M.; Hajishirzi, H.; and Farhadi, A. 2016. A diagram is worth a dozen images. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, 235–251. Springer.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, 12888–12900. PMLR.
- Li, W.; Yuan, Y.; Liu, J.; Tang, D.; Wang, S.; Zhu, J.; and Zhang, L. 2024a. TokenPacker: Efficient Visual Projector for Multimodal LLM. *arXiv preprint arXiv:2407.02392*.
- Li, Y.; Huang, Y.; Yang, B.; Venkitesh, B.; Locatelli, A.; Ye, H.; Cai, T.; Lewis, P.; and Chen, D. 2024b. Snapkv: Llm knows what you are looking for before generation. *arXiv preprint arXiv:2404.14469*.

- Lin, Z.; Lin, M.; Lin, L.; and Ji, R. 2024. Boosting Multimodal Large Language Models with Visual Tokens Withdrawal for Rapid Inference. *arXiv preprint arXiv:2405.05803*.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26296–26306.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024b. Llava-next: Improved reasoning, ocr, and world knowledge.
- Liu, H.; You, Q.; Han, X.; Liu, Y.; Huang, H.; He, R.; and Yang, H. 2024c. Visual Anchors Are Strong Information Aggregators For Multimodal Large Language Model. *arXiv preprint arXiv:2405.17815*.
- Liu, Y.; Duan, H.; Zhang, Y.; Li, B.; Zhang, S.; Zhao, W.; Yuan, Y.; Wang, J.; He, C.; Liu, Z.; et al. 2023. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- Lu, P.; Mishra, S.; Xia, T.; Qiu, L.; Chang, K.-W.; Zhu, S.-C.; Tafjord, O.; Clark, P.; and Kalyan, A. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35: 2507–2521.
- Luo, G.; Zhou, Y.; Zhang, Y.; Zheng, X.; Sun, X.; and Ji, R. 2024. Feast Your Eyes: Mixture-of-Resolution Adaptation for Multimodal Large Language Models. *arXiv preprint arXiv:2403.03003*.
- Norouzi, N.; Orlova, S.; de Geus, D.; and Dubbelman, G. 2024. ALGM: Adaptive Local-then-Global Token Merging for Efficient Semantic Segmentation with Plain Vision Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15773–15782.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, 2641–2649.
- Pope, R.; Douglas, S.; Chowdhery, A.; Devlin, J.; Bradbury, J.; Heek, J.; Xiao, K.; Agrawal, S.; and Dean, J. 2023. Efficiently scaling transformer inference. *Proceedings of Machine Learning and Systems*, 5: 606–624.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Reid, M.; Savinov, N.; Teplyashin, D.; Lepikhin, D.; Lillcrap, T.; Alayrac, J.-b.; Soricut, R.; Lazaridou, A.; Firat, O.; Schrittwieser, J.; et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Shang, Y.; Cai, M.; Xu, B.; Lee, Y. J.; and Yan, Y. 2024. Llava-prumerge: Adaptive token reduction for efficient large multimodal models. *arXiv preprint arXiv:2403.15388*.
- Shi, D.; Tao, C.; Rao, A.; Yang, Z.; Yuan, C.; and Wang, J. 2023. Crossget: Cross-guided ensemble of tokens for accelerating vision-language transformers. *arXiv preprint arXiv:2305.17455*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, W.; Lv, Q.; Yu, W.; Hong, W.; Qi, J.; Wang, Y.; Ji, J.; Yang, Z.; Zhao, L.; Song, X.; et al. 2023. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*.
- Wu, S.; Fei, H.; Li, X.; Ji, J.; Zhang, H.; Chua, T.-S.; and Yan, S. 2024. Towards Semantic Equivalence of Tokenization in Multimodal LLM. *arXiv preprint arXiv:2406.05127*.
- Wu, X.; Zeng, F.; Wang, X.; Wang, Y.; and Chen, X. 2023. Ppt: Token pruning and pooling for efficient vision transformers. *arXiv preprint arXiv:2310.01812*.
- Xiao, G.; Tian, Y.; Chen, B.; Han, S.; and Lewis, M. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- Yao, L.; Li, L.; Ren, S.; Wang, L.; Liu, Y.; Sun, X.; and Hou, L. 2024. DeCo: Decoupling Token Compression from Semantic Abstraction in Multimodal Large Language Models. *arXiv preprint arXiv:2405.20985*.
- Ye, X.; Gan, Y.; Huang, X.; Ge, Y.; Shan, Y.; and Tang, Y. 2024. VoCo-LLaMA: Towards Vision Compression with Large Language Models. *arXiv preprint arXiv:2406.12275*.
- Yue, X.; Ni, Y.; Zhang, K.; Zheng, T.; Liu, R.; Zhang, G.; Stevens, S.; Jiang, D.; Ren, W.; Sun, Y.; et al. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9556–9567.
- Zhang, P.; Dong, X.; Zang, Y.; Cao, Y.; Qian, R.; Chen, L.; Guo, Q.; Duan, H.; Wang, B.; Ouyang, L.; et al. 2024a. InternLM-XComposer-2.5: A Versatile Large Vision Language Model Supporting Long-Contextual Input and Output. *arXiv preprint arXiv:2407.03320*.
- Zhang, P.; Wang, X. D. B.; Cao, Y.; Xu, C.; Ouyang, L.; Zhao, Z.; Ding, S.; Zhang, S.; Duan, H.; Yan, H.; et al. 2023. Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112*.
- Zhang, Z.; Sheng, Y.; Zhou, T.; Chen, T.; Zheng, L.; Cai, R.; Song, Z.; Tian, Y.; Ré, C.; Barrett, C.; et al. 2024b. H2o: Heavy-hitter oracle for efficient generative inference of large language models. *Advances in Neural Information Processing Systems*, 36.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.
- Zhuang, J.; Hu, J.; Mu, L.; Hu, R.; Liang, X.; Ye, J.; and Hu, H. 2025. FALIP: Visual Prompt as Foveal Attention Boosts CLIP Zero-Shot Performance. In *European Conference on Computer Vision*, 236–253. Springer.