

A-VL: Adaptive Attention for Large Vision-Language Models

Junyang Zhang¹, Mu Yuan¹, Ruiguang Zhong², Puhan Luo¹, Huiyou Zhan¹, Ningkang Zhang¹,
Chengchen Hu², Xiang-Yang Li^{1*}

¹University of Science and Technology of China, Hefei, China

²NIO Inc., Shanghai, China

{zhangjunyang, ym0813, luopuhan, zhangnk}@mail.ustc.edu.cn, taurus.zhong@nio.com, zhanhuiyou@foxmail.com,
huc@ieee.org, xiangyangli@ustc.edu.cn

Abstract

The Large Vision-Language Model (LVLM) integrates computer vision and natural language processing techniques, offering substantial application potential. However, these models demand extensive resources during inference. Adaptive attention techniques can dynamically reduce computational redundancy and thus improve efficiency. Although current adaptive attention methods significantly reduce the memory requirements of Transformer-based language models, they are not tailored for LVLMs. We observe that LVLMs generate responses from both remote image tokens and local text tokens, and different modalities have different attention patterns. This observation inspires us to manage the attention for each modality separately. Specifically, for visual input, we store the cache of potentially useful information but only compute the most critical parts. For language input, we care more about local information. Based on our observation and analysis of vision-language attention patterns, we develop A-VL, a plug-and-play adaptive attention tailored for LVLM inference. Extensive evaluations on three vision-language tasks and five datasets show the effectiveness of our designs. Our approach A-VL outperforms existing adaptive attention methods in reducing memory usage and computational load without compromising performance.

Introduction

In recent years, large vision-language models (LVLMs) have exhibited exceptional capabilities in vision-language understanding and reasoning (Liu et al. 2024a; Bai et al. 2023; Wang et al. 2023). LVLMs are increasingly integrated into many applications, such as personal intelligent assistants and vehicle cockpit systems. Despite their impressive performance, deploying LVLMs in real-world systems still faces practical challenges due to expensive resource overheads. Consequently, reducing computational overheads and improving inference speed have become critical problems.

In LVLM inference, visual and text inputs are separately encoded and fed into the pretrained large language model (LLM). As each token generated through LLM autoregression depends on all preceding tokens, this process consumes considerable time and memory resources. This issue is exacerbated in LVLMs since high-resolution images lead to

rapidly expanding token sequences. Recent studies explored adaptive attention and KV cache compression techniques for single-modal language models, achieving significant advancements. For instance, StreamingLLM (Xiao et al. 2023) maintains only fixed cache positions, while H₂O (Zhang et al. 2024b) preserves the most important caches based on the accumulated attention. These methods reduce memory usage and enhance the inference speed of LLMs. However, these methods are designed specifically for single-modal language models. Our experiments demonstrate the uniqueness of the attention pattern of LVLMs, which has inspired us to develop an adaptive attention scheme tailored to the unique needs of different modalities within LVLMs.

Adaptive attention methods must address two critical aspects: identifying redundancy and maintaining performance. To develop an effective adaptive attention method for LVLMs, we must first examine the specific characteristics of attention in LVLMs. Subsequently, we can design a strategy based on these characteristics to reduce memory usage and computational demands while maintaining performance. In this paper, we reveal that different modalities of LVLMs exhibit different attention patterns, necessitating separate analysis and processing for each modality within LVLM’s adaptive attention mechanism. Vision inputs, though positioned at the beginning of the sequence, consistently receive a substantial proportion of attention during generation. In contrast, text inputs at the sequence’s end exhibit rapid attention decay. Vision inputs also show patterns of attention sparsity and drift. Specially, the attention at the token level for vision inputs presents a surprisingly sparse distribution. Additionally, we find that some useful vision caches are not required at every step. We observe a gradual shift in attention among these useful vision caches.

To address this issue, we develop distinct methods for different modalities. In response to the sparsity and attention drift observed in vision attention, we categorize attention importance and dynamically compute the most critical parts of the KV cache. Specifically, we first preserve caches likely to be useful in subsequent inference and evict those deemed useless. We then select the most critical caches in the current step from these useful caches for computation. Simultaneously, we regularly update this critical cache to ensure accurate computations. For text attention, which rapidly loses attention, we concentrate on the local cache and retain only

*Xiang-Yang Li is the corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the most essential remote text caches. Notably, our approach does not require fine-tuning the model and offers plug-and-play compatibility with LVLMs.

Our main contributions are listed as follows:

- We investigate and identify different attention patterns across different modalities within LVLMs.
- We propose A-VL, a plug-and-play adaptive attention method for LVLMs, designed to significantly reduce memory usage and computational load during inference without compromising performance.
- We evaluate the effectiveness of A-VL across three representative vision-language tasks on different LVLMs. Experimental results show that A-VL outperforms existing adaptive attention methods in memory and computational efficiency.

Background and Related Work

Large Vision-Language Modal.

Large vision-language models (LVLMs) integrate the capabilities of visual encoders and large language models to understand multimodal information and generate textual responses related to images. LVLMs such as LLaVA (Liu et al. 2024b), Qwen-VL (Bai et al. 2023) and CogVLM (Wang et al. 2023) have demonstrated impressive performance across various vision-language tasks. The LVLM typically consists of a visual encoder, an adapter, and a pre-trained LLM. The visual encoder processes and converts vision inputs into hidden features. Then the adapter transforms features into tokens for LLM. Text inputs are also converted into tokens through the embedding model. The LLM generates responses based on the combined visual and textual tokens. The inference of an LLM is divided into two phases: the prefill phase and the decode phase. During the prefill phase, all input tokens are processed and KV caches are simultaneously generated. The KV caches store the intermediate results of the processed tokens. In the decode phase, the model computes using the newly generated token and the KV cache from all previous tokens. As the sequence of tokens extends, the size of the KV cache grows linearly, directly affecting the memory usage and computational demands during inference. However, the vision tokens from higher-resolution images output by the visual encoder significantly increases the length of token sequences. For instance, LLaVA 1.5 processes 576 image tokens, while LLaVA 1.6 increases this number to over 2000. This increase presents challenges to the deployment and utilization of LVLMs.

Inference Optimization for LLM and LVLM.

The decoder-only transformer-based LLM necessitates autoregressive iteration for output generation, where the generation of each token depends on all preceding tokens. This dependency significantly increases memory requirements and computational overhead. To address these challenges, existing methods are categorized into two main approaches: inference process optimization and computation

reduction. Inference process optimization includes strategies such as page attention (Kwon et al. 2023) and speculative sampling (Xia et al. 2024) to achieve flexible and efficient management of the entire inference process. Computation reduction involves methods like adaptive attention, early exit (Bae et al. 2023), and structured pruning (An et al. 2024). Our primary focus is on adaptive attention technology, which dynamically reduces computational redundancy according to inputs and has produced remarkable results in LLMs. StreamingLLM (Xiao et al. 2023) is designed to retain a fixed-position cache, which can train LLMs with long contexts on a limited cache window. Anagnostidis et al. (Anagnostidis et al. 2024) propose a new sigmoid function to dynamically sparse attention, but the model needs to be fine-tuned. FastGen (Ge et al. 2023) combines multiple strategies to retain recent caches, special characters, and caches that have been historically assigned heavier attention. H₂O (Zhang et al. 2024b) leverages historical accumulated attention and combined with the local cache for KV cache reduction, and has achieved remarkable results. Although these methods improve the inference efficiency of LLM, they are designed for single-modal text inputs. In addition, there are some studies that focus on enhancing the efficiency of transformer-based visual models. SPViT (Kong et al. 2022) prunes tokens in Vision Transformers (ViT), and PuMer (Cao, Paranjape, and Hajishirzi 2023) prunes and merge tokens in cross attention; however, these methods are designed for encoder models. FastV (Chen et al. 2024) is an optimization solution designed for autoregressive LVLMs, revealing that without using KV cache, LVLM can potentially use only 50% of tokens after the second layer. Nevertheless, the KV cache is an essential optimization tool for autoregressive inference. FastV has not been verified in inference with KV cache, because token pruning after the second layer may lead to the loss of some KV cache, making it impossible to select tokens in subsequent steps. And the pruning in FastV with the KV cache may lead to the loss of potentially important tokens in the subsequent steps.

Observations and Insights

In this section, we employ experiments to uncover key insights into the inference of large vision-language models, facilitating optimization of these processes.

Attention of Different Modalities

The large language model in LVLMs processes three types of input tokens: system input, image tokens and user instruction, as shown in Figure 2. System input typically comprises control information, such as role settings and task modes. Image tokens represent vision information encoded by the visual encoder. User instruction is the query posed by users about the image. Attention in LLM quantifies the relevance between two tokens by assigning weights. The attention directed toward image tokens in user queries implies the importance of these image tokens to the current question (Cao, Paranjape, and Hajishirzi 2023), which we call text-aware vision attention.

We utilize the attention score to demonstrate the attention

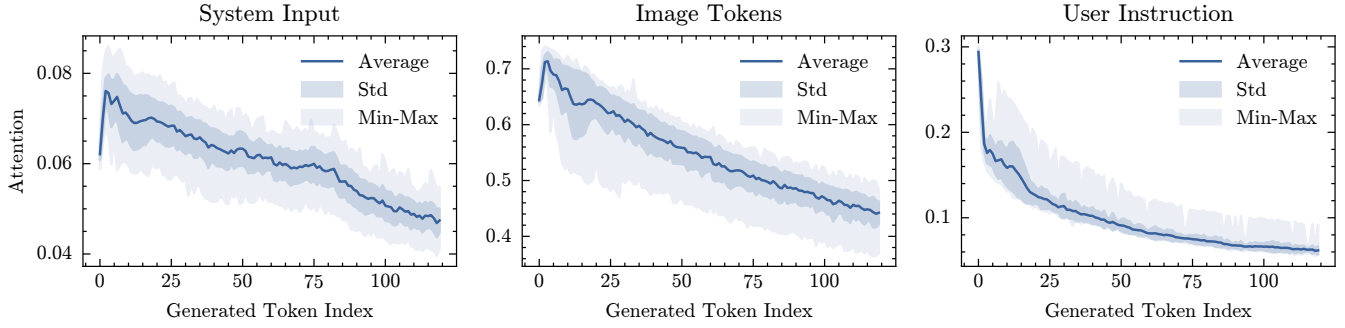


Figure 1: Variations in attention scores among different types of tokens during the decode phase.




Figure 2: The token sequence is composed of encoded inputs from different modalities.

allocated to input tokens during the generation of a new token. The attention score is calculated as the average of the multi-head attention weights. The formula of attention score is as follows:

$$s_t^l = \sum_{h=1}^H A_t^{l,h} / H, \quad (1)$$

where A represents the attention weights assigned from the last token to all preceding tokens; l denotes the index of the transformer decoder layers; and h indicates the attention head. To investigate attention changes in LLaVA, we randomly select 100 samples from the Flickr30k dataset (Young et al. 2014) that can output longer text sequences. We quantify the attention score changes in LLaVA-1.5 7B model across three token types during the decode phase. The results are depicted in Figure 1. Surprisingly, the attention allocated to system input and image tokens gradually decreases as the token sequence lengthens, while the attention to user instructions decreases sharply. Note that user instructions are often of similar length to system inputs or even longer, but still exhibit different attention patterns. This suggests that the remote system inputs and image tokens are crucial for inference. These results indicate that the attention allocated to each modality may need to be managed separately.

Heterogeneity of Vision Attention

We subsequently explore how vision attention changes across different decoder layers. For an intuitive visualization, we visualize the attention scores across various layers for image tokens during the inference with the LLaVA-1.5 7B model, as illustrated in Figure 3. The visualization shows that most image tokens are allocated little attention after first decoder and the vision attention is highly sparse. To quantitatively illustrate the ubiquity of this sparsity, we use the

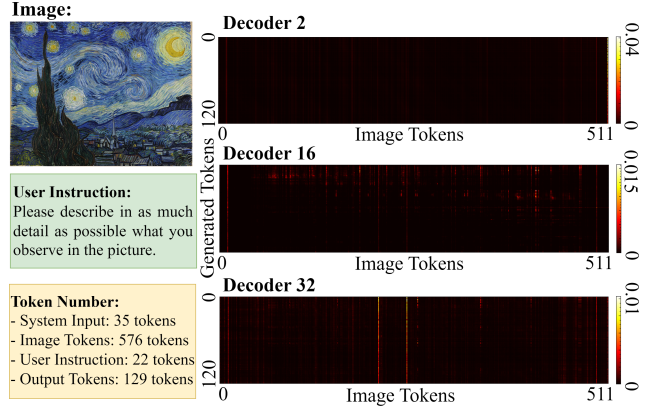


Figure 3: The vision attention during the generation of new tokens is very sparse.

LLaVA-1.5 7B model to evaluate the OCRVQA dataset and find that 22.88% of the tokens receive over 80% of the total attention in all layers. This indicates that image information is highly concentrated in a few tokens, resulting in inefficient inference. It is similar to the Pareto principle. Based on the observation that a small number of image tokens are allocated a large amount of attention, we define a metric termed the p -percentile concordance index (PPCI). This metric compares the correlation between two attention weights on the tokens, specifically focusing on key tokens. The PPCI measures the proportion of the top $p\%$ of tokens that are identical in two different attention weights. Specifically, for two sequences of attention scores, S_1 and S_2 , on the N tokens, the sets T_1 and T_2 represent the sets of top $p\%$ tokens in S_1 and S_2 respectively, sorted by their attention scores. The p -percentile concordance index (PPCI) is defined as follows:

$$p\% \text{ PPCI} = \frac{|T_1 \cap T_2|}{p\% \cdot N}. \quad (2)$$

We use LLaVA-1.5 7B model to evaluate the OCRVQA dataset, and analyze the PPCI between the attention across each decoder layer and that in the first decoder layer. The results are depicted in Figure 4. As illustrated by the figure, the 50% PPCI between the first decoder layer and subsequent layers is less than 70%, and the 50% PPCI between




Figure 4: The correlation of attention between the first decoder layer and each subsequent decoder layer.

the first layer and the last decoder layer is below 60%. We also calculate the 50% PPCI between the first three layers and the last layer, finding it to be less than 70% on average. Based on this observation, we refer to this phenomenon where different decoder layers focus on different image tokens as the *heterogeneity* of vision attention. Consequently, it is essential to design the adaptive attention method based on the attention distribution for each layer.

Continuity of Vision Attention

We subsequently analyze the attention from the temporal dimension with the LLaVA-1.5 7B model and the OCRVQA dataset. During the autoregressive generation of tokens, we observe a high correlation in attention within the same decoder layer across different steps, as depicted in Figure 5. The figure illustrates that during the generation of 20 new tokens (20 steps) in the same decoder layer, the $p\%$ PPCI in image tokens between each step and the first step. As depicted in the figure, each layer consistently allocates high attention to the top 50% of image tokens. For the top 30% image tokens, each layer allocates high attention in the short term, although this may shift after three steps. We refer to this short-term correlation of attention among image tokens within the same layer as the *continuity* of vision attention. This indicates that attention may gradually drift within partial image tokens. Given the gradual nature of these changes, we can reasonably assume that the set of core tokens remains stable in the short term.

Based on the heterogeneity and continuity of vision attention, we can select important image tokens for each layer according to the attention scores, thereby reducing computational load. In fact, we retain only the top 30% of image tokens for each decoder layer using the LLaVA-1.5 7B model and OCRVQA dataset to ensure that 99% of the samples generate three tokens identical to the original model output. However, for longer outputs, because attention drifts, relying on only 30% image tokens is not enough to predict. Thus, we can preserve potentially useful image tokens for future use, but compute only the most critical token at each step. This specific design will be elaborated in the following section.

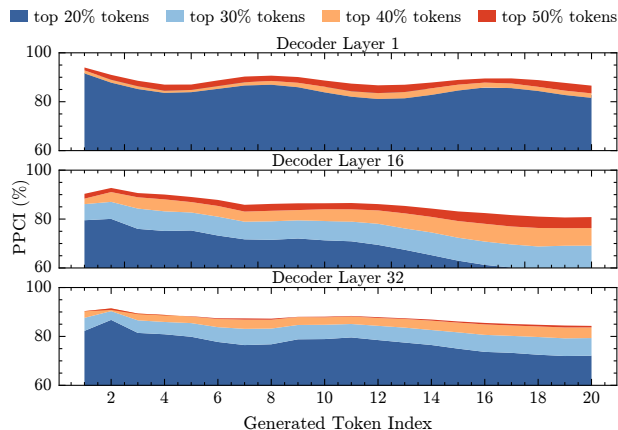


Figure 5: The correlation of attention between the first step and each subsequent step in the same decoder layer.

Locality of Text Attention

As illustrated in Figure 1, the attention of text tokens decays rapidly. Therefore, it is advisable to focus on local text attention and a limited number of critical remote text attention. The H₂O method can be precisely applied to achieve this objective. Specially, we maintain local text cache caches and evict useless text caches based on accumulated historical attention. We also analyze the attention patterns of Llama 2 (Touvron et al. 2023), a large language model. We observe that the attention patterns in Llama 2 are similar to those in the text token of LVLMs, both characterized by rapid attention decay. This similarity further underscores the suitability of employing an LLM’s adaptive attention method to the text tokens of LVLMs.

A-VL Design

Based on the above insights, we propose A-VL, a plug-and-play adaptive attention method tailored for LVLMs.

Adaptive Text-Aware Vision Attention

Decode phase. We design a hierarchical adaptive attention mechanism for vision attention based on the attention scores from the perspective of text tokens, as shown in Figure 6. In the prefill phase, we categorize image tokens into three sets: core, secondary, and minor, according to the attention score of the last text token. This classification uses the attention score from the text token to emphasize the semantic relationship between text and images. Specifically, after computing the attention score in each decoder layer, tokens within the top $C\%$ of scores are designated as core tokens, those within the top $S\%$ as secondary tokens, and the remainder as minor tokens. We ensure that $S > C$, thereby making core tokens a subset of secondary tokens. Minor tokens are immediately evicted to minimize KV cache usage. Secondary tokens, which may be useful in the future, are retained in memory; however, they are not all used in computations. Only core tokens are involved in the computations of subsequent attention modules. However, our observations suggest

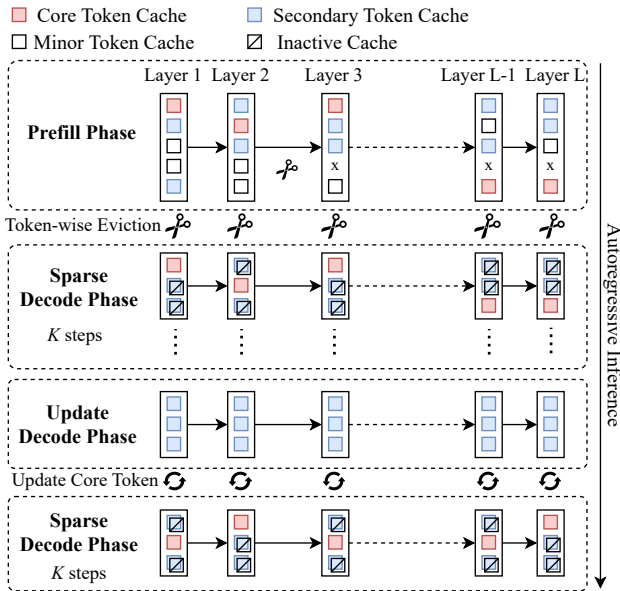


Figure 6: The design of adaptive vision attention.

that vision attention may shift gradually among secondary tokens, potentially altering the set of core tokens. Thus, every K steps, we utilize all secondary tokens for inference and compute the attention scores to update the core token set, ensuring that attention remains focused on the most significant tokens. We refer to this step as the update decode phase. This design is based on our insights into the continuity of vision attention, allowing the model to minimize cache usage and further reduce computational load of attention. It should be noted that our method is applied independently to each decoder layer, and the results are not utilized until running to the corresponding layer in the next step, allowing for parallel processing with the original model’s inference.

Prefill phase. The related work, FastV, has shown that inference without using KV cache may result in the deletion of half image tokens based on attention scores at each step. However, employing KV cache with this method leads to performance degradation, as potentially useful tokens are evicted, causing cache misses. Our experiments illustrate this phenomenon in evaluation section. Inspired by this study, we propose simplifying this technology for only use in the prefill phase. Specifically, we only retain $P\%$ image tokens based on attention score after the second decoder layer in the prefill phase. And then we use only the KV cache of these retained image tokens for inference in the decode phase. We can reduce the parameter P to prevent the performance from decreasing. This technology reduces computational load in the prefill stage and has been experimentally verified to be compatible with the above adaptive vision attention technology.

Adaptive Text Attention

Based on the insight of text attention, when generating new text tokens, both remote image tokens and local text tokens

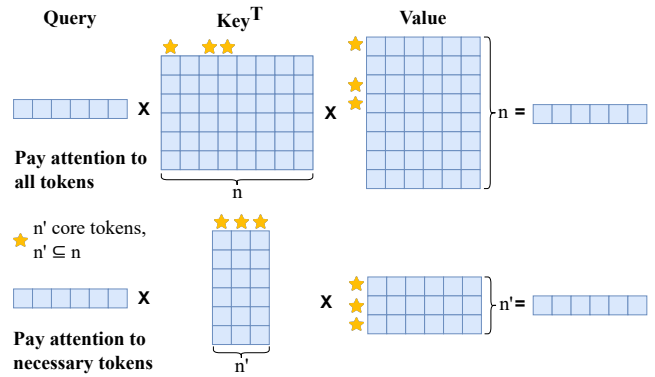


Figure 7: Adaptive vision attention uses only core KV cache. The details such as softmax are omitted in the figure.

are used to predict. Since image tokens are processed separately, we use H₂O’s accumulated attention method for text cache eviction. Specially, we set a text cache window size of $T\%$ of the maximum text token length. After computing attention weights at each layer, we separate the attention of vision and text, and store the accumulated text attention scores. If the text cache size exceeds the set text cache window size, we evict the redundant caches from the first half of the window based on these scores. This method, utilized by H₂O, has proven effective in LLM. We incorporate this approach into the text attention of LVLMs and verify its effectiveness through experiments. We list all the parameters involved in our method and their description in Table 1.

Implementation of A-VL

Our method is plug-and-play, requiring no fine-tuning of the original model. Cache selection and eviction can be processed in parallel with the original model inference, and the results are not utilized until the next step. Besides, in our adaptive vision attention method, core caches are selected from secondary caches for computation, as depicted in Figure 7. Slicing the cache before matrix multiplication introduces a performance bottleneck, because the core caches are not stored contiguously in memory (Liu et al. 2023b). Sometimes increasing latency of slicing beyond original matrix multiplication. To address this, we develop a specialized CUDA operator that allows direct multiplication with selected rows or columns of the second matrix, eliminating the need for slicing. This CUDA operator employs block processing and shared memory to enhance speed further.

| Param. | Description |
|--------|--|
| S% | The proportion of secondary image tokens |
| C% | The proportion of core image tokens |
| K | Update the core token set every K steps |
| P% | Image tokens retained in prefill phase |
| T% | The proportion of retained text cache |

Table 1: Description of parameters in A-VL.

| Model | Method | DocVQA | TextVQA | OCRBench | Nocaps | Flickr30k | VQAv2 | Avg. |
|---------------|------------|--------|---------|----------|--------|-----------|-------|--------------|
| LLaVA 1.5 7B | Original | 31.20 | 48.67 | 20.60 | 101.88 | 73.81 | 76.52 | 58.78 |
| | FastV | 27.35 | 47.16 | 19.30 | 101.06 | 72.99 | 74.86 | 57.12 |
| | H2O | 30.34 | 48.22 | 20.40 | 102.23 | 73.53 | 76.49 | 58.54 |
| | Our | 31.26 | 48.44 | 20.90 | 102.67 | 73.81 | 76.52 | 58.93 |
| LLaVA 1.5 13B | Original | 35.39 | 52.86 | 22.50 | 107.53 | 79.69 | 77.97 | 62.66 |
| | FastV | 30.27 | 50.64 | 20.60 | 106.41 | 78.88 | 76.65 | 60.57 |
| | H2O | 33.86 | 52.72 | 22.30 | 107.11 | 79.15 | 77.96 | 62.18 |
| | Our | 35.22 | 52.97 | 22.40 | 107.52 | 79.79 | 77.97 | 62.65 |
| LLaVA 1.6 7B | Original | 74.43 | 64.65 | 53.10 | 88.18 | 68.45 | 80.01 | 71.47 |
| | FastV | 67.69 | 63.15 | 48.30 | 86.21 | 66.73 | 79.58 | 68.61 |
| | H2O | 72.01 | 64.15 | 51.10 | 88.63 | 68.61 | 79.97 | 70.75 |
| | Our | 74.46 | 64.69 | 52.00 | 88.14 | 68.38 | 80.02 | 71.28 |
| LLaVA 1.6 13B | Original | 77.23 | 67.06 | 53.90 | 88.12 | 66.67 | 80.93 | 72.32 |
| | FastV | 70.26 | 65.14 | 49.10 | 87.78 | 65.96 | 80.48 | 69.79 |
| | H2O | 75.85 | 66.58 | 53.00 | 88.33 | 66.71 | 80.91 | 71.90 |
| | Our | 77.30 | 66.97 | 53.90 | 87.68 | 66.52 | 80.94 | 72.22 |
| Qwen-VL | Original | 68.65 | 61.10 | 49.10 | 58.92 | 59.01 | 79.17 | 62.66 |
| | FastV | 51.45 | 52.85 | 38.60 | 58.73 | 57.84 | 76.05 | 55.92 |
| | H2O | 65.84 | 60.59 | 46.60 | 59.06 | 59.39 | 79.16 | 61.77 |
| | Our | 68.09 | 60.66 | 49.10 | 58.46 | 58.88 | 79.16 | 62.39 |

Table 2: Performance comparison of different models, different tasks and different methods. Except for the original model, each method only retains 50% of the KV cache. The first and second highest average metrics are bolded.




Figure 8: The latency to calculate 30% caches in attention.

We measure the latency during the process in Figure 7 with PyTorch on an NVIDIA A40 GPU under the configure of LLaVA-1.6 7B. Latency variations across different batch sizes are illustrated in Figure 8. It is evident that slicing is slow and negates the benefits of reduced computation. Our CUDA operator demonstrates effective acceleration.

Evaluation

In this section, we demonstrate the effectiveness of our method through experiments.

Evaluation Setup

We implement our evaluation system based on the multi-modal evaluation suite, LMMs-Eval (Zhang et al. 2024a).

Evaluation Tasks We employ multiple vision-language tasks to evaluate our method. Since our method primarily optimizes the KV cache in the decode phase, we do

not choose any multiple-choice datasets that generate only one token during the prefill phase. For example, SEED-Bench (Li et al. 2024) is a fully multiple-choice benchmark, so we did not select it. The tasks and datasets used in our evaluation are as follows:

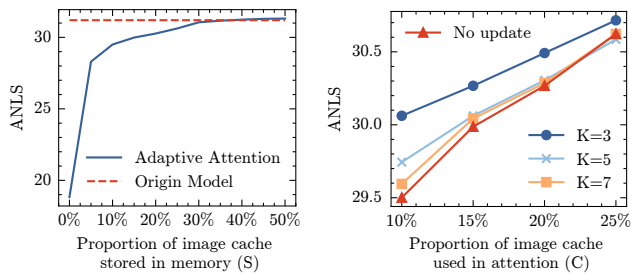
- **Image Caption.** This task involves automatically generating textual descriptions for visual content. We employ two datasets, Nocaps (Agrawal et al. 2019) and Flickr30k (Plummer et al. 2015). The metric is the CIDEr score (Vedantam, Lawrence Zitnick, and Parikh 2015).
- **Visual Question Answering (VQA).** This task requires models to answer questions based on visual information from images. We use three datasets, DocVQA (Mathew, Karatzas, and Jawahar 2021), TextVQA (Singh et al. 2019) and VQAv2 (Goyal et al. 2017), measuring performance with the ANLS (Biten et al. 2019) and accuracy metrics.
- **Optical Character Recognition (OCR).** OCR involves recognizing text content within images. We select the OCRBench dataset (Liu et al. 2023a), which is specifically designed for LVLMS. The metric is the score designed in OCRBench.

Among these tasks, some emphasize the overall content of the image (e.g., image caption), and some focus on fine-grained details (e.g., OCR). This distinction enables the evaluation across varying levels of granularity.

Models. We use mainstream LVLMS and their different sizes, including LLaVA 1.5 7B, LLaVA 1.5 13B, LLaVA 1.6 7B, LLaVA 1.6 13B and Qwen-VL.

| Model | Settings | | | | | Stored Cache | Used Cache | KV Cache Memory | DocVQA | TextVQA | OCRBench | Nocaps |
|------------------|----------|----|---|----|----|--------------|------------|-----------------|--------|---------|----------|--------|
| | S% | C% | K | P% | T% | | | | | | | |
| LLaVA 1.5 7B | - | - | - | - | - | 100% | 100% | 339 MB | 31.20 | 48.67 | 20.60 | 101.88 |
| | 40 | - | - | - | - | 46% | 46% | 156 MB | 31.26 | 48.44 | 20.90 | 102.67 |
| | 40 | 30 | 3 | - | - | 46% | 39% | 156 MB | 31.02 | 48.44 | 21.00 | 102.67 |
| | 45 | 30 | 3 | 90 | - | 50% | 40% | 170 MB | 31.02 | 48.46 | 20.70 | 102.34 |
| | 45 | 30 | 3 | 90 | 50 | 48% | 38% | 163 MB | 30.92 | 48.46 | 20.50 | 102.22 |
| LLaVA 1.5 13B | - | - | - | - | - | 100% | 100% | 529 MB | 35.39 | 52.86 | 22.50 | 107.53 |
| | 40 | - | - | - | - | 46% | 46% | 243 MB | 35.37 | 52.86 | 22.20 | 107.61 |
| | 45 | 30 | 3 | - | - | 50% | 40% | 265 MB | 35.22 | 52.97 | 22.40 | 107.52 |
| | 45 | 30 | 3 | 90 | - | 50% | 40% | 265 MB | 35.46 | 52.83 | 22.30 | 107.60 |
| | 45 | 30 | 3 | 90 | 60 | 48% | 38% | 254 MB | 35.43 | 52.84 | 22.30 | 107.59 |
| LLaVA 1.6 7B | - | - | - | - | - | 100% | 100% | 1179 MB | 74.43 | 64.65 | 53.10 | 88.18 |
| | 40 | - | - | - | - | 43% | 43% | 507 MB | 74.29 | 64.76 | 51.20 | 88.14 |
| | 45 | 30 | 3 | 90 | - | 47% | 36% | 554 MB | 74.46 | 64.71 | 51.90 | 87.43 |
| LLaVA 1.6 13B | 45 | 30 | 3 | 90 | 70 | 45% | 35% | 531 MB | 74.44 | 64.69 | 52.00 | 87.45 |
| | - | - | - | - | - | 100% | 100% | 1841 MB | 77.23 | 67.06 | 53.90 | 88.12 |
| LLaVA 1.6 13B | 40 | - | - | - | - | 43% | 43% | 792 MB | 77.30 | 66.97 | 53.90 | 87.68 |
| | 45 | 30 | 3 | 90 | - | 47% | 36% | 865 MB | 76.97 | 67.06 | 54.40 | 87.59 |
| | 45 | 30 | 3 | 90 | 70 | 45% | 35% | 828 MB | 76.97 | 67.06 | 54.50 | 87.59 |

Table 3: Performance under different settings with our method. KV Cache Memory refers to the memory usage when the batch size is one. As the batch size increases, memory consumption scales proportionally, resulting in greater memory savings.



(a) different secondary cache. (b) different core cache.

Figure 9: Performance under different parameters S, C, K.

Baselines. We compare the output of the original model, FastV designed for LVLMS, and H₂O designed for LLMs. For the H₂O method, set the cache window to 75% recent and 25% heavy hitter. For FastV, in order to compare performance with the KV cache, we delete redundant tokens according to the FastV method in the prefill phase, and use the generated KV cache for inference in the decode phase.

Performance Guarantee

We evaluate the performance of various models across different tasks, as shown in Table 2. FastV, H₂O, and our method retains 50% of the KV cache. Our method achieves nearly lossless performance using only 50% of the KV cache. And our average performance is better than the other two baselines. FastV shows reduced performance when utilizing the KV cache, due to the loss of potentially useful tokens during token pruning. The H₂O method designed for LLMs shows effectiveness, but its average performance is still inferior to our method. We observe that while H₂O can




Figure 10: Performance under different parameters P and T.

sometimes maintain performance like our method in image caption tasks that focus on the overall image information, our method demonstrates better performance in OCR and VQA tasks, which require attention to finer details. This indicates that reducing the number of image tokens may have minimal impact on tasks emphasizing general image content, but for detail-oriented tasks, our method, specifically designed for LVLMS is essential. Our method demonstrates good results across various models, confirming its effectiveness. In all methods, 50% of the KV Cache is retained, resulting in similar decoding inference latency. Detailed memory savings from the KV Cache are shown in Table 3.

Memory Usage of KV Cache

We first analyze the relationship between KV cache memory usage and performance across various parameters with LLaVA-1.5 7B model and DocVQA dataset. For vision attention, we evaluate the performance with different parameters, as shown in Figure 9. Figure 9a illustrates that retaining 40% of the image cache as the secondary cache guarantees no performance loss, and retaining only 20% results can

ensure performance decline within 3%. According to Figure 9b, computing with only the core cache yields better results than simply eliminating these caches. Additionally, increasing the update frequency (lowering K) ensures that the core cache consistently contains the most critical elements. These results confirm the insights previously discussed. Performance under varying parameters P and T are presented in Figure 10. Specifically, maintaining over 60% of the text cache (Parameter T) ensures stable performance. In contrast, retaining less than 80% of the image tokens during the prefill phase (Parameter P) leads to performance degradation. This underscores the underperformance of the FastV method with the KV cache, as previously analyzed. However, integration with our method can reduce computational load by nearly 10% during the prefill phase.

Our comprehensive analysis of different models and datasets under various settings is shown in Table 3. As computational and memory demands decrease, the performance decreases slightly. Nevertheless, our method maintains near-lossless performance with less than 50% cache stored and only 35% used in computations. Notably, our method does not utilize the entire stored cache for computations, thus reducing the computational load compared to other methods, offering a unique advantage not provided by the baselines.

Decoder Inference Latency

Our method utilizes fixed eviction ratios for each layer, facilitating batch inference. We measure the inference latency of the Transformer Decoder in LLaVA-1.6 7B under various batch sizes on NVIDIA A40. By compressing the KV cache to 50%, we achieve a 1.8x increase in decoder inference. Furthermore, by utilizing our CUDA operator to only calculate the core cache, we can further enhance the speed by an additional 1.1x. Consequently, the total decoder latency required is only 50.5% of the original. It is evident that our method can substantially reduce inference latency. While our custom CUDA operator does increase speed, the improvement is less pronounced than in isolated tests due to the influence of other operations in PyTorch on its performance. We plan to further improve our CUDA operator in future work.

Conclusion

In this paper, we propose A-VL, a plug-and-play adaptive attention method tailored specifically for large vision-language models. Our analysis of attention patterns across different modalities indicates that remote vision tokens consistently receive high attention, while textual tokens exhibit rapid attention decay. Moreover, vision attention patterns are characterized by sparsity and drift, enabling dynamic selection of the most critical cache for computation. Consequently, we design the adaptive attention method tailored to each modality separately, which reduces the KV cache stored and concentrates computational load on the most critical cache. Our experiments confirm the efficacy of A-VL, demonstrating substantial reductions in memory usage and computation load without compromising performance.

Acknowledgments

The research is partially supported by USTC-NIO Smart Electric Vehicle Joint Lab, National Key R&D Program of China under Grant No. 2021ZD0110400, Innovation Program for Quantum Science and Technology 2021ZD0302900 and China National Natural Science Foundation with No. 62132018, 62231015, 623B2093, "Pioneer" and "Leading Goose" R&D Program of Zhejiang, 2023C01029 and 2023C01143.

References

- 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.
- An, Y.; Zhao, X.; Yu, T.; Tang, M.; and Wang, J. 2024. Fluctuation-based adaptive structured pruning for large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 10865–10873.
- Anagnostidis, S.; Pavlo, D.; Biggio, L.; Noci, L.; Lucchi, A.; and Hofmann, T. 2024. Dynamic context pruning for efficient and interpretable autoregressive transformers. *Advances in Neural Information Processing Systems*, 36.
- Bae, S.; Ko, J.; Song, H.; and Yun, S.-Y. 2023. Fast and Robust Early-Exiting Framework for Autoregressive Language Models with Synchronized Parallel Decoding. arXiv:2310.05424.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond. *arXiv preprint arXiv:2308.12966*.
- Biten, A. F.; Tito, R.; Mafla, A.; Gomez, L.; Rusiñol, M.; Jawahar, C.; Valveny, E.; and Karatzas, D. 2019. Scene Text Visual Question Answering. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 4290–4300.
- Cao, Q.; Paranjape, B.; and Hajishirzi, H. 2023. PuMer: Pruning and merging tokens for efficient vision language models. *arXiv preprint arXiv:2305.17530*.
- Chen, L.; Zhao, H.; Liu, T.; Bai, S.; Lin, J.; Zhou, C.; and Chang, B. 2024. 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*.
- Ge, S.; Zhang, Y.; Liu, L.; Zhang, M.; Han, J.; and Gao, J. 2023. Model tells you what to discard: Adaptive kv cache compression for llms. *arXiv preprint arXiv:2310.01801*.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, 6325–6334. IEEE Computer Society.
- Kong, Z.; Dong, P.; Ma, X.; Meng, X.; Niu, W.; Sun, M.; Shen, X.; Yuan, G.; Ren, B.; Tang, H.; et al. 2022. Spvit: Enabling faster vision transformers via latency-aware soft token pruning. In *European conference on computer vision*, 620–640. Springer.

- Kwon, W.; Li, Z.; Zhuang, S.; Sheng, Y.; Zheng, L.; Yu, C. H.; Gonzalez, J.; Zhang, H.; and Stoica, I. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, 611–626.
- Li, B.; Ge, Y.; Ge, Y.; Wang, G.; Wang, R.; Zhang, R.; and Shan, Y. 2024. SEED-Bench: Benchmarking Multimodal Large Language Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13299–13308.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024a. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024b. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Liu, Y.; Li, Z.; Yang, B.; Li, C.; Yin, X.; Liu, C.-I.; Jin, L.; and Bai, X. 2023a. On the hidden mystery of ocr in large multimodal models. *arXiv preprint arXiv:2305.07895*.
- Liu, Z.; Wang, J.; Dao, T.; Zhou, T.; Yuan, B.; Song, Z.; Shrivastava, A.; Zhang, C.; Tian, Y.; Re, C.; et al. 2023b. Deja vu: Contextual sparsity for efficient llms at inference time. In *International Conference on Machine Learning*, 22137–22176. PMLR.
- Mathew, M.; Karatzas, D.; and Jawahar, C. 2021. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 2200–2209.
- 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.
- Singh, A.; Natarjan, V.; Shah, M.; Jiang, Y.; Chen, X.; Parikh, D.; and Rohrbach, M. 2019. Towards VQA Models That Can Read. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8317–8326.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Vedantam, R.; Lawrence Zitnick, C.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4566–4575.
- 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*.
- Xia, H.; Yang, Z.; Dong, Q.; Wang, P.; Li, Y.; Ge, T.; Liu, T.; Li, W.; and Sui, Z. 2024. Unlocking efficiency in large language model inference: A comprehensive survey of speculative decoding. *arXiv preprint arXiv:2401.07851*.
- Xiao, G.; Tian, Y.; Chen, B.; Han, S.; and Lewis, M. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- Young, P.; Lai, A.; Hodosh, M.; and Hockenmaier, J. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2: 67–78.
- Zhang, K.; Li, B.; Zhang, P.; Pu, F.; Cahyono, J. A.; Hu, K.; Liu, S.; Zhang, Y.; Yang, J.; Li, C.; et al. 2024a. LMMs-Eval: Reality Check on the Evaluation of Large Multimodal Models. *arXiv preprint arXiv:2407.12772*.
- 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.