

AirCache: Activating Inter-modal Relevancy KV Cache Compression for Efficient Large Vision-Language Model Inference

Kai Huang, Hao Zou, Bochen Wang, Ye Xi, Zhen Xie, Hao Wang
Alibaba Group

Abstract

Recent advancements in Large Visual Language Models (LVLMs) have gained significant attention due to their remarkable reasoning capabilities and proficiency in generalization. However, processing a large number of visual tokens and generating long-context outputs impose substantial computational overhead, leading to excessive demands for key-value (KV) cache. To address this critical bottleneck, we propose AirCache, a novel KV cache compression method aimed at accelerating LVLMs inference. This work systematically investigates the correlations between visual and textual tokens within the attention mechanisms of LVLMs. Our empirical analysis reveals considerable redundancy in cached visual tokens, wherein strategically eliminating these tokens preserves model performance while significantly accelerating context generation. Inspired by these findings, we introduce an elite observation window for assessing the importance of visual components in the KV cache, focusing on stable inter-modal relevancy modeling with enhanced multi-perspective consistency. Additionally, we develop an adaptive layer-wise budget allocation strategy that capitalizes on the strength and skewness of token importance distribution, showcasing superior efficiency compared to uniform allocation. Comprehensive evaluations across multiple LVLMs and benchmarks demonstrate that our method achieves comparable performance to the full cache while retaining only 10% of visual KV cache, thereby reducing decoding latency by 29% to 66% across various batch size and prompt length of inputs. Notably, as cache retention rates decrease, our method exhibits increasing performance advantages over existing approaches.

1. Introduction

The past few years have witnessed remarkable advancements in Large Visual-Language Models (LVLMs) [1, 6, 16, 22, 23, 32, 35, 36], both in research and practical applications. Although LVLMs exhibit increasingly advanced visual processing capabilities, such as handling

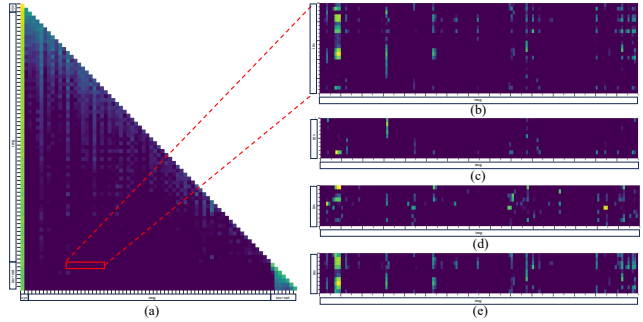


Figure 1. Illustration of our motivation. (a) An example of attention maps during the decoding process of LLaVA-OV-7B [16]. (b) Partial visual attention maps obtained using all text tokens as the observation window. (c) Using the last 16 text tokens as the observation window. (d) Using the last 16 visual tokens as the observation window. (e) Using the proposed method as the observation window, demonstrating more consistent attention distribution.

high-resolution, multiple images, and video sequences, these improvements entail substantial computational costs. The exponential increase in visual tokens, coupled with the demands of long-context generation tasks, results in unsustainable memory overhead from key-value (KV) cache storage [30]. This not only significantly increases GPU memory consumption but also severely degrades computational efficiency due to heightened memory bandwidth contention.

Recent approaches [5, 13, 14, 19, 39, 42] aim to prune visual tokens during inference, which can be categorized into token pruning and KV cache compression. Token pruning [11, 13, 14, 18, 20, 39] involves reducing visual tokens during the prefill phase, which decreases both the number of tokens processed in subsequent transformer layers and the tokens stored in the KV cache. Although it significantly enhances inference speed, the aggressive elimination of visual tokens during prefill leads to a substantial loss of visual information, severely degrading model performance. KV cache compression [19, 21, 25, 34, 42] reduces token solely during the decoding phase by pruning stored data in the KV cache. This technique has been extensively validated in LLMs. Since all tokens undergo a complete forward pass and the causal attention mechanism in LLMs differentiates the importance distribution among tokens, selec-

tively deleting certain tokens has minimal impact on model performance. These methods typically assess the importance of visual tokens by evaluating their strength or sparsity within the attention maps computed under the causal attention mechanism. Given the strict unidirectionality of causal attention, using cross-modality attention scores between instruction tokens and visual tokens provides a more comprehensive evaluation of visual token importance than using full attention scores. Moreover, we observe significant variations in the distribution of visual token strength across different layers, indicating that uniformly allocating the budget is not the optimal strategy.

To address the aforementioned issues, we propose AirCache, a novel KV cache compression method designed for efficient LVLMS inference. AirCache comprises two main components: visual token importance scoring and layer-wise KV cache budget allocation. To effectively evaluate the importance of visual tokens in the KV cache, we introduce an elite observation window that selects critical instruction tokens by leveraging self-attention scores among instruction tokens rather than using all or continuous local instruction tokens. As illustrated in Figure 1, compared to mainstream methods that assess visual token importance based on observation windows guided by all or partial text instruction tokens, the elite observation window demonstrates stronger consistency. Most text tokens within the elite observation window tend to provide more similar evaluations for the same visual token, thereby enhancing the effectiveness and stability of subsequent voting-based ranking. We further propose to quantify the compression budget for different layers from two perspectives. The first aspect is the emphasis a layer places on visual information, measured by the sum of attention allocated to all visual tokens, as well as the strength of importance score distribution. Additionally, analysis of attention distributions across layers reveals a distinct head effect, where a few visual tokens receive high importance scores, while the majority exhibit mediocre and average importance characteristics. Our experimental results show that preserving only 10% of the important visual KV cache results in less than 1% average model performance degradation, corroborating the significant head effect in the importance distribution of visual tokens. Combining these two aspects, we allocate the budget for visual KV cache based on both the strength and skewness of importance scores distributions.

Our primary contributions include: (1) We conduct a comprehensive and detailed exploration of the differences in inter-modal relevancy patterns based on attention interaction in evaluating visual KV cache importance. We introduce an elite observation window with carefully selected key textural instruction tokens to achieve more effective and stable importance evaluations. (2) We propose to quantify the differentiated compression budget allocation across dif-

ferent layers using the strength and skewness of the importance scores distribution, based on the prominence of visual information and the efficacy of attention allocation. (3) Experimental results on widely used LVLMS and benchmark datasets demonstrate that the proposed method achieves superior performance compared to other existing methods.

2. Related Work

Visual Token Pruning. Visual token pruning [4, 5, 12, 13, 20, 39] accelerates model inference by reducing visual tokens during the prefill phase. BLIP-2 [17] and Qwen-VL [1] employ cross-attention with learnable query embeddings to aggregate input visual tokens, while HoneyBee [4] and MobileVLM [7] use convolutional neural networks. These approaches necessitate model retraining, which significantly increases migration costs. In contrast, methods like [5, 39] introduce no extra parameters but leverage attention maps to score and evict visual tokens based on importance. IVTP [13] implements two-stage pruning by deleting tokens in both Vision Transformer (ViT) and LLM, enabling text-guided pruning through instruction integration. However, it is important to note that visual token pruning inevitably leads to visual information loss by discarding tokens prior to or within early layers, thereby substantially degrading model performance.

KV cache compression. KV cache strategies accelerate inference by storing the key and value states of previous tokens, thereby avoiding redundant computations during decoding. However, as the number of tokens increases, the pressure on memory and bandwidth also rises, limiting the model’s application in long content understanding or generation. To address this challenge, mainstream methods primarily focus on two directions: intra-layer KV cache pruning [15, 29, 40] and inter-layer compression budget allocation [21, 38]. H2O [42] and SnapKV [19] propose evicting tokens in the KV cache using global and local cumulative attention scores, while PyramidKV [3] constructs a pyramid-like hierarchical budget allocation to achieve more refined KV cache compression management. The input to LVLMS is a cross-modal construct primarily composed of visual tokens with auxiliary text tokens, necessitating additional consideration of the impact of different modality differences when extending KV cache compression to LVLMS. Elastic [25] merges non-critical tokens with important tokens, PrefixKV [34] introduces an adaptive layer-wise KV retention recipe based on binary search for maximal preservation of contextual information. VL-Cache [33] utilizes visual token sparsity for layer-wise budget allocation and token eviction through dedicated scoring policies.

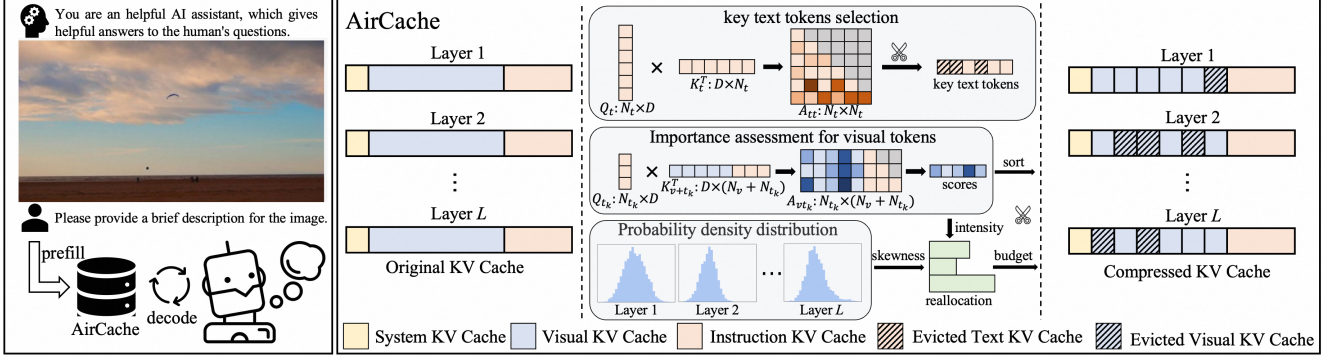


Figure 2. The left panel illustrates the inference workflow of standard LLMs. The proposed method integrates during the KV cache storage process following the prefill stage, maintaining compatibility with mainstream LLMs architecture. The right panel presents the overview of AirCache. We first employ self-attention across the instruction text tokens to identify the key text tokens, establishing an elite observation window for assessing visual token significance. We further reallocate the compression budgets between layers based on the strength and skewness of the importance distribution. Finally, pruning is applied to the KV cache according to the importance ranking of visual tokens in the KV Cache and the compression budget to obtain the final compressed KV cache.

3. Methodology

3.1. Preliminary

Similar to most LLMs, the inference process of LVLMs is primarily divided into prompt prefill and token decoding. **Prompt Prefill.** Given the hidden states $X \in \mathbb{R}^{N \times D}$ of the input prompt, where N is the total number of visual and text prompt tokens, and D denotes hidden dimension. We simplify the expression by omitting the indices of the head and layer of the hidden states. The query, key, and value states of the attention block can be represented as:

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_Q, \quad \mathbf{K} = \mathbf{X}\mathbf{W}_K, \quad \mathbf{V} = \mathbf{X}\mathbf{W}_V, \quad (1)$$

where $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{D \times D}$ are learnable projection matrices. Following the completion of the attention interaction, the key and value are stored in the KV cache to support subsequent token generation.

Token Decoding. In this stage, the model efficiently generates new tokens by iteratively utilizing and updating the KV cache. Specifically, during the i th iteration, we only need to compute the key and value associated with the newly generated token $\mathbf{x}_i \in \mathbb{R}^{1 \times D}$. By employing cache indexing, the model avoids recomputing the key and value for tokens that have already been processed. Subsequently, the KV cache is updated with the newly computed key and value for the latest token:

$$\mathbf{K} = \text{Concat}(\mathbf{K}, \mathbf{x}_i \mathbf{W}_K), \quad \mathbf{V} = \text{Concat}(\mathbf{V}, \mathbf{x}_i \mathbf{W}_V). \quad (2)$$

Although the KV cache mitigates the issue of redundant computation during the decoding process, its characteristic of linear growth in computational requirements relative to the length of the input sequence poses challenges regarding latency and memory usage, especially in scenarios involving longer inputs or outputs. This issue is further exacerbated as recent LVLMs tend to prioritize high-resolution, high-frame-rate visual inputs.

3.2. Overview

The AirCache is primarily divided into two components. The first component evaluates the significance of visual tokens based on the elite observation window. Section 3.3 details how to optimize the relevancy between visual and language modalities to guide the compression of the visual segment in the KV cache. The second component quantifies the KV cache budget requirements across various layers by analyzing the overall strength and skewness of the importance score distribution, thereby facilitating dynamic layer-wise KV cache budget allocation, as elaborated in Section 3.4. The architecture and algorithm underlying the method are illustrated in Figure 2 and Algorithm 1.

3.3. Elite Observation Window

In LVLMs, the presence of multiple modalities complicates the selection of observation windows, suggesting that adhering to the relevant practices established in LLMs may not be optimal. For instance, as illustrated in Figure 1, the visual token importance measures derived from the continuous last segment of the prompt exhibit poor consistency. This variability indicates that different text tokens within the observation window yield significantly diverse visual tokens importance distributions. Consequently, the hit rate statistics derived from the voting mechanism are prone to substantial noise, and the disparities among modalities may further amplify this noise interference. Based on this consideration, we propose refining the observation windows by incorporating carefully selected text tokens. By utilizing the attention hit rate within the text modality, we aim to enhance the consistency of visual token evaluation, thereby improving the performance of KV cache compression. To this end, we reorganize the hidden states of the input prompt by:

$$\mathbf{X} = \text{Concat}(\mathbf{X}_v, \mathbf{X}_t) \in \mathbb{R}^{(N_v + N_t) \times D}, \quad (3)$$

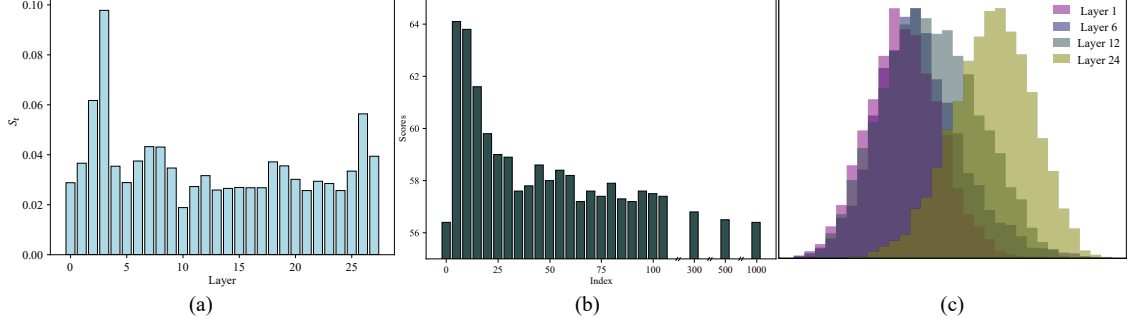


Figure 3. (a) The strength of the distribution of visual token importance scores across different layers, which is defined as the sum of all visual token importance scores. (b) Performance with selecting only one visual token each time based on the sorted results of visual token importance scores, the index of 0 indicates the setting of removing all visual tokens. (c) The distribution of visual token importance scores across different layers, with only selected layers displayed for clarity. All results are derived from LLaVA-OV-7B [16] on the ChatQA [26].

where \mathbf{X}_v and \mathbf{X}_t represent the hidden states corresponding to visual tokens and instruction text tokens, respectively. N_v and N_t denote the respective quantities of these tokens, satisfying the condition $N_v + N_t = N$. For convenience, the system prompts are omitted. We compute the text unimodal attention matrix composed of instruction text tokens as:

$$\mathbf{A}_{tt} = \text{Softmax} \left(\frac{\mathbf{Q}_t \mathbf{K}_t^T}{\sqrt{D}} \right) \in \mathbb{R}^{N_t \times N_t}, \quad (4)$$

where \mathbf{Q}_t and \mathbf{K}_t denote the query and key of the text tokens, respectively. Considering the characteristics of causal attention, we select the last token of the instruction text as reference and filter out key text tokens that receive high attention scores based on their interactions with other text tokens. Consistent with [42], we employ a relevance threshold to regulate the sparsity of key text tokens:

$$k = \{j \mid \mathbf{A}_{tt}[N_t - 1, j] \geq \alpha \cdot \max \mathbf{A}_{tt}[N_t - 1, :]\}, \quad (5)$$

where $\alpha \in [0, 1]$ represents the relevance threshold, and $j \in \{0, 1, \dots, N_t - 1\}$ is the index of the text token. When α is equal to 0, all text tokens are retained, and when α is equal to 1, only the text token with the highest score is preserved. We redefine the key for the aggregated visual tokens and key text tokens as \mathbf{K}_{vt_k} , along with the query corresponding to the key text tokens \mathbf{Q}_{t_k} , as follows:

$$\begin{aligned} \mathbf{Q}_{t_k} &= \mathbf{Q}_t[k] \in \mathbb{R}^{N_{t_k} \times D}, \\ \mathbf{K}_{vt_k} &= \text{Concat}(\mathbf{K}_v, \mathbf{K}_t[k]) \in \mathbb{R}^{(N_v + N_{t_k}) \times D}, \end{aligned} \quad (6)$$

where \mathbf{K}_v represents the key of the visual token, and N_{t_k} denotes the number of selected key text tokens. Furthermore, the attention matrix between the visual tokens and the key text tokens can be computed as follows:

$$\mathbf{A}_{vt_k} = \text{Softmax} \left(\frac{\mathbf{Q}_{t_k} \mathbf{K}_{vt_k}^T}{\sqrt{D}} \right) \in \mathbb{R}^{N_{t_k} \times (N_v + N_{t_k})}. \quad (7)$$

The elite observation window, comprising only these key text tokens, offers a more comprehensive and focused semantic representation compared to those formed by all text

tokens or fixed continuous sequences. It provides enhanced stability and accuracy in assessing the importance of visual tokens. The final importance scores of the visual tokens can be obtained by performing average pooling of the attention matrix along the text dimension:

$$\mathbf{I}_v = \frac{1}{N_{t_k}} \sum_{j=0}^{N_{t_k}-1} \mathbf{A}_{vt_k}[j, : N_v]. \quad (8)$$

By sorting the importance scores, the KV cache of visual tokens with lower rankings can be discarded in accordance with the compression budget. Furthermore, the elite observation window we employ exhibits reduced computational complexity, and the additional overhead of key text token selection is negligible in visual token-heavy LVLMs.

3.4. Layer-wise KV Cache Budget Allocation

We further investigate the characteristics of the visual token importance score distribution based on the elite observation window across various layers. Specifically, we consider the sum of the importance scores of all visual tokens as an indicator of the layer’s attention to visual information. As illustrated in Figure 3(a), there is a significant disparity in the attention given to visual tokens across different layers. This observation highlights the necessity of performing layer-wise token budget allocation, where layers dedicating more attention to visual information merit a larger token budget. As shown in Figure 3(b), we select one visual token at a time based on the importance scores. It can be observed that only approximately 10% of all visual tokens have a significantly positive impact on the final result. This suggests a head effect in the model’s attention distribution over visual tokens, with a small subset receiving higher attention, while the majority have lower and more balanced scores. Extending to the importance distribution across different layers in Figure 3(c), it shows significant differences in importance distribution among the layers, particularly regarding the head effect of visual tokens. Consequently,

we aim to allocate a larger token budget to the layers that demonstrate a pronounced head effect in their visual tokens importance distribution. This indicates a more accurate and nuanced understanding of the visual content by those layers.

As aforementioned, we quantify the allocation of visual token budget across different layers from two perspectives: the strength and skewness of the importance distribution. The strength of the importance distribution is calculated as:

$$s_t = \sum_{i=0}^{N_v-1} \mathbf{I}_v[i]. \quad (9)$$

We exclude the attention between text tokens and instead aggregate the attention scores between text tokens and all visual tokens to quantify the importance of visual information in each layer. The greater the importance, the more budget is allocated. Meanwhile, we use the skewness of importance distribution to assess how effectively each layer aligns with the distribution characteristics of attention to visual tokens:

$$s_k = \frac{N_v}{(N_v-1)(N_v-2)} \sum_{i=1}^{N_v} \left(\frac{\mathbf{I}_v[i] - \mu_{I_v}}{\sigma_{I_v}} \right)^3, \quad (10)$$

where μ_{I_v} and σ_{I_v} represent the mean and standard deviation of the importance as well as attention scores distribution of visual tokens, which essentially reflect the distribution of attention scores. Ultimately, the reallocated token budget integrates these factors into the original budget r :

$$\hat{r} = \frac{1}{2}(s'_t + s'_k)r. \quad (11)$$

where s'_t and s'_k represent the normalized strength and skewness across different layers, respectively.

4. Experiments

4.1. Setup

Models. The selection of LVLMS is primarily based on two criteria: model architecture and parameter size. In terms of architecture, the objective is to validate the effectiveness of the proposed method across various architectures, including LLaVA-OV-7B [16], InternVL2-8B [6], and Qwen2-VL-7B [35]. These models exhibit significant differences in visual token encoding, ViT, and LLM. For example, both LLaVA-OV-7B [16] and InternVL2-8B [6] utilize the AnyRes strategy with distinct ViT to increase the number of visual tokens, while Qwen2-VL-7B [35] supports native resolution with a large quantity of visual tokens. In terms of parameter size, models of InternVL2 series are selected for their diverse range of parameter scales, including InternVL2-1B [6], InternVL2-4B [6], and InternVL2-26B [6]. Due to space constraints, results for this section are provided in the appendix. Additionally, LLaVA-Video-7B-Qwen2 [41] serves as the base model for the video tasks evaluation.

Algorithm 1 AirCache for KV Cache Compression.

- 1: **Input:** Total cache budget r , key, query, value of text prompt tokens $K_t, Q_t, V_t \in \mathbb{R}^{N_t \times D}$, key, query, value of visual prompt tokens $K_v, Q_v, V_v \in \mathbb{R}^{N_v \times D}$, number of layers L .
- 2: **for** layer $l \leftarrow 0$ to $N - 1$ **do**
- 3: $A_{tt} \leftarrow \text{Softmax} \left(\frac{\mathbf{Q}_t \mathbf{K}_t^T}{\sqrt{D}} \right)$
- 4: Obtain the index of key text tokens k using Eq. (5).
- 5: $Q_{t_k}, K_{v_{t_k}} \leftarrow Q_t[k], \text{Concat}(K_v, K_t[k])$
- 6: $\mathbf{A}_{v_{t_k}} \leftarrow \text{Softmax} \left(\frac{\mathbf{Q}_{t_k} \mathbf{K}_{v_{t_k}}^T}{\sqrt{D}} \right)$
- 7: $\mathbf{I}_v \leftarrow \frac{1}{N_{t_k}} \sum_{j=0}^{N_{t_k}-1} \mathbf{A}_{v_{t_k}}[j, : N_v] \triangleright$ Importance Scores
- 8: $s_t \leftarrow \sum_{i=0}^{N_v-1} \mathbf{I}_v[i] \triangleright$ Strength
- 9: $s_k \leftarrow \frac{N_v}{(N_v-1)(N_v-2)} \sum_{i=1}^{N_v} \left(\frac{\mathbf{I}_v[i] - \mu_{I_v}}{\sigma_{I_v}} \right)^3 \triangleright$ Skewness
- 10: $C^l \leftarrow (\{(K_v, K_t), (V_v, V_t), (\mathbf{I}_v, s_t, s_k)\})$
- 11: **end for**
- 12: **for** layer $l \leftarrow 0$ to $N - 1$ **do**
- 13: $K'_v, K'_t, V_v, V_t, \mathbf{I}_v, s_t, s_k \leftarrow C^l$
- 14: $s_t, s_k \leftarrow s_t, s_k \triangleright$ Normalized Across Different Layers
- 15: $\hat{r} = \frac{1}{2}(s'_t + s'_k)r \triangleright$ Obtain Layer-wise Budget
- 16: $K' \leftarrow \text{Concat}(K_v(\mathbf{I}[N_v \times \hat{r}]), K_t) \triangleright$ Compress Key
- 17: $V' \leftarrow \text{Concat}(V_v(\mathbf{I}[V_v \times \hat{r}]), V_t) \triangleright$ Compress Value
- 18: $C^l \leftarrow (\{K', V'\})$
- 19: **end for**
- 20: **Return:** $C \triangleright$ Compressed KV Cache

Datasets. The performance is primarily evaluated across four types of datasets: 1) Visual Question Answering (VQA) datasets, which include ChatQA [26], InfoVQA [27], DocVQA [28], and TextVQA [31]. Since KV cache compression mainly affects the decoding stage, VQA datasets that require longer output tokens are preferred for evaluating the method's efficacy. The similarity between generated and reference answers is typically measured by the Average Normalized Levenshtein Similarity (ANLS) metric [2]. 2) Multiple-choice question datasets, such as MMBench-EN [24], involves selecting the correct option from several choices, with accuracy serving as the ultimate metric. 3) Judgment datasets like MME [10], also utilize judgment accuracy as the evaluation criterion. 4) Video VQA datasets, such as MMBench-Video [9], measure the alignment between generated and reference answers using metrics provided by a large model evaluator.

Baselines. The comparison focuses on two main categories of methods. The first category encompasses classic and prominent KV cache compression techniques in LLMs, such as H2O [42], and SnapKV [19]. The second category comprises methods specifically optimized for LVLMS, namely PrefixKV [34] and Elastic Cache [25]. H2O [42] employs a full-range approach, while SnapKV [19] adopts a window-based strategy using textual tokens as the observation window for evaluating token importance. PrefixKV [34] utilizes the Lorenz curve to characterize the impor-

Table 1. The comparison of the KV cache compression methods on multimodal VQA benchmarks. The best result is highlighted in bold.

Models	Methods	ChatQA [26]				InfoVQA [27]				DocVQA [28]				TextVQA [31]			
		50%	10%	5%	1%	50%	10%	5%	1%	50%	10%	5%	1%	50%	10%	5%	1%
LLaVA-OV-7B [16]	Full	80.3	80.3	80.3	80.3	66.1	66.1	66.1	66.1	87.0	87.0	87.0	87.0	76.0	76.0	76.0	76.0
	H2O [42]	79.8	77.4	74.0	71.0	65.4	59.2	54.2	52.0	86.4	74.2	60.8	55.3	75.5	70.1	67.9	60.4
	Elastic [25]	80.0	77.9	76.5	71.0	65.3	60.5	57.2	52.0	86.6	74.2	60.8	55.3	75.5	72.0	68.7	60.2
	PrefixKV [34]	80.1	78.2	76.8	70.9	65.4	61.1	58.0	51.9	86.6	80.5	74.7	55.4	75.8	72.7	70.0	60.2
	SnapKV [19]	80.1	79.3	78.3	72.9	66.1	64.2	63.0	57.8	86.9	84.4	80.6	64.1	75.8	73.4	69.5	58.2
	Ours	80.4	79.9	79.4	76.4	66.0	65.7	64.6	62.5	86.8	85.5	83.3	73.2	76.0	75.3	73.4	67.1
InternVL2-8B [6]	Full	82.2	82.2	82.2	82.2	73.0	73.0	73.0	73.0	91.0	91.0	91.0	91.0	77.7	77.7	77.7	77.7
	H2O [42]	81.9	81.6	79.3	76.8	72.3	70.6	65.1	61.2	90.5	79.4	70.3	48.9	77.1	70.0	66.0	59.6
	Elastic [25]	81.8	81.0	79.4	76.8	72.3	67.0	65.1	60.8	90.5	79.4	70.3	53.5	77.0	70.2	66.0	59.0
	PrefixKV [34]	81.9	80.7	79.3	76.7	72.3	67.4	65.3	61.0	90.6	79.5	70.2	55.1	77.1	70.3	66.2	58.6
	SnapKV [19]	82.2	80.4	79.5	75.0	72.8	72.3	68.4	64.9	90.9	90.1	83.5	71.1	77.4	73.9	71.0	61.7
	Ours	82.2	81.7	81.0	78.4	73.2	72.6	71.8	68.2	91.0	90.0	88.5	78.5	77.5	77.0	75.4	68.6
Qwen2-VL-7B [35]	Full	83.2	83.2	83.2	83.2	76.2	76.2	76.2	76.2	93.8	93.8	93.8	93.8	84.4	84.4	84.4	84.4
	H2O [42]	83.2	78.9	75.0	66.4	75.4	65.7	61.0	55.1	90.5	82.3	70.7	48.9	84.2	75.6	68.2	54.5
	Elastic [25]	83.2	79.7	75.4	68.0	75.5	67.1	62.3	55.5	93.5	83.0	71.9	49.2	84.2	76.7	68.8	54.4
	PrefixKV [34]	83.2	79.4	75.9	68.0	75.7	67.9	62.7	55.7	93.7	83.7	72.8	49.8	84.2	76.9	69.4	54.3
	SnapKV [19]	83.3	81.6	80.0	69.1	76.3	74.9	73.0	66.2	93.7	87.2	87.1	71.2	84.3	83.1	74.6	60.2
	Ours	83.4	82.3	80.8	75.9	76.2	75.2	74.5	70.7	93.8	92.9	90.9	78.3	84.4	83.4	80.9	68.6

Table 2. Performance comparison on other types datasets. Best-performing results are highlighted in bold.

Methods	MMBench-EN [24]			MME [10]			MMBench-Video [9]		
	50%	10%	1%	50%	10%	1%	50%	10%	1%
Full	82.3	82.3	82.3	1585	1585	1585	1.81	1.81	1.81
H2O [42]	0.82	0.82	0.82	1585	1585	1585	1.66	1.63	1.47
Elastic [25]	0.82	0.82	0.82	1585	1585	1585	1.68	1.62	1.51
PrefixKV [34]	0.82	0.82	0.82	1585	1585	1585	1.72	1.68	1.50
SnapKV [19]	0.82	0.82	0.82	1585	1585	1585	1.75	1.70	1.55
Ours	0.82	0.82	0.82	1585	1585	1585	1.80	1.78	1.67

tance distribution across different layers, guided by the retained important KV matrix. Elastic Cache [25] employs a fixed-point elimination strategy for the token importance with cache merging. All KV cache compression methods are applied exclusively to the visual tokens, as experiments have demonstrated that it offers greater stability and better performance compared to compressing the entire tokens.

Implementation details. All evaluation results are derived from VLMEvalKit [8], which is an open-source evaluation toolkit for LVLMs. The relevance threshold in Equation 5 is set to 0.9. The referee model for MMBench-Video [9] is Qwen2.5-72B [37]. The proposed method performs a one-time KV cache compression only after the prefill stage is completed. Specifically, once the prefill ends for each layer, the complete KV cache is fully saved. After computing all layers, we then calculate the budget for each layer and compress the KV cache accordingly. All experiments are conducted on $8 \times A100$ -80G GPUs.

4.2. Main Results

We report the results in four sections: VQA datasets, other types of datasets including multiple-choice questions, judgment, and video VQA, inference performance, and the comparison with token pruning methods.

VQA datasets. Table 1 presents the comparative results on the ChatQA [26], InfoVQA [27], DocVQA [28], and TextVQA [31] datasets. These datasets encompass tasks such as conventional image question answering, image in-

formation extraction, document image understanding, and text recognition within images. It can be observed that the proposed method achieves the best results under most reduction ratios. When 50% of visual KV cache is pruned, our method performs almost equivalently to the full KV cache, indicating near lossless performance. Even when reduced to only 10% of visual KV cache, the performance gap with the full cache is limited to approximately 1%. As the token compression ratio increases, the advantage of the proposed method becomes more pronounced. For instance, when only 1% of visual tokens are retained, we outperform SnapKV [19] by an average of 6.5% across four VQA datasets with the LLaVA-OV-7B [16]. Moreover, the performance differences across different types of VQA datasets indicate that image documents or text understanding, which require higher local visual perception, have stricter demands on KV cache compression than general image understanding. The proposed method demonstrates a more significant advantage over other methods with these datasets, providing strong evidence of its superiority. Furthermore, the above conclusions also apply to different LVLMs, we achieve a significant advantage with InternVL2-8B [6] and Qwen2-VL-7B [35].

Other types datasets. Table 2 presents the results of various methods on other types of multimodal datasets. For datasets involving multiple-choice questions and judgment types, such as MMBench-EN [24] and MME [10], the final result heavily depends on the accuracy of decoding the first token. Since the KV cache only influences the tokens decoded after the initial token, all methods achieve the same performance as the full KV cache for these types of datasets. The results from the MMBench-Video [9] dataset indicate that our proposed method achieves superior performance, further demonstrating its generality and superiority.

Inference efficiency. We maintain a constant output token count of 512 and compare the full cache with the proposed

Table 3. Quantitative results on inference latency and throughput. The number of tokens output is consistently set to 512.

Batch Size	Prompt Length	Prefill Latency (s)		Decoding Latency (s)			Throughput (token/s)		
		Full	Ours	Full	50%	10%	Full	50%	10%
8	2k	1.6	1.8 _{+12.5%}	11.6	9.4 _{-19.0%}	8.2 _{-29.3%}	353	436 _{+23.5%}	500 _{+41.6%}
	8k	4.9	5.3 _{+8.2%}	15.0	11.9 _{-20.7%}	9.4 _{-37.3%}	273	344 _{+26.0%}	436 _{+59.7%}
	16k	11.3	11.9 _{+5.3%}	21.7	15.0 _{-30.9%}	10.4 _{-47.5%}	189	273 _{+44.4%}	394 _{+108.5%}
	32k	25.2	26.4 _{+4.8%}	35.9	22.0 _{-38.7%}	12.3 _{-65.7%}	114	186 _{+63.2%}	333 _{+192.1%}
16	2k	2.3	2.5 _{+8.7%}	14.4	11.2 _{-22.2%}	9.7 _{-32.6%}	569	731 _{+28.5%}	845 _{+48.5%}
	8k	9.4	9.8 _{+4.3%}	22.6	15.6 _{-31.0%}	11.8 _{-47.8%}	362	525 _{+45.0%}	694 _{+91.7%}
	16k	20.3	21.4 _{+5.4%}	36.9	23.0 _{-37.7%}	12.8 _{-65.3%}	222	356 _{+60.4%}	640 _{+188.3%}

Table 4. Performance comparison with token pruning methods under various compression ratio. The prefill and decoding latency are measured under the settings of batch size of 16, total input token of 8k, and the compression ratio of 10%.

Methods	ChatQA [26]			MMBench-EN [24]			MME [10]			MMBench-Video [9]			Latency (s)	
	50%	10%	1%	50%	10%	1%	50%	10%	1%	50%	10%	1%	Prefill	Decoding
Full	80.3	80.3	80.3	82.3	82.3	82.3	1585	1585	1585	1.81	1.81	1.81	9.4	22.6
FastV [5]	73.6	47.7	16.9	81.4	74.8	33.6	1574	1378	786	1.69	1.36	1.02	5.4	11.7
FasterVLM [39]	75.5	50.1	18.8	81.4	77.0	26.5	1559	1441	731	1.71	1.42	1.17	5.2	11.3
IVTP [13]	74.9	55.8	22.5	81.7	77.5	36.2	1563	1469	849	1.70	1.47	1.24	5.8	12.6
Ours	80.4	79.9	76.4	82.3	82.3	82.3	1585	1585	1585	1.79	1.72	1.60	9.8	11.8

AirCache across different batch sizes and input token counts in terms of prefill latency, decoding latency, and throughput. To facilitate testing, the text instruction token count is set to 64. As shown in Table 3, the proposed method significantly improves inference decoding speed under various inputs compared to the full cache. When compressing 50% of the visual KV cache, the proposed method minimally impacts model performance and prefill latency, while reducing decoding latency by 19% to 39% and increasing model throughput by 24% to 63%. When compressing 10% of the visual KV cache, the proposed method slightly decreases model performance by less than 1%, while further reducing decoding latency by 29% to 66% and increasing model throughput by 42% to 192%. These results demonstrate that our proposed method can maintain strong model performance while ensuring significant decoding acceleration, indicating the effectiveness of visual token importance assessment and hierarchical budget allocation.

Compared with token pruning methods. In addition to KV cache compression, another important method for accelerating LVLMs inference is to prune visual tokens before the prefill stage. Table 4 compares the proposed method with these type of acceleration methods, including FastV [5], FasterVLM [39], and IVTP [13]. A notable characteristic is that the token pruning methods exhibit acceleration effects during the prefill stage as they preemptively prune visual tokens. In contrast, all KV cache compression methods, including the proposed approach, demonstrate prefill stage inference speeds comparable to those of the full cache. However, comparisons across various evaluation sets clearly show that the proposed method achieves an absolute advantage, particularly on VQA datasets, even when retaining only 1% of the visual tokens. In contrast, token pruning methods experience a significant performance

Table 5. Ablation results on the components of the AirCache.

Setting	ChatQA [26]	InfoVQA [27]	DocVQA [28]	TextVQA [31]
Continuous window (16)	70.4	56.6	61.3	55.9
Continuous window (32)	72.9	57.8	64.1	58.2
All text tokens	72.2	58.4	65.7	57.0
Visual window (32)	68.8	55.1	59.2	53.7
Average allocation	72.2	57.5	69.9	62.4
Pyramid allocation	69.6	54.9	55.8	52.6
Only w/ s_t	74.2	59.8	71.1	64.9
Only w/ s_k	74.7	61.4	71.9	63.6
Ours	76.4	62.5	73.2	67.1

decline under these conditions. The prefill stage facilitates sufficient interaction between visual and textual information, as the attention mechanism aggregates visual information to key text tokens. The proposed method effectively preserves these key text tokens during KV cache token pruning. Consequently, even when only 1% of visual tokens are retained, the model maintains good inference performance. In contrast, most token pruning methods do not involve cross-modal interaction, resulting in direct visual information loss that cannot be transferred to the subsequent decoding stage in any form. This leads to insufficient and inaccurate visual guidance for generating, significantly degrading performance.

4.3. Ablation Studies

In this section, we evaluate our method under various settings to verify the proposed modules from several perspectives. The base model is LLaVA-OV-7B [16].

Elite Observation Window. We introduce three common types of observation windows to assess the effectiveness of the proposed elite observation window. The first type employs a sequence of continuous text tokens as the observation window for evaluating the significance of visual tokens. The second type utilizes all text tokens, whereas the third employs the continuous visual token from the last visual

Table 6. Comparison results with different compression audiences, where * indicates uniform compression without distinguishing between visual tokens and text tokens.

Models	50%	10%	5%	1%
H2O* [42]	74.6/407s	46.5/559s	32.0/733s	27.8/1048s
H2O [42]	79.8/323s	77.4/251s	74.0/239s	71.0/215s
Ours*	75.8/374s	53.4/530s	39.7/688s	30.1/963s
Ours	80.4/306s	79.9/235s	79.4/221s	76.4/207s

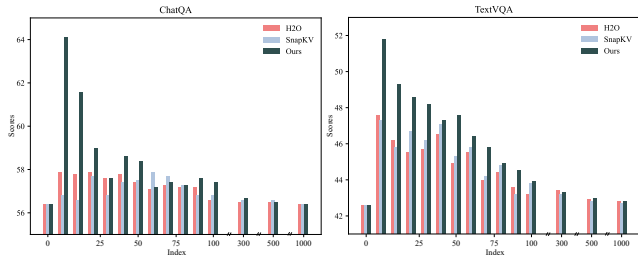


Figure 4. Comparison on ChatQA [26] and TextVQA [31] by retaining only one visual token, which is selected based on the sorting of visual token importance scores using different methods.

segment as a reference. As shown in the first column of Table 5, the proposed elite observation window yields the best performance compared to the other three aforementioned approaches. The method using only visual tokens performs significantly worse due to the absence of textual instruction guidance, highlighting the importance of activating cross-modal associations in KV cache compression. Additionally, compared to methods that directly use all or portions of continuous text token sequences without considering variations in visual information perception among text tokens, the proposed method selects key text tokens to enhance the consistency of the observation window. This selection ensures greater stability and accuracy in evaluating the significance of visual tokens.

Layer-wise Token Budget Allocation Module. The second column of Table 5 presents the comparison using various budget allocation strategies across layers. Aside from the basic equal allocation, we also incorporate the strategy from PyramidKV [3], which allocates budgets in a descending ratio from shallow to deep layers. The results demonstrate that the proposed budget allocation, based on the strength and skewness of importance scores, delivers superior outcomes. Although PyramidKV [3] perform well in LLMs, its suboptimal results in LVLMs highlight the distinct characteristics of multimodal models and their heightened sensitivity to hierarchical budget allocation. An inappropriate allocation strategy is more likely to adversely affect the final results. We also investigate the separate use of strength and skewness for layer-wise budget quantification, revealing that their combination produces the best results.

Visual Token Importance Evaluation. Figure 4 visually demonstrates the comparative effectiveness of various methods in evaluating the importance of visual tokens. To

directly assess the efficacy of these methods in ranking visual token importance, we removed all non-essential modules and selected only one visual token at a time from the ranked sorting results, ordered from high to low importance. Theoretically, as the importance ranking of tokens declines, the model’s performance is also expected to decrease. The graph clearly shows that the proposed method aligns more closely with the theoretical trend, where the visual token ranked as TOP1 achieves the highest metric, followed by a continuous and smooth decline. Notably, after around 10%, the performance of the visual tokens approximates that of the case with all visual tokens removed. This indicates significant redundancy among visual tokens in the KV cache of LVLMs. The results in Table 1, which demonstrate that retaining over 10% of the visual tokens limits performance decline to no more than 1%, further corroborate this observation. In contrast, the performance derived from other methods exhibit significant noise, indicating instability and deficiencies in their importance assessment.

Unified Compression or Vision-only Compression.

Given that the KV cache in LVLMs encompasses both visual and textual modalities, in contrast to the single textual modality of LLMs, we compared Unified Compression, which does not differentiate between modal differences during KV cache compression, with Vision-only Compression, which compresses only visual tokens. As shown in Table 6, the results indicate that unified compression performs significantly worse than vision-only compression in terms of both inference performance and speed. Moreover, as the compression ratio increases, the inference time for unified compression actually continues to rise. The primary reason is that unified compression disrupts the effective representation of textual instructions, leading the model to produce incorrect outputs and exhibit a tendency for repetition, which in turn increases inference time.

5. Conclusion

In this work, we propose Activating Inter-modal Relevancy (AirCache), a novel KV cache compression for efficient LVLMs inference. By leveraging inter-modal relevancy between textual and visual tokens to guide KV cache eviction, we identified significant differences in the attention distribution of visual tokens corresponding to different textual tokens. Based on this finding, we propose an elite observation window with refined key text tokens, which provides a more stable and effective assessment of visual token importance. Furthermore, by exploring the diversity of importance score distribution across different layers, we propose to achieve comprehensive quantification in terms of distribution strength and skewness for adaptive layer-wise compression budget allocation. Extensive experiments across multiple benchmarks and LVLMs demonstrate the effectiveness of AirCache in KV cache compression.

References

- [1] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. [1](#), [2](#)
- [2] Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gomez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. Scene text visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4291–4301, 2019. [5](#)
- [3] Zefan Cai, Yichi Zhang, Bofei Gao, Yuliang Liu, Tianyu Liu, Keming Lu, Wayne Xiong, Yue Dong, Baobao Chang, Junjie Hu, et al. Pyramidkv: Dynamic kv cache compression based on pyramidal information funneling. *arXiv preprint arXiv:2406.02069*, 2024. [2](#), [8](#)
- [4] Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. Honeybee: Locality-enhanced projector for multimodal llm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13817–13827, 2024. [2](#)
- [5] 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. In *European Conference on Computer Vision*, pages 19–35. Springer, 2024. [1](#), [2](#), [7](#)
- [6] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 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*, pages 24185–24198, 2024. [1](#), [5](#), [6](#), [11](#), [12](#)
- [7] Xiangxiang Chu, Limeng Qiao, Xinyang Lin, Shuang Xu, Yang Yang, Yiming Hu, Fei Wei, Xinyu Zhang, Bo Zhang, Xiaolin Wei, et al. Mobilevlm: A fast, reproducible and strong vision language assistant for mobile devices. *arXiv preprint arXiv:2312.16886*, 2023. [2](#)
- [8] Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Jiaqi Wang, et al. Vlmevalkit: An open-source toolkit for evaluating large multi-modality models. In *Proceedings of the 32nd ACM international conference on multimedia*, pages 11198–11201, 2024. [6](#)
- [9] Xinyu Fang, Kangrui Mao, Haodong Duan, Xiangyu Zhao, Yining Li, Dahua Lin, and Kai Chen. Mmbench-video: A long-form multi-shot benchmark for holistic video understanding. *arXiv preprint arXiv:2406.14515*, 2024. [5](#), [6](#), [7](#), [12](#), [13](#)
- [10] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, and Rongrong Ji. Mme: A comprehensive evaluation benchmark for multimodal large language models. *ArXiv*, abs/2306.13394, 2023. [5](#), [6](#), [7](#)
- [11] Yefei He, Feng Chen, Jing Liu, Wenqi Shao, Hong Zhou, Kaipeng Zhang, and Bohan Zhuang. Zipvl: Efficient large vision-language models with dynamic token sparsification and kv cache compression. *arXiv preprint arXiv:2410.08584*, 2024. [1](#)
- [12] Wenbo Hu, Zi-Yi Dou, Liunian Harold Li, Amita Kamath, Nanyun Peng, and Kai-Wei Chang. Matryoshka query transformer for large vision-language models. *arXiv preprint arXiv:2405.19315*, 2024. [2](#)
- [13] Kai Huang, Hao Zou, Ye Xi, BoChen Wang, Zhen Xie, and Liang Yu. Ivtp: Instruction-guided visual token pruning for large vision-language models. In *European Conference on Computer Vision*, pages 214–230. Springer, 2024. [1](#), [2](#), [7](#)
- [14] Chaoya Jiang, Haiyang Xu, Wei Ye, Qinghao Ye, Chenliang Li, Ming Yan, Bin Bi, Shikun Zhang, Fei Huang, and Songfang Huang. Bus: Efficient and effective vision-language pre-training with bottom-up patch summarization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2900–2910, 2023. [1](#)
- [15] Jang-Hyun Kim, Junyoung Yeom, Sangdoon Yun, and Hyun Oh Song. Compressed context memory for online language model interaction. *arXiv preprint arXiv:2312.03414*, 2023. [2](#)
- [16] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024. [1](#), [4](#), [5](#), [6](#), [7](#), [11](#), [12](#), [14](#), [15](#)
- [17] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. [2](#)
- [18] Wentong Li, Yuqian Yuan, Jian Liu, Dongqi Tang, Song Wang, Jie Qin, Jianke Zhu, and Lei Zhang. Tokenpacker: Efficient visual projector for multimodal llm. *arXiv preprint arXiv:2407.02392*, 2024. [1](#)
- [19] Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. Snapkv: Llm knows what you are looking for before generation. *arXiv preprint arXiv:2404.14469*, 2024. [1](#), [2](#), [5](#), [6](#), [11](#), [12](#), [13](#)
- [20] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models. In *European Conference on Computer Vision*, pages 323–340. Springer, 2024. [1](#), [2](#)
- [21] Akide Liu, Jing Liu, Zizheng Pan, Yefei He, Reza Haf-fari, and Bohan Zhuang. Minicache: Kv cache compression in depth dimension for large language models. *Advances in Neural Information Processing Systems*, 37:139997–140031, 2025. [1](#), [2](#)
- [22] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306, 2024. [1](#), [11](#)
- [23] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023. [1](#)
- [24] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He,

- Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In *European conference on computer vision*, pages 216–233. Springer, 2024. [5](#), [6](#), [7](#)
- [25] Zuyan Liu, Benlin Liu, Jiahui Wang, Yuhao Dong, Guangyi Chen, Yongming Rao, Ranjay Krishna, and Jiwen Lu. Efficient inference of vision instruction-following models with elastic cache. In *European Conference on Computer Vision*, pages 54–69. Springer, 2024. [1](#), [2](#), [5](#), [6](#), [11](#), [12](#), [13](#)
- [26] Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022. [4](#), [5](#), [6](#), [7](#), [8](#), [11](#), [12](#), [13](#)
- [27] Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1697–1706, 2022. [5](#), [6](#), [7](#), [11](#), [12](#), [13](#), [14](#)
- [28] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209, 2021. [5](#), [6](#), [7](#), [11](#), [12](#), [13](#), [14](#)
- [29] Piotr Nawrot, Adrian Łańcucki, Marcin Chochowski, David Tarjan, and Edoardo M Ponti. Dynamic memory compression: Retrofitting llms for accelerated inference. *arXiv preprint arXiv:2403.09636*, 2024. [2](#)
- [30] Reiner Pope, Sholto Douglas, Aakanksha Chowdhery, Jacob Devlin, James Bradbury, Jonathan Heek, Kefan Xiao, Shivani Agrawal, and Jeff Dean. Efficiently scaling transformer inference. *Proceedings of Machine Learning and Systems*, 5:606–624, 2023. [1](#)
- [31] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019. [5](#), [6](#), [7](#), [8](#), [11](#), [12](#), [13](#)
- [32] Peter Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Adithya Jairam Vedagiri IYER, Sai Charitha Akula, Shusheng Yang, Jihan Yang, Manoj Middepogu, Ziteng Wang, et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *Advances in Neural Information Processing Systems*, 37:87310–87356, 2025. [1](#)
- [33] Dezhan Tu, Danylo Vashchilenko, Yuzhe Lu, and Panpan Xu. V1-cache: Sparsity and modality-aware kv cache compression for vision-language model inference acceleration. *arXiv preprint arXiv:2410.23317*, 2024. [2](#)
- [34] Ao Wang, Hui Chen, Jianchao Tan, Kefeng Zhang, Xunliang Cai, Zijia Lin, Jungong Han, and Guiguang Ding. Prefixkv: Adaptive prefix kv cache is what vision instruction-following models need for efficient generation. *arXiv preprint arXiv:2412.03409*, 2024. [1](#), [2](#), [5](#), [6](#), [12](#), [13](#)
- [35] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. [1](#), [5](#), [6](#), [11](#), [12](#)
- [36] Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, et al. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding. *arXiv preprint arXiv:2412.10302*, 2024. [1](#)
- [37] An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024. [6](#)
- [38] Yifei Yang, Zouying Cao, Qiguang Chen, Libo Qin, Dongjie Yang, Hai Zhao, and Zhi Chen. Kvsharer: Efficient inference via layer-wise dissimilar kv cache sharing. *arXiv preprint arXiv:2410.18517*, 2024. [2](#)
- [39] Qizhe Zhang, Aosong Cheng, Ming Lu, Zhiyong Zhuo, Minqi Wang, Jiajun Cao, Shaobo Guo, Qi She, and Shanghang Zhang. [cls] attention is all you need for training-free visual token pruning: Make vlm inference faster. *arXiv preprint arXiv:2412.01818*, 2024. [1](#), [2](#), [7](#)
- [40] Yuxin Zhang, Yuxuan Du, Gen Luo, Yunshan Zhong, Zhenyu Zhang, Shiwei Liu, and Rongrong Ji. Cam: Cache merging for memory-efficient llms inference. In *Forty-first International Conference on Machine Learning*, 2024. [2](#)
- [41] Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*, 2024. [5](#)
- [42] Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark Barrett, et al. H2o: Heavy-hitter oracle for efficient generative inference of large language models. *Advances in Neural Information Processing Systems*, 36, 2024. [1](#), [2](#), [4](#), [5](#), [6](#), [8](#), [12](#), [13](#)

Roadmap of Appendix

In this supplementary material, we first state the limitations of the proposed method and potential future work in Section A. Next, we provide more details on the method’s application in Section B. After that, additional main comparative experiments with more models are discussed in Section C. Furthermore, we present additional ablation experiments in Section D. Finally, the visualization of chat generation is shown in Section E.

A. Limitations and Future Works

In performing dynamic allocation of the layer-wise compression budget, the proposed method requires obtaining the strength and skewness of the importance distribution of visual tokens for all layers before determining the allocable budget for each layer. This necessitates storing the complete KV cache after the prefill stage and executing the reduction only once the final compression budget is determined. Consequently, the proposed method is at a disadvantage in terms of peak memory consumption, a challenge also faced by most hierarchical budget allocation methods. Addressing how to maintain a peak memory advantage while supporting dynamic allocation of budgets across layers will be a key focus of our future work. In parallel, we will continue to explore the information flow mechanisms of different modalities in the inference process of LVLMS to further optimize the proposed method.

B. Implementation Details

For most methods, we adhere to their initial setup and perform the reduction of the visual KV cache based on the obtained importance ranking of visual tokens and the specified compression ratio. Our findings indicate that merging the dropped KV cache into the KV cache that needs to be retained works effectively on the LLaVA-v1.5 [22]. However, this approach tends to cause repetition issues in the LLaVA-OV series [16], InternVL2 series [6], and Qwen2-VL series [35], which results in a decline in model performance. Consequently, for Elastic Cache [25], we omitted the merge operation in the main experiments to achieve optimal performance results. In practical applications, and as observed in most existing multimodal evaluation datasets, visual tokens constitute the majority, while text tokens remain concise and short. The redundancy in the KV cache primarily resides in the visual part. As demonstrated in Table 7, a comparison of the actual number of visual tokens and text tokens in these multimodal VQA datasets shows that the visual component accounts for more than 97%. Thus, compressing only the visual KV cache eliminates redundant cache without affecting the complete expression of text instructions. Unless otherwise specified, all methods

Table 7. The number of visual tokens and text tokens across different models and evaluation sets.

Models	ChatQA [26]		InfoVQA [27]		DocVQA [28]		TextVQA [31]	
	N_v	N_t	N_v	N_t	N_v	N_t	N_v	N_t
LLaVA-OV-7B [16]	4763	47	6382	45	7224	42	5183	39
InternVL2-8B [6]	1828	32	3740	31	3230	28	1668	25
Qwen2-VL-7B [35]	1302	36	4450	34	4669	31	1325	28

and experiments perform KV cache compression solely on the visual part.

C. Additional Main Results

Comparison with various model parameter sizes. Table 8 displays the comparison results of different parameter-sized InternVL2 [6] series models across various VQA datasets as the compression ratio varies. Similar to the conclusions drawn from experiments with different model architectures, the proposed method achieves superior results on models with different parameter sizes compared to existing methods. For instance, when retaining only 1% of the visual KV cache, the proposed method outperforms the SnapKV [19] method by an average of approximately 4.3% to 6.0% across four VQA evaluation datasets as the model parameter size varies. By synthesizing experiments on different architectures and base models with varying parameter sizes, we observe that the proposed method not only achieves better compression results but also demonstrates good general applicability. Furthermore, a comparison of models with different parameter sizes reveals that as the model parameter size decreases, the impact of KV cache compression on model performance becomes more significant. This trend indicates that smaller parameter-sized models are less effective at integrating information within tokens, thereby placing a greater emphasis on the KV cache compression method’s ability to select important visual tokens. The proposed method demonstrates a superior capability in assessing the importance of visual tokens, thereby more effectively reducing model performance loss.

Inference efficiency on more LVLMS. Inference efficiency on more LVLMS. Table 9 further presents the comparison of inference latency between the proposed method and the full cache on InternVL2-8B [6] and Qwen2-VL-7B [35]. From the table, it can be observed that when the input demand is relatively low, the model’s need for the KV cache is reduced, thus limiting the gains from the KV cache compression method. Nevertheless, there is at least a 21% reduction in decoding latency and a 27% increase in throughput. As the input demand continues to rise, the benefits from KV cache compression become more significant. For example, in the case of the Qwen2-VL-7B [35] with a batch size of 16 and a prompt length of 16k, the proposed method can reduce decoding latency by 42% and increase throughput by 73% with almost no impact on model performance,

Table 8. The comparison of the KV cache compression methods on multimodal VQA benchmarks. The best result is highlighted in bold.

Models	Methods	ChatQA [26]				InfoVQA [27]				DocVQA [28]				TextVQA [31]			
		50%	10%	5%	1%	50%	10%	5%	1%	50%	10%	5%	1%	50%	10%	5%	1%
InternVL2-1B [6]	Full	67.9	67.9	67.9	67.9	50.1	50.1	50.1	50.1	80.0	80.0	80.0	80.0	70.8	70.8	70.8	70.8
	H2O [42]	67.7	62.0	57.7	53.5	50.0	42.7	38.8	33.7	79.8	73.8	68.2	57.0	70.1	58.4	52.3	47.6
	Elastic [25]	67.5	61.8	57.6	54.1	50.1	42.6	39.7	32.2	79.8	74.0	68.8	57.5	70.3	59.8	54.7	47.2
	PrefixKV [34]	67.9	62.1	58.0	53.3	49.9	43.8	40.6	34.2	79.7	74.3	70.4	59.6	70.1	60.7	55.3	48.6
	SnapKV [19]	67.8	63.7	59.8	56.7	50.1	47.3	44.4	39.7	79.8	76.5	72.1	61.2	70.4	64.5	60.7	52.2
	Ours	67.8	65.7	63.5	60.8	50.1	49.6	47.5	45.8	80.0	77.7	74.3	68.5	70.6	68.9	66.5	59.3
InternVL2-4B [6]	Full	81.1	81.1	81.1	81.1	65.9	65.9	65.9	65.9	88.1	88.1	88.1	88.1	74.7	74.7	74.7	74.7
	H2O [42]	81.1	79.2	77.6	72.1	65.9	61.1	57.4	51.8	79.9	80.1	75.4	69.2	74.2	66.4	58.8	51.3
	Elastic [25]	81.1	79.4	77.9	73.6	65.8	61.8	59.2	53.3	79.7	80.7	75.9	69.6	74.0	67.0	60.4	52.6
	PrefixKV [34]	81.0	79.5	77.8	73.2	65.9	62.6	59.1	53.4	88.0	81.4	76.4	70.3	74.4	67.5	61.4	53.7
	SnapKV [19]	81.1	79.3	78.5	74.6	65.9	64.3	61.8	56.7	88.0	84.3	79.7	73.2	74.4	70.3	65.8	60.3
	Ours	81.1	80.4	79.6	77.7	66.0	65.5	64.2	62.1	88.1	86.8	84.3	81.5	74.5	73.6	70.7	67.4
InternVL2-26B [6]	Full	85.4	85.4	85.4	85.4	75.4	75.4	75.4	75.4	92.1	92.1	92.1	92.1	82.5	82.5	82.5	82.5
	H2O [42]	84.9	82.4	80.4	78.6	75.0	71.8	65.1	62.5	91.9	84.4	81.6	75.1	82.3	75.2	70.3	65.2
	Elastic [25]	84.6	82.8	81.6	78.3	74.8	73.6	65.5	62.7	91.8	83.8	81.2	74.3	82.4	75.6	70.7	65.7
	PrefixKV [34]	84.8	82.2	81.5	78.8	75.2	73.8	66.4	63.2	92.0	84.2	81.5	74.7	82.4	75.5	70.6	65.4
	SnapKV [19]	85.3	83.5	83.0	80.1	75.4	74.1	69.5	65.6	91.9	86.6	86.3	82.5	82.5	78.6	74.3	71.1
	Ours	85.5	84.7	84.2	82.3	75.4	74.8	72.8	70.7	92.1	91.4	89.0	86.8	82.4	81.7	78.2	76.6

Table 9. Quantitative results on inference latency and throughput. The number of tokens output is consistently set to 512.

Models	Batch Size	Prompt Length	Prefill Latency (s)		Decoding Latency (s)			Throughput (token/s)			
			Full	Ours	Full	50%	10%	Full	50%	10%	
			InternVL2-8B [6]	8	2k	1.2	1.4 ^{+16.7%}	9.3	7.0 ^{+24.7%}	5.9 ^{+36.6%}	440
Qwen2-VL-7B [35]	8	8k	4.6	5.0 ^{+8.7%}	13.3	10.5 ^{+21.1%}	8.8 ^{+33.8%}	308	390 ^{+26.6%}	465 ^{+51.0%}	
		16k	9.8	10.5 ^{+7.1%}	24.8	15.7 ^{+36.7%}	11.9 ^{+52.0%}	165	261 ^{+58.2%}	344 ^{+52.0%}	
		32k	23.8	24.8 ^{+4.2%}	46.2	28.1 ^{+39.2%}	18.2 ^{+60.6%}	89	146 ^{+64.0%}	225 ^{+60.4%}	
		16	2k	2.5	2.7 ^{+8.0%}	12.2	7.4 ^{+39.3%}	6.0 ^{+50.8%}	671	1107 ^{+65.0%}	1365 ^{+50.8%}
	Qwen2-VL-7B [35]	8	8k	9.9	10.6 ^{+7.1%}	21.6	12.9 ^{+40.3%}	8.9 ^{+58.8%}	379	635 ^{+67.5%}	920 ^{+58.8%}
			16k	21.6	22.4 ^{+3.7%}	28.5	16.2 ^{+43.2%}	10.0 ^{+64.9%}	287	506 ^{+76.3%}	819 ^{+65.0%}
16			2k	1.1	1.3 ^{+18.2%}	8.4	6.6 ^{+21.4%}	5.2 ^{+38.1%}	488	621 ^{+27.3%}	788 ^{+38.1%}
			8k	4.3	4.7 ^{+27.3%}	12.6	9.7 ^{+23.0%}	8.1 ^{+35.7%}	325	422 ^{+29.8%}	506 ^{+55.7%}
		16k	9.4	10.2 ^{+8.5%}	23.7	14.9 ^{+37.1%}	11.2 ^{+52.7%}	173	275 ^{+59.0%}	366 ^{+112.9%}	
		32k	22.7	24.0 ^{+5.7%}	45.0	26.2 ^{+41.8%}	14.6 ^{+67.6%}	91	156 ^{+71.4%}	214 ^{+135.2%}	
Qwen2-VL-7B [35]	16	2k	2.1	2.3 ^{+9.5%}	10.1	6.8 ^{+32.7%}	5.2 ^{+48.5%}	811	1205 ^{+48.6%}	1575 ^{+94.2%}	
		8k	8.6	9.2 ^{+7.0%}	19.7	12.1 ^{+38.6%}	8.4 ^{+57.4%}	426	677 ^{+37.1%}	975 ^{+128.9%}	
		16k	19.2	20.1 ^{+4.7%}	27.0	15.6 ^{+42.2%}	9.7 ^{+64.1%}	303	525 ^{+73.3%}	845 ^{+179.0%}	

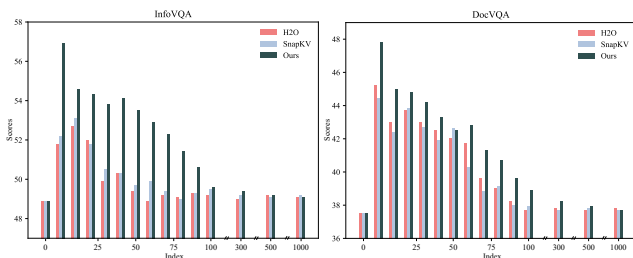


Figure 5. Comparison on InfoVQA [27] and DocVQA [28] by retaining only one visual token, which is selected based on the sorting of visual token importance scores using different methods.

while only adding 5% to prefill latency.

Detailed results of MMBench-Video. Table 10 presents the breakdown scores of different methods applied to the LLaVA-OV-7B [16] on the MMBench-Video [9] evaluation dataset. The proposed method outperforms existing methods in most subcategories, demonstrating its superior performance and stability. Notably, as the compression ratio

increases, the advantages of the proposed method become more pronounced, especially for perceptual items that are more sensitive to visual information. By more accurately assessing the importance of visual tokens, the proposed method retains the critical visual KV cache, thereby minimizing the loss of model performance.

D. Additional Ablation Studies

KV Cache Merge vs. KV Cache Drop. Table 11 compares the model performance using merge and drop strategies for KV cache compression. The drop strategy clearly outperforms the merge strategy, with performance decline becoming more significant as the merge proportion increases. This phenomenon persists across different VQA evaluation sets and various compression ratios, indicating that directly dropping less important visual KV cache in LLMs is a wiser choice. While the drop operation results in loss of visual information in the corresponding KV cache, the full to-

Table 10. The detailed comparison of the KV cache compression methods on MMBench-Video [9]. CP (coarse perception), FP-S (single-instance fine-grained perception), FP-C (cross-instance fine-grained perception), HL (Hallucination), LR (logic reasoning), AR (attribute reasoning), RR (relation reasoning), CSR (commonsense reasoning), TR (temporal reasoning).

Ratio	Methods	Overall	Perception	Reasoning	CP	FP-S	FP-C	HL	LR	AR	RR	CSR	TR
	Full	1.81	1.86	1.70	1.90	1.94	1.70	0.81	1.63	1.84	1.64	1.85	1.57
50%	H2O [42]	1.66	1.68	1.56	1.76	1.74	1.66	0.71	1.56	1.76	1.65	1.63	1.42
	Elastic [25]	1.68	1.71	1.60	1.79	1.75	1.66	0.73	1.58	1.77	1.67	1.66	1.44
	PrefixKV [34]	1.72	1.75	1.63	1.84	1.79	1.68	0.82	1.58	1.79	1.68	1.68	1.48
	SnapKV [19]	1.75	1.80	1.67	1.88	1.88	1.59	0.77	1.58	1.83	1.73	1.85	1.47
	Ours	1.80	1.84	1.69	1.89	1.94	1.66	0.81	1.61	1.84	1.77	1.79	1.53
10%	H2O [42]	1.63	1.64	1.61	1.72	1.71	1.52	0.74	1.54	1.81	1.64	1.73	1.43
	Elastic [25]	1.62	1.64	1.58	1.77	1.67	1.51	0.79	1.46	1.71	1.61	1.68	1.48
	PrefixKV [34]	1.68	1.74	1.58	1.81	1.73	1.63	0.80	1.56	1.79	1.68	1.69	1.47
	SnapKV [19]	1.70	1.76	1.60	1.83	1.83	1.62	0.82	1.55	1.70	1.67	1.78	1.44
	Ours	1.78	1.80	1.65	1.85	1.91	1.64	0.82	1.58	1.80	1.75	1.79	1.48
1%	H2O [42]	1.47	1.45	1.48	1.62	1.51	1.40	0.72	1.20	1.51	1.45	1.55	1.42
	Elastic [25]	1.51	1.52	1.51	1.65	1.54	1.40	0.75	1.22	1.54	1.50	1.60	1.45
	PrefixKV [34]	1.50	1.48	1.50	1.63	1.52	1.39	0.74	1.20	1.51	1.48	1.59	1.44
	SnapKV [19]	1.55	1.56	1.54	1.70	1.60	1.45	0.82	1.30	1.63	1.57	1.69	1.50
	Ours	1.67	1.70	1.58	1.78	1.76	1.65	0.72	1.58	1.76	1.65	1.64	1.52

Table 11. The results that the dropped KV cache is merged with the nearest preserved KV cache at different proportions. 100% means complete merging is used, while 0% means complete dropping is used.

Datasets	Ratio	100%	80%	60%	40%	20%	0%
ChatQA [26]	10%	76.4	77.0	77.4	78.6	79.2	79.9
	1%	72.2	73.6	74.2	75.5	76.1	76.4
InfoVQA [27]	10%	61.1	63.7	64.5	64.9	65.5	65.7
	1%	56.4	57.6	58.8	60.4	61.8	62.5
DocVQA [28]	10%	80.3	81.6	82.7	83.8	84.3	85.5
	1%	67.0	68.9	70.5	71.8	72.6	73.2
TextVQA [31]	10%	68.8	70.6	72.2	73.7	74.6	75.3
	1%	59.7	61.8	63.2	65.3	66.6	67.1

Table 12. Comparison of results across different evaluation sets and compression ratios with varying relevance thresholds.

Datasets	Ratio	0.99	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6
ChatQA [26]	10%	78.4	79.8	79.9	79.6	78.7	78.6	78.1	77.5	77.7
	1%	75.8	76.5	76.4	76.2	75.5	74.1	73.3	72.9	71.5
InfoVQA [27]	10%	65.1	65.5	65.7	65.7	65.0	63.7	62.2	61.6	59.4
	1%	61.9	62.3	62.5	62.4	61.2	58.6	55.4	53.7	52.2
DocVQA [28]	10%	82.4	85.6	85.5	84.7	81.3	78.8	76.1	75.5	74.8
	1%	71.2	72.8	73.2	72.9	70.0	65.5	61.8	58.9	55.6
TextVQA [31]	10%	73.4	75.0	75.3	75.4	74.6	74.0	73.3	72.5	71.4
	1%	65.5	67.2	67.1	66.8	65.2	64.7	63.5	62.4	61.2

ken attention interaction during the prefill stage allows the remaining important visual tokens and text tokens to potentially absorb this missing information. This absorption helps mitigate the information loss caused by dropping KV cache during subsequent decoding. Conversely, although the merge operation appears to preserve all visual information, the model lacks the ability to decode the original visual information from the merged visual KV cache. This operation may disrupt the representation of important vi-

Table 13. Layer-wise budget Jensen-Shannon divergence across different datasets and compression ratios, where Avg. corresponds to using the average allocation strategy.

Ratio	Setting	ChatQA [26]	InfoVQA [27]	DocVQA [28]	TextVQA [31]
10%	Avg. & s_t	0.46	0.54	0.61	0.57
	Avg. & s_k	0.53	0.49	0.58	0.55
	s_t & s_k	0.62	0.58	0.63	0.52
1%	Avg. & s_t	0.51	0.58	0.63	0.59
	Avg. & s_k	0.57	0.52	0.60	0.58
	s_t & s_k	0.68	0.61	0.59	0.57

sual information, ultimately leading to a decline in model performance.

Relevance Threshold α . Table 12 compares model performance under different relevance threshold. A relevance threshold around 0.9 achieves the best overall performance across various evaluation sets and compression ratios. If the relevance threshold is set too high or too low, it can lead to incomplete expression of instruction information or the inclusion of noise, respectively. This degrades the quality of visual KV cache importance assessment, thereby affecting the model’s performance after KV cache compression. Comparing a higher relevance threshold with a lower one reveals that introducing more noisy text significantly affects model performance. This emphasizes the importance of filtering out irrelevant text tokens within the observation window when compressing KV cache in LVLMS.

The Consistency of Strength and Skewness. To illustrate the difference between the dynamic budget and the average budget derived from the distribution strength and skewness used in this method, we recorded the budget distribution differences across various evaluation instances. The differences are quantified using Jensen-Shannon (JS) divergence, which ranges from 0 to 1. A JS divergence closer to 0 in-

User:

Social Media Platform Cheat Sheet

<p style="text-align: center;">Pinterest</p> <p>Key differentiating trait: More of a niche platform, despite very large usage. Great for businesses that involve planning, design, creation, making. Heavily female audience.</p> <p>Best for: Link clicks</p> <p>Resource and expertise level needed: High</p> <p>Good for these business types: Restaurants, Interior design, Wedding venues</p>	<p style="text-align: center;">Twitter</p> <p>Key differentiating trait: Fast-moving, real-time, news-driven platform. Strong customer service element, be prepared to engage and respond to customers here.</p> <p>Best for: Link clicks, engagement, and foot traffic</p> <p>Resource and expertise level needed: Low</p> <p>Good for these business types: Retail, Nonprofits, Newspapers & media outlets</p>
<p style="text-align: center;">LinkedIn</p> <p>Key differentiating trait: Unlike almost every other platform, which is designed for the user to focus on their personal interests, this is a profession-driven platform.</p> <p>Best for: Link clicks, event registrations</p> <p>Resource and expertise level needed: Medium</p> <p>Good for these business types: Luxury goods & services, Software providers, B2B Companies</p>	<p style="text-align: center;">Instagram</p> <p>Key differentiating trait: There are very few ways to link from Instagram, so this is a brand-building platform. You need to have lots of high-quality images and videos to support this platform.</p> <p>Best for: Foot traffic, word-of-mouth awareness, engagement</p> <p>Resource and expertise level needed: Medium to high</p> <p>Good for these business types: Bakeries & Coffee Shops, Travel agencies, Art Museums</p>

Which three business types is Pinterest good for?

LLaVA-OV w/ Full Cache: Restaurants, Interior design, Wedding venues

LLaVA-OV w/ H2O: Restaurants, Hotels, Retail

LLaVA-OV w/ Elastic: Restaurants, Hotels, Retail

LLaVA-OV w/ AirCache: Restaurants, Interior design, Wedding venues

Figure 6. Chat example applying KV cache compression methods on LLaVA-OV-7B [16].

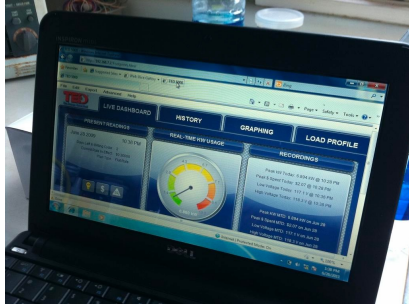
indicates smaller differences between the two distributions, while a value closer to 1 indicates larger differences. As shown in Table 13, the hierarchical budgets allocated based solely on the strength and skewness of the importance distribution are similar to those with an average allocation, indicating a complementary relationship. The former examines the layer’s emphasis on visual information, while the latter focuses on the layer’s ability to understand and interpret visual information. Combining both approaches can lead to better model performance.

Ablation Results of the Visual KV Cache Importance Evaluation. Figure 5 shows a comparison of selecting a single visual token based on the importance ranking of visual KV cache obtained by various methods on InfoVQA [27] and DocVQA [28]. As the importance of the selected visual token decreases, the performance of our proposed method also decreases reasonably. Additionally, for the same importance ranking, the performance of our proposed method is superior to that of existing methods.

E. Visualization of Chat Generation

Figures 6, 7, 8, and 9 illustrate a comparison of different methods applied to real chat generation while retaining only 1% of the visual KV cache. It is evident that the answers generated by the proposed method are more accurate.

User:



What date is under present readings?

LLaVA-OV w/ Full Cache: june 28, 2009

LLaVA-OV w/ H2O: june 20

LLaVA-OV w/ Elastic: june 28

LLaVA-OV w/ AirCache: june 28, 2009

Figure 7. Chat example applying KV cache compression methods on LLaVA-OV-7B [16].

User:

11:14 to 11:39 a.m.	Coffee Break Coffee will be served for men and women in the lobby adjacent to exhibit area. Please move into exhibit area. (Exhibits Open)
11:39 a.m.	TRRF GENERAL SESSION (PART I) Presiding: Lee A. Waller TRRF Vice President
11:39 to 11:44 a.m.	"Introductory Remarks" Lee A. Waller, TRRF Vice President
11:44 a.m. to 12:25 p.m.	Individual Interviews with TRRF Public Board Members and Scientific Advisory Council Members Conducted by TRRF Treasurer Philip G. Kuehn to get answers which the public refrigerated warehousing industry is looking for. Plus questions from the floor. Dr. Emil M. Mrak, University of California, Chairman, TRRF Board; Sam R. Cecil, University of Georgia College of Agriculture; Dr. Stanley Charn, Tufts University School of Medicine; Dr. Robert H. Cotton, ITT Continental Baking Company; Dr. Owen Fennema, University of Wisconsin; Dr. Robert E. Hardenburg, USDA.
12:25 to 12:58 p.m.	Questions and Answers
12:58 to 4:00 p.m.	Exhibits Open Capt. Jack Stoney Room
2:00 to 5:00 p.m.	TRRF Scientific Advisory Council Meeting Ballroom Foyer

What time is 'question and answers' session?

LLaVA-OV w/ Full Cache: 12:25 to 12:58 p.m.

LLaVA-OV w/ H2O: 12:25-12:38

LLaVA-OV w/ Elastic: 12:25-12:50

LLaVA-OV w/ AirCache: 12:25 to 12:58 p.m.

Figure 8. Chat example applying KV cache compression methods on LLaVA-OV-7B [16].


User:	 <p data-bbox="397 724 673 745">What is unusual about this image?</p>
LLaVA-OV w/ Full Cache:	<p>The image shows a man ironing clothes on the back of a moving vehicle, which is an unusual and potentially dangerous activity. Typically, ironing is done indoors, away from traffic and moving vehicles, to ensure safety and focus on the task. The man's actions in this image are unconventional and could pose risks to himself, other road users, and the passengers in the vehicle.</p>
LLaVA-OV w/ H2O:	<p>The image shows a person ironing clothes outside on a street, which is an unusual and unexpected activity. Typically, ironing is done indoors, in a home or a dedicated laundry area, to protect the clothes from outdoor elements and to maintain privacy. Performing this task outdoors breaks with conventional norms and could be due to various reasons such as a lack of indoor space, a desire for fresh air, or it might be part of a public demonstration or event.</p>
LLaVA-OV w/ Elastic:	<p>The image shows a person ironing clothes in an unusual setting: it appears to be inside a car, which is not a typical place for doing laundry. Ironing typically takes place in a home or a dedicated laundry room, where there is space and a suitable environment for the activity. The presence of an ironing board and iron in what looks like a vehicle's interior is unexpected and humorous.</p>
LLaVA-OV w/ AirCache:	<p>The image shows a man ironing clothes on the back of a moving vehicle, which is an unusual and unexpected sight. Typically, ironing is done indoors, usually on a stable surface like a table or ironing board, away from traffic and outdoor elements. The man's activity here poses risks to his safety and could also be considered a disruption to traffic flow.</p>

Figure 9. Chat example applying KV cache compression methods on LLaVA-OV-7B [16]. Important information is highlighted in red and blue.