

StreamKV: Streaming Video Question-Answering with Segment-based KV Cache Retrieval and Compression

Yilong Chen^{1,2,3}, Xiang Bai^{1,2,3}, Zhibin Wang^{3*}, Chengyu Bai^{1,2}, Yuhan Dai³, Ming Lu^{1†}

¹State Key Laboratory of Multimedia Information Processing, School of Computer Science, Peking University

²School of Software and Microelectronics, Peking University

³TaoBao & Tmall Group of Alibaba

chenyl@stu.pku.edu.cn

Abstract

Video Large Language Models (Video-LLMs) have demonstrated significant potential in the areas of video captioning, search, and summarization. However, current Video-LLMs still face challenges with long real-world videos. Recent methods have introduced a retrieval mechanism that retrieves query-relevant KV caches for question answering, enhancing the efficiency and accuracy of long real-world videos. However, the compression and retrieval of KV caches are still not fully explored. In this paper, we propose **StreamKV**, a training-free framework that seamlessly equips Video-LLMs with advanced KV cache retrieval and compression. Compared to previous methods that used uniform partitioning, StreamKV dynamically partitions video streams into semantic segments, which better preserves semantic information. For KV cache retrieval, StreamKV calculates a summary vector for each segment to retain segment-level information essential for retrieval. For KV cache compression, StreamKV introduces a guidance prompt designed to capture the key semantic elements within each segment, ensuring only the most informative KV caches are retained for answering questions. Moreover, StreamKV unifies KV cache retrieval and compression within a single module, performing both in a layer-adaptive manner, thereby further improving the effectiveness of streaming video question answering. Extensive experiments on StreamingVQA benchmarks demonstrate that StreamKV significantly outperforms existing Online Video-LLMs, achieving superior accuracy while substantially improving both memory efficiency and computational latency.

Code — <https://github.com/soulpower/StreamKV>

Introduction

Recent advancements in large language models (LLMs)(Touvron et al. 2023a,b; Ouyang et al. 2022; OpenAI 2023; Bai et al. 2023) have significantly enhanced the capabilities of Video-LLMs across various video understanding tasks, such as video captioning(Qasim, Horsch, and Prasad 2023), search(Luo et al. 2021), and summarization(Alaa et al. 2024). Although significant progress has been made, most current Video-LLMs are

*Project Leader

†Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

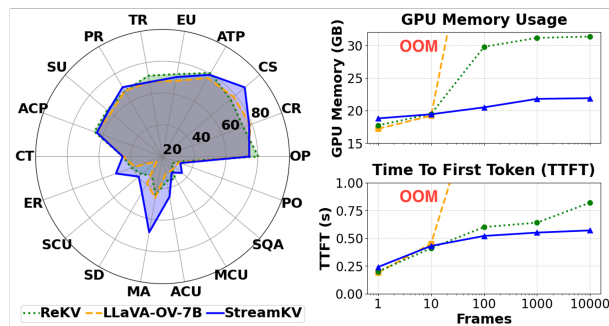


Figure 1: Comparison with ReKV on StreamingBench in terms of VideoQA accuracy, memory usage, and latency

primarily designed for offline usage. The model processes the entire video along with all questions at once, which leads to substantial GPU memory usage and increased response latency as video lengths grow.

Several pioneering studies have explored the paradigm of online Video-LLMs, where models continuously process video streams and answer user queries based on previous observations. Online Video-LLMs involve three key challenges: (1) How to efficiently and effectively process streaming video; (2) How to balance the retention of visual context with memory consumption; and (3) How to quickly and accurately retrieve relevant historical information when responding to user queries.

Recently, ReKV (Di et al. 2025) has introduced a retrieval mechanism that retrieves only query-relevant KV-caches, enhancing the efficiency and accuracy of question answering. However, this approach divides the video stream into uniform segments, disrupting the continuous semantic structure of the video and storing the entire historical visual context, leading to significant memory consumption. Moreover, its retrieval mechanism remains inflexible.

To this end, we propose StreamKV, a training-free framework that seamlessly equips Video-LLMs with advanced KV cache retrieval and compression. Unlike previous methods that uniformly segment video streams and risk arbitrarily breaking semantic boundaries, StreamKV dynamically partitions the video stream into semantic segments, which better preserves semantic information. For KV cache

retrieval, StreamKV calculates a summary vector for each segment to retain segment-level information. The KV pairs generated from this summary vector are excluded from further compression, which is crucial for accurately answering related questions. For KV cache compression, existing methods mainly target the decoding stage to reduce memory and computational overhead during inference. Although effective in offline settings, they are not suitable for StreamingVQA. In contrast, our approach treats semantic segments as the basic unit of compression and applies compression immediately after each segment is encoded. In the StreamingVQA scenario, since user questions are typically unknown when performing KV compression and multi-turn dialogues are expected, the compression process should focus on video semantics rather than specific user questions. To achieve this, StreamKV introduces a guidance prompt to capture the key semantic elements within each segment, ensuring only the most informative KV caches are retained for answering questions, thereby effectively removing redundancy. To further enhance both KV compression and retrieval, StreamKV introduces a unified layer-adaptive KV selection module, which dynamically allocates selection budgets across all transformer layers under an overall budget. This approach leverages the distinct information distributions in transformer layers instead of uniform allocation. This allows for optimal budget allocation, increasing the informative content retained under a fixed budget.

We compare StreamKV with the latest method, ReKV (Di et al. 2025), on StreamingBench (Lin et al. 2024) using the same foundation model. As shown in Figure 1, StreamKV outperforms ReKV in terms of VideoQA accuracy, memory usage, and latency.

In summary, our main contributions are as follows:

- We propose StreamKV, a training-free framework that seamlessly equips Video-LLMs with advanced KV cache retrieval and compression.
- To better preserve the semantic continuity of video content, StreamKV adopts a semantic partitioning and summary vector mechanism. This approach facilitates both subsequent compression and retrieval.
- To enable KV cache compression in streaming scenarios, we introduce a guidance prompt to capture key semantic elements within each segment, ensuring essential information is retained even under aggressive compression.
- To further improve KV cache retrieval and compression, we propose a Unified Layer-Adaptive KV Selection Module that allocates the selection budget optimally across transformer layers, maximizing informative content under a fixed total budget.
- Comprehensive experiments demonstrate that StreamKV significantly outperforms existing Online Video-LLMs, achieving superior accuracy while substantially improving both memory efficiency and latency.

Related Work

Video Large Language Models. With the rapid progress of Multimodal Large Language Models (MLLMs) (Li et al.

2023; Wang et al. 2024; Li et al. 2024a; Zhang et al. 2025b), Video LLMs (Zhang, Li, and Bing 2023; Lin et al. 2023; Li et al. 2024b; Zhang et al. 2025a) have drawn attention in recent years. Typically, these models use a visual encoder (Radford et al. 2021; Zhai et al. 2023; Tschannen et al. 2025) to extract video features, followed by a modality projector (e.g., MLP (Liu et al. 2023a) and Q-former (Li et al. 2023)) to map visual features into language space. The mapped features are combined with text tokens as input to the LLM to generate a contextual response. While exhibiting strong performance on offline video understanding benchmarks (Zhou et al. 2024; Fu et al. 2024), these models are intrinsically not well-suited for streaming video understanding owing to the memory bottleneck and complexity of information in long videos. To bridge this gap, our work equips Video-LLMs with streaming capabilities.

Streaming Video Understanding. Streaming video understanding (Wu et al. 2024; Li et al. 2025; Ding et al. 2025; Liu et al. 2025) requires Video LLMs to process real-time video frames and answer user questions based on all content up to a specified timestamp. VideoLLM-Online (Chen et al. 2024a) proposes the LIVE framework for streaming dialogue. However, it does not provide an effective solution for long-term video input processing. Subsequent works such as Flash-Vstream (Zhang et al. 2024a) and Dispidder (Qian et al. 2025) focus on managing complex video content and improving model performance through designed memory-augmented architectures. Recently, ReKV (Di et al. 2025) introduced a retrieval mechanism that retrieves query-relevant KV caches for question answering. However, it retains all generated KV caches, resulting in substantial memory consumption, and its retrieval strategy requires further optimization. To overcome these limitations, we propose StreamKV, which substantially improves memory efficiency and retrieval effectiveness, achieving strong performance in streaming video understanding.

KV Cache Compression for Video LLMs. Efficient KV cache compression (Cai et al. 2025; Yang et al. 2024; Wang et al. 2025) is essential in MLLMs to manage memory and latency overhead. FastV (Chen et al. 2024b) accelerates the prefill phase by pruning visual tokens in specific layers, utilizing attention scores from the final query token to guide the selection. Similarly, SparseVLM (Zhang et al. 2024b) employs cross-attention to identify visual tokens relevant to user queries. However, most existing methods (Li et al. 2024c; Zhang et al. 2023; Fu et al. 2025) are heavily based on a given user question, limiting their robustness and applicability in StreamingVQA scenarios. To address this challenge, StreamKV introduces a guidance prompt to capture the key semantic elements within each segment rather than relying on specific user questions.

Method

We first introduce the workflow of StreamKV, as illustrated in Figure 3. StreamKV partitions the video stream into semantic segments and calculates a summary vector for each segment. These segments are sequentially encoded to generate frame-level KV blocks, which are then compressed - except for those derived from the summary vector - and stored

in the KV Bank. Upon receiving a question, StreamKV retrieves query-relevant KV blocks from the KV Bank to generate responses. Both the compression and retrieval of KV blocks are performed using our proposed unified layer-adaptive selection module.

Semantic Segment Partitioning and Encoding

Semantic Segment Partitioning. We sample video frames at a regular interval and extract embedding $f_t \in \mathbb{R}^{P^2 \times D}$ for each frame using a vision encoder, where P^2 is the number of ViT patch tokens and D is the hidden dimension of ViT (Dosovitskiy et al. 2021). To detect significant visual changes and potential semantic boundaries, we compute the cosine similarity between adjacent embeddings:

$$s_t = \frac{f_{t-1} \cdot f_t}{\|f_{t-1}\| \|f_t\|}. \quad (1)$$

As illustrated in Figure 2a, frames with low similarity scores are identified as semantic boundaries, each indicating the start of a new segment.

To avoid excessively short segments, we apply an exclusion window of size m around each boundary, ensuring that the resulting segments are of sufficient length to contain relevant information. To limit segment length, we introduce a segment merging technique to exploit the temporal redundancies inherent in videos, as illustrated in Figure 2b. If the current segment exceeds a threshold M , we merge the most similar adjacent frame pair, based on previously computed cosine similarities (Xu et al. 2024; Jin et al. 2024).

Through semantic partitioning, the video stream is continually divided into a sequence of semantic segments $[S^i]$. Each segment is defined as $S^i := [f_t]_{t=1}^{T_i}$, where the segment length T_i satisfies $T_i \in [m, M]$. To preserve segment-level information, we calculate a summary vector $f_s^i = \frac{1}{T_i} \sum_{t=1}^{T_i} f_t^i$ for each segment by averaging frame-level features at each spatial location.

Segment-based Sliding-window Encoding. We encode the video stream \mathcal{V}^T incrementally, processing it segment by segment. At each step, given the current segment S^i , its summary vector f_s^i , and a local window of past KV pairs \mathbf{L} , we compute attention (Vaswani et al. 2017)

$$\mathbf{O} = \text{Attn}(\mathbf{W}_Q \mathbf{X}^i, [\mathbf{L}_k, \mathbf{W}_K \mathbf{X}^i], [\mathbf{L}_v, \mathbf{W}_V \mathbf{X}^i]), \quad (2)$$

where \mathbf{W}_Q , \mathbf{W}_K , and \mathbf{W}_V are the attention parameters, \mathbf{L}_k and \mathbf{L}_v denote the key and value vectors of \mathbf{L} , and $\mathbf{X}^i = [S^i \| f_s^i]$ is the concatenation of S^i and f_s^i .

For each segment S^i , we derive a collection of frame-level KV blocks. Specifically, the m -th KV block is defined as $\mathbf{b}_m^i = [(\mathbf{k}_{m,p}^i, \mathbf{v}_{m,p}^i)]_{p=1}^{P^2}$, where $\mathbf{k}_{m,p}^i$, $\mathbf{v}_{m,p}^i$ denote the patch-level key and value vectors of the m -th frame. The representative key vector for \mathbf{b}_m^i is computed by averaging its patch-wise key vectors:

$$\mathbf{r}_m^i = \frac{1}{P^2} \sum_{p=1}^{P^2} \mathbf{k}_{m,p}^i \in \mathbb{R}^{D'}. \quad (3)$$

For computational efficiency, we do not distinguish between attention heads and instead concatenate them into a single

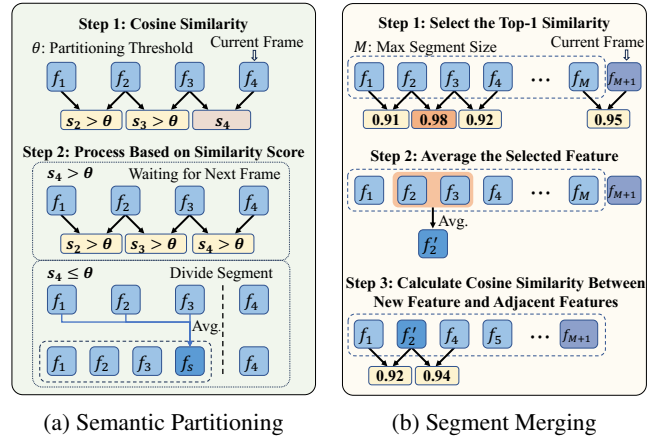


Figure 2: Video Segment Processing

vector of dimension D' . The summary vector f_s^i similarly yields a summary KV block \mathbf{b}_s^i and representative key \mathbf{r}_s^i .

We extend this notation to all L network layers: for each segment i and each layer l , we denote the frame-level KV blocks and representative key vectors as $\mathbf{b}_m^{i,l}$ and $\mathbf{r}_m^{i,l}$ respectively. Likewise, the summary KV block and its representative key vector are denoted as $\mathbf{b}_s^{i,l}$ and $\mathbf{r}_s^{i,l}$. At each layer l , the collection of KV blocks and their representative keys for segment i are given by:

$$\mathbf{B}_l^i := [\mathbf{b}_m^{i,l}]_{m=1}^{T_i}, \quad \mathbf{R}_l^i := [\mathbf{r}_m^{i,l}]_{m=1}^{T_i}. \quad (4)$$

Here, $\mathbf{b}_s^{i,l}$ is specifically preserved to retain segment-level information, while \mathbf{B}_l^i is subject to further compression.

Unified Layer-Adaptive KV Selection Module

We formulate both KV compression and retrieval as the problem of selecting the most relevant KV entries from a selection range according to a specific selection criterion (Lewis et al. 2021). Specifically, for compression, we use a guidance prompt to select the most informative KV blocks within a segment; for retrieval, we use user questions to select query-relevant KV blocks from the KV Bank. To efficiently address this, we introduce a Unified Layer-Adaptive KV Selection Module. Unlike uniform allocation strategies, this approach adaptively allocates selection budgets across transformer layers according to their information distributions, thereby increasing the overall retained content under a fixed budget.

Step-1: Calculate Cosine Similarity Scores. Given L transformer layers, for each layer l , let \mathbf{R}_l be the selection range, i.e., a sequence of candidate representative key vectors $[\mathbf{r}_j^l]_{j \in \text{idx}(\mathbf{R}_l)}$, where $\text{idx}(\mathbf{R}_l)$ denotes the indices of \mathbf{R}_l , and let \mathbf{c}^l represent the selection criterion vector for layer l . We define $\text{Sim}_l(j)$ as the cosine similarity between each candidate \mathbf{r}_j^l and the criterion \mathbf{c}^l .

The selection problem is thus formulated as selecting an index subset \mathcal{I}_l for each layer, corresponding to the top K_l candidates with the highest similarity,

$$\mathcal{I}_l = \text{Top } K_{\text{idx}}([\text{Sim}_l(j)]_{j \in \text{idx}(\mathbf{R}_l)}, K_l), \quad (5)$$

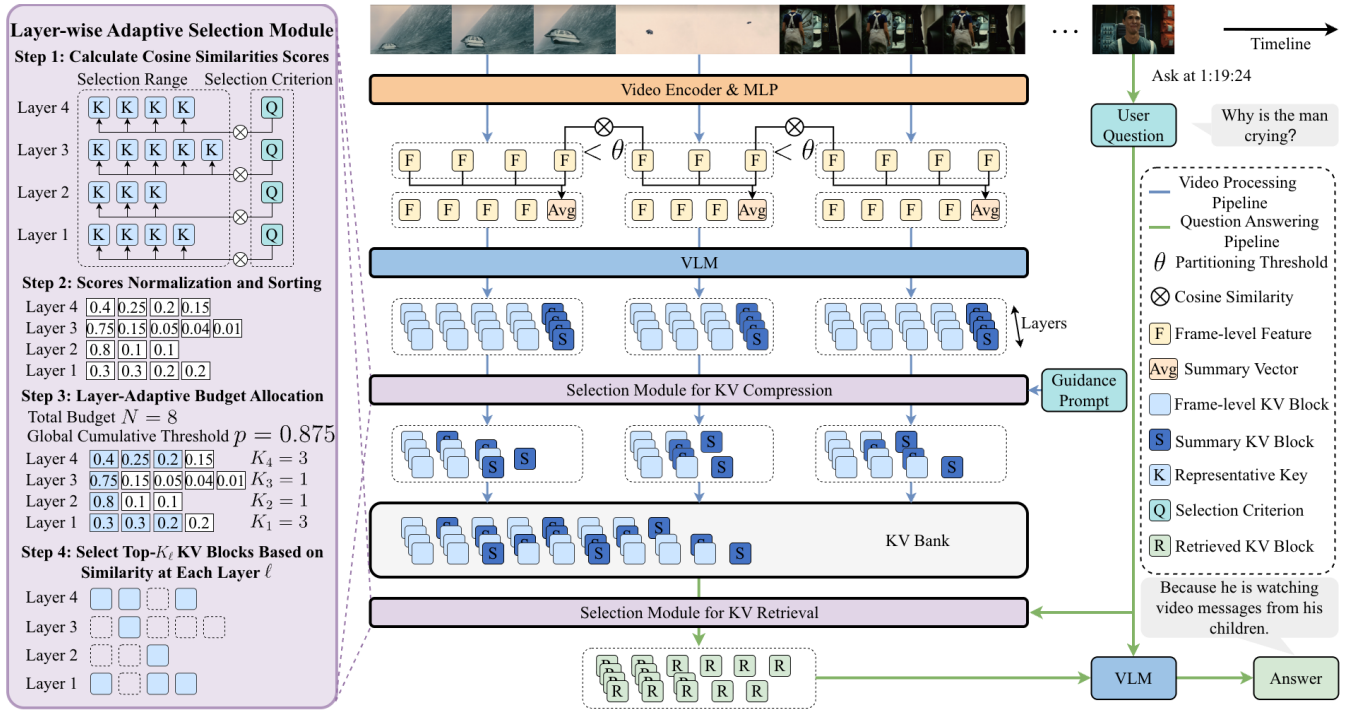


Figure 3: StreamKV workflow. StreamKV dynamically partitions video streams into semantic segments, and calculates a summary vector for each segment. These segments are sequentially encoded to generate frame-level KV blocks. KV compression is applied to all blocks except for those generated by the summary vector. We store the compressed KV blocks and summary KV block in the KV Bank. Upon receiving a question, StreamKV retrieves query-relevant KV blocks from the KV Bank to generate responses. Both the compression and retrieval are performed using our proposed unified layer-adaptive selection module.

where $K_l \leq |\mathbf{R}_l|$ and $\sum_{l=1}^L K_l = N$. Here, $|\mathbf{R}_l|$ is the size of \mathbf{R}_l and N denotes the total selection budget.

In contrast to uniform allocation, we adaptively assign K_l based on the similarity distribution in each layer, increasing the total informative content retained under the budget N .

Step-2: Scores Normalization and Sorting. For each layer l , we first calculate its normalized similarity score distribution (Rumelhart, Hinton, and Williams 1986):

$$\widetilde{\text{Sim}}_l(j) = \frac{\exp(\text{Sim}_l(j))}{\sum_{k=1}^{|\mathbf{R}_l|} \exp(\text{Sim}_l(k))}, \quad j \in \text{idx}(\mathbf{R}_l). \quad (6)$$

We then sort the normalized scores in descending order to derive the priority sequence $[\widetilde{\text{Sim}}_l(s_l(j))]_{j \in \text{idx}(\mathbf{R}_l)}$, where $s_l(j)$ represents the index of the j -th largest score in \mathbf{R}_l .

Step-3: Layer-Adaptive Budget Allocation. Let p be a global cumulative score threshold, which serves as an intermediate variable to facilitate the determination of the allocation. For each layer l , we define K_l^p as the minimal prefix length such that the cumulative sum of the top K_l^p normalized scores reaches or exceeds p :

$$K_l^p = \min\{k \mid \sum_{j=1}^k \widetilde{\text{Sim}}_l(s_l(j)) \geq p\}. \quad (7)$$

The global threshold p is determined such that the sum of selected candidates across all layers satisfies the total selection budget constraint, i.e., $\sum_{l=1}^L K_l^p = N$. This can be efficiently solved via a binary search over possible p values. Once the

Algorithm 1: Binary Search for Global Threshold p

- 1: **Input:** Total budget N ,
- 2: Priority sequence $[\widetilde{\text{Sim}}_l(s_l(j))]_{j \in \text{idx}(\mathbf{R}_l)}$, $\forall l$
- 3: Initialize $p_1 \leftarrow 0$, $p_2 \leftarrow 1$
- 4: **while** $p_2 - p_1 > \epsilon$ **do**
- 5: $p \leftarrow \frac{p_1 + p_2}{2}$
- 6: $K_l^p = \min\{k \mid \sum_{j=1}^k \widetilde{\text{Sim}}_l(s_l(j)) \geq p\}$, $\forall l$
- 7: $\delta \leftarrow \sum_l K_l^p - N$
- 8: **if** $\delta = 0$ **then return** p
- 9: **else if** $\delta < 0$ **then** $p_1 \leftarrow p$
- 10: **else** $p_2 \leftarrow p$
- 11: **end while**
- 12: **return** p

appropriate p is identified, we obtain the corresponding allocation $\{K_l^p\}_{l=1}^L$, which serves as the final budget allocation $\{K_l\}_{l=1}^L$. Finally, based on this budget allocation, we derive the selected index subsets $\{\mathcal{I}_l\}_{l=1}^L$ for each layer.

In summary, we encapsulate the above procedure into a unified selection function:

$$\{\mathcal{I}_l\}_{l=1}^L = \text{SelectKV}(\{\mathbf{R}_l, \mathbf{c}^l\}_{l=1}^L, N), \quad (8)$$

where SelectKV denotes the layer-adaptive KV selection module, with inputs $\{\mathbf{R}_l, \mathbf{c}^l\}_{l=1}^L$ (i.e., the candidate representative KV vectors and the corresponding selection criteria

for each layer) and the total selection budget N . The output is the set of selected representative indices $\{\mathcal{I}_l\}_{l=1}^L$.

KV Compression via Guidance Prompt

In StreamingVQA, video segments are encoded sequentially, and the resulting KV caches are stored for subsequent tasks. As input length increases, storing all KV caches becomes infeasible. Existing KV compression methods mainly target the decoding stage to reduce memory and computational overhead during inference. Although effective in offline settings, they are not suitable for StreamingVQA. In contrast, our compression method is applied immediately after each segment is encoded, retaining only the most informative KV caches from each segment in our KV bank, significantly reducing memory consumption. Moreover, since user questions are typically unknown when performing KV compression and multi-turn dialogues (Zhu, Wang et al. 2023) are expected, the compression process should focus on video semantics rather than specific user questions.

To address these requirements, StreamKV introduces a guidance prompt designed to capture the key semantic elements within each segment, such as **salient entities** (e.g., people, objects, locations, key visual concepts), **key events and actions** (what happened, when, and where), **temporal and causal relationships** (how events unfold and cause-effect chains), **contextual cues** (scene changes, dialogue, narrative shifts) and **important numerical or factual details** (for tasks like counting, summarization, or fact-based QA). This approach enables segments to preserve comprehensive visual semantics and contextual coherence, thereby ensuring that essential information is retained even under aggressive compression. For reference, we present examples of our guidance prompts in Appendix A.

For each segment, we use the unified layer-adaptive KV selection module to compress redundant KV blocks. Specifically, for each layer l , the selection range is the set of representative key vectors \mathbf{R}_l^i (as defined in Eq. (4)), and the selection criterion is the guidance prompt vector $\mathbf{g}^l = \frac{1}{N_g} \sum_{k=1}^{N_g} \mathbf{g}_k^l$, where N_g is the number of tokens in the guidance prompt, and \mathbf{g}_k^l is the k -th query vector for layer l . Here, $\mathbf{g}^l \in \mathbb{R}^{D'}$, where D' is the dimension of both the guidance vector and the representative key vectors in \mathbf{R}_l^i . Given the compression ratio θ , the total selection budget $N = \lceil (1 - \theta) T_i \rceil \times L$, where T_i is the frame count of segment i and L is the number of layers. The indices of the most informative KV blocks for each layer are selected as:

$$\{\mathcal{I}_l^i\}_{l=1}^L = \text{SelectKV}(\{\mathbf{R}_l^i, \mathbf{g}^l\}_{l=1}^L, N). \quad (9)$$

The compressed frame-level KV blocks and their representative key vectors for each segment and layer are then constructed as $\tilde{\mathbf{B}}_l^i = [\mathbf{b}_m^{i,l} \mid m \in \mathcal{I}_l^i]$ and $\tilde{\mathbf{R}}_l^i = [\mathbf{r}_m^{i,l} \mid m \in \mathcal{I}_l^i]$, where $[\cdot]$ denotes ordered concatenation. To preserve segment-level information, we also explicitly retain the segment summary KV block $\mathbf{b}_s^{i,l}$. For each layer l , the KV block bank \mathcal{B}_l and the set of its corresponding representative key vectors \mathcal{R}_l are continuously updated in parallel:

$$\mathcal{B}_l \leftarrow [\mathcal{B}_l, \tilde{\mathbf{B}}_l^i, \mathbf{b}_s^{i,l}], \quad \mathcal{R}_l \leftarrow [\mathcal{R}_l, \tilde{\mathbf{R}}_l^i, \mathbf{r}_s^{i,l}], \quad (10)$$

where $\tilde{\mathbf{B}}_l^i$ and $\mathbf{b}_s^{i,l}$ are the compressed KV blocks and the segment summary block, respectively, and $\tilde{\mathbf{R}}_l^i$ and $\mathbf{r}_s^{i,l}$ are their corresponding representative key vectors.

Streaming Video Question-Answering

KV Retrieval. Upon receiving a user question, we use the layer-adaptive KV Selection Module to retrieve query-relevant KV-caches. For each layer l , the selection range is the set of representative key vectors \mathcal{R}_l , and the selection criterion is the question vector $\mathbf{q}^l = \frac{1}{N_q} \sum_{k=1}^{N_q} \mathbf{q}_k^l \in \mathbb{R}^{D'}$, where N_q is the number of tokens in the question and \mathbf{q}_k^l is the k -th query vector at layer l . The total selection budget is set to $N = N_r \times L$, where N_r is the desired average number of KV blocks to retrieve for each layer. The indices of the retrieved KV blocks for each layer are obtained as:

$$\{\mathcal{I}_l\}_{l=1}^L = \text{SelectKV}(\{\mathcal{R}_l, \mathbf{q}^l\}_{l=1}^L, N). \quad (11)$$

The retrieved KV blocks are aggregated as:

$$\mathcal{P}_l = [\mathcal{B}_l[j] \mid j \in \mathcal{I}_l], \quad (12)$$

where \mathcal{P}_l denotes the set of KV blocks retrieved from the layer- l KV block bank \mathcal{B}_l , with $\mathcal{B}_l[j]$ representing the j -th block. The collection $\{\mathcal{P}_l\}_{l=1}^L$ is used for subsequent question answering.

Question-Answering Using Retrieved KV. The retrieved Video KV-caches $\{\mathcal{P}_l\}_{l=1}^L$ serve as the context for video question-answering. Formally, the attention calculation is formulated as:

$$\mathbf{O} = \text{Attn}(\mathbf{W}_Q \mathbf{X}, [\mathbf{C}_k, \mathbf{W}_K \mathbf{X}], [\mathbf{C}_v, \mathbf{W}_V \mathbf{X}]), \quad (13)$$

where \mathbf{X} represents either the question tokens or the current token being decoded, and \mathbf{C}_k and \mathbf{C}_v are the key and value vectors from the context, which includes the retrieved KV caches $\{\mathcal{P}_l\}_{l=1}^L$, question, and previously generated tokens. **Positional Encoding.** Rotary Position Embeddings (RoPE)(Su et al. 2023) are widely adopted in Video LLMs for temporal encoding, but their effectiveness degrades in long sequences due to suppressed attention between distant tokens. To alleviate this, StreamKV employs distinct positional encoding for video encoding and question answering. For segment encoding, inspired by LM-Infinite (Han et al. 2024), RoPE is applied only within the local window. For question answering, StreamKV treats retrieved tokens as consecutive and applies RoPE based on their relative positions rather than absolute positions.

Experiments

Implementation details

Experimental setup. We select the LLaVA-OneVision-Qwen2-7B-OV (Li et al. 2024a) as the baseline model due to its simplicity and strong performance. All experiments are conducted on NVIDIA H20 GPUs (96G) with FP16 precision. StreamKV processes video streams at 0.5 FPS, and the local window size is set to 15K, following the same experimental setup as ReKV. For dynamic partitioning, the minimum and maximum segment lengths are set to 4 and 64 frames, respectively, with a partitioning threshold of 0.99 employed to determine partitioning points. For KV-Cache retrieval, we set the number of retrieved frames to 8.

Model	Frames	Real-Time										Omni-Source				Contextual				Overall
		OP	CR	CS	ATP	EU	TR	PR	SU	ACP	CT	ER	SCU	SD	MA	ACU	MCU	SQA	PO	
Proprietary MLLMs																				
Gemini 1.5 pro	1 fps	79.0	80.5	83.5	79.7	80.0	84.7	77.8	64.2	72.0	48.7	46.8	39.6	74.9	80.0	51.4	40.7	54.8	45.1	67.1
GPT-4o	64	77.1	80.5	83.9	76.5	70.2	83.8	66.7	62.2	69.1	49.2	41.2	37.2	43.6	56.0	41.2	38.4	32.8	56.9	60.2
Claude 3.5 Sonnet	20	80.5	77.3	82.0	81.7	72.3	75.4	61.1	61.8	69.3	43.1	31.6	34.0	32.8	48.8	38.4	34.8	34.4	64.7	57.7
Open-Source Video MLLMs																				
Qwen2-VL-7B	0.2-1fps	75.2	82.8	73.2	77.5	68.3	71.0	72.2	61.2	61.5	46.1	41.2	22.0	32.8	43.6	31.2	26.0	39.6	22.7	54.1
MiniCPM-V-2.6-8B	32	71.9	71.1	77.9	75.8	64.6	65.7	70.4	56.1	62.3	53.4	40.8	24.0	34.0	41.2	34.0	31.6	41.9	22.2	53.9
InternVL-V2-8B	16	68.1	60.9	69.4	77.1	67.7	62.9	59.3	53.3	55.0	56.5	37.6	26.4	37.2	42.0	32.0	31.2	32.3	40.9	51.4
Kangaroo-7B	64	71.1	84.4	70.7	73.2	67.1	61.7	56.5	55.7	62.0	38.9	37.6	31.2	28.8	39.2	32.8	26.4	33.8	16.0	51.1
LongVA-7B	128	70.0	63.3	61.2	70.9	62.7	59.5	61.1	53.7	54.7	34.7	39.6	32.4	28.0	41.6	32.8	29.6	30.3	15.9	48.7
VILA-1.5-8B	14	53.7	49.2	71.0	56.9	53.4	53.9	54.6	48.8	50.1	17.6	41.6	26.4	28.4	36.0	26.8	34.0	23.2	17.7	43.2
Video-CCAM-14B	96	56.4	57.8	65.3	62.8	64.6	51.4	42.6	48.0	49.6	31.6	33.6	22.0	28.4	34.8	27.6	24.4	16.7	22.7	42.5
Video-LLaMA2-7B	32	55.9	55.5	57.4	58.2	52.8	43.6	39.8	42.7	45.6	35.2	30.4	32.4	30.4	36.0	24.8	26.8	18.7	0.0	40.4
LLaVA-OV-7B	32	80.4	74.2	76.0	80.7	72.7	71.7	67.6	65.5	65.7	45.1	40.8	37.2	33.6	44.8	35.6	36.0	27.3	29.6	56.4
Streaming MLLMs																				
Flash-VStream-7B	-	25.9	43.6	24.9	23.9	27.3	13.1	18.5	25.2	23.9	48.7	25.9	24.9	25.6	28.4	24.8	25.2	26.8	2.0	24.0
VideoLLM-online-8B	2fps	39.1	40.1	34.5	31.1	46.0	32.4	31.5	34.2	42.5	27.9	31.2	26.5	24.1	32.0	24.2	29.2	30.8	3.9	32.5
Dispider-7B	1fps	74.9	75.5	74.1	73.1	74.4	59.9	76.1	62.9	62.2	45.8	35.5	25.3	38.6	43.3	39.6	27.7	34.8	25.3	53.1
ReKV-7B	0.5fps	74.4	78.9	78.6	77.1	68.3	67.9	67.6	62.6	64.3	44.6	38.8	24.8	39.6	46.4	31.2	30.4	30.4	30.8	53.5
Ours																				
StreamKV-7B (↓60%)	0.5fps	74.7	78.1	87.7	79.4	70.8	67.6	70.4	64.6	64.0	45.1	51.2	39.7	46.0	68.4	45.6	31.7	36.0	32.0	58.9
StreamKV-7B (↓80%)	0.5fps	73.6	77.3	85.8	77.8	72.7	64.8	68.5	63.4	63.7	44.6	48.4	36.4	45.6	66.4	42.5	31.2	33.7	31.9	57.4
StreamKV-7B (↓90%)	0.5fps	73.8	77.3	85.9	77.5	73.3	63.9	69.4	61.4	63.2	35.8	48.4	36.4	44.0	66.1	43.6	30.0	33.2	31.4	56.7

Table 1: Performance comparison on StreamingBench.

Streaming Video Question Answering

We evaluate StreamKV on StreamingBench to assess its capability in streaming video understanding. Table 1 presents comprehensive question-answering accuracies on StreamingBench, covering 18 subtasks organized into three categories: Real-Time Visual Understanding, Omni-Source Understanding, and Contextual Understanding. StreamKV significantly outperforms existing Online Video-LLMs, achieving state-of-the-art performance even when retaining only 10% of the key-value pairs. Our experiments reveal several interesting observations:

1) In the Clips Summarization (CS) subtask, which involves summarizing the content of specific video segments, StreamKV with a 60% compression ratio achieves a notably high accuracy of 87.7%, representing an improvement of 9.1% over ReKV and 11.7% over the foundation LLaVA-OneVision model, and even outperforming all three proprietary MLLMs. This significant gain mainly stems from the semantic partitioning and summary vector mechanism, which effectively preserves key segment-level information, as well as our compression method, which effectively captures essential semantic elements within each segment.

2) For all four Omni-Source Understanding subtasks, StreamKV surpasses the performance of two proprietary MLLMs. These tasks evaluate a model’s ability to simul-

taneously process visual and audio content in video streams. While most existing Video LLMs cannot directly process audio, they instead infer visual scenes based on textual descriptions. StreamKV’s superior performance in these subtasks demonstrates its ability to accurately capture fine-grained visual information over lengthy video stream. This advantage primarily stems from our layer-adaptive KV retrieval strategy, which enables the precise retrieval of the most relevant video KVs from the KV bank.

3) Similarly, StreamKV achieves superior performance on the Anomaly Context Understanding (ACU) task, which assesses the ability of MLLMs to detect and accurately identify unusual or unexpected events within a video stream. These results indicate that StreamKV can effectively capture subtle semantic changes and reliably recognize anomalies, thereby enabling precise understanding in dynamic and unpredictable environments.

Ablations

Unless otherwise specified, all of our ablation studies are conducted on StreamingBench, and the reported scores represent the overall performance across three task categories.

Effectiveness of segment-level compression. To verify that our semantic partitioning approach benefits the compression process, we compare the performance of compression

Method	↓0%	↓10%	↓20%	↓30%	↓40%	↓50%	↓60%	↓70%	↓80%	↓90%
Uniform	60.49	59.57	58.94	58.57	57.65	57.32	56.33	55.56	53.02	51.41
Semantic	61.20	60.18	60.18	59.36	59.54	59.07	58.89	58.11	57.43	56.72

Table 2: Ablation results on semantic partitioning across varying compression rates.

Method	↓0%	↓10%	↓20%	↓30%	↓40%	↓50%	↓60%	↓70%	↓80%	↓90%
w/o. summary	60.52	59.46	59.25	58.44	57.95	57.42	56.21	54.66	54.65	53.85
w/. summary	61.20	60.18	60.18	59.36	59.54	59.07	58.89	58.11	57.43	56.72

Table 3: Ablation results on the impact of the summary vector across varying compression rates.

Com.	Ret.	↓50%	↓60%	↓70%	↓80%	↓90%
Uni.	Uni.	58.12	57.83	57.16	56.74	55.91
Ada.	Uni.	58.49	58.44	57.53	57.11	56.35
Uni.	Ada.	58.52	58.36	57.41	57.06	56.42
Ada.	Ada.	59.07	58.89	58.11	57.43	56.72

Table 4: Ablation results for the Layer-Adaptive KV Selection Module across varying compression rates.

on semantic segments versus uniform segments across various compression ratios. As shown in Table 2, compression performed on semantic segments consistently achieves superior performance across all compression ratios, indicating its greater effectiveness in preserving comprehensive visual information and contextual continuity.

Effectiveness of the segment summary vector. To assess the effectiveness of the summary vector, we evaluate StreamKV with and without incorporating the summary vector across various compression ratios. As illustrated in Table 3, the performance of StreamKV is significantly better when the summary vector is included. This result suggests that the summary vector effectively preserves segment-level information and maintains the structural integrity of the video content, thereby enabling the model to generate more accurate answers.

Effectiveness of Layer-Adaptive KV Selection Module. We evaluate our proposed layer-adaptive selection module for both KV compression and retrieval. As shown in Table 4, both compression and retrieval can use a uniform or adaptive selection budget across layers. The result demonstrates that the fully adaptive strategy consistently achieves the best performance across all compression ratios. Additionally, applying the adaptive strategy to either compression or retrieval alone also outperforms the fully uniform approach. This demonstrates the effectiveness of our layer-adaptive retrieval module.

Number of retrieved frames. We set the compression ratio at 60%, i.e., discarding 60% of the KV caches, and evaluate the impact of varying the number of retrieved frames. As illustrated in Figure 4, we observe that increasing the number of retrieved frames in StreamKV actually leads to a decreased performance. This suggests that StreamKV is able to accurately retrieve the most relevant frames, and retrieving additional frames introduces extra irrelevant information (Liu et al. 2023b), which can in turn hinder the sub-

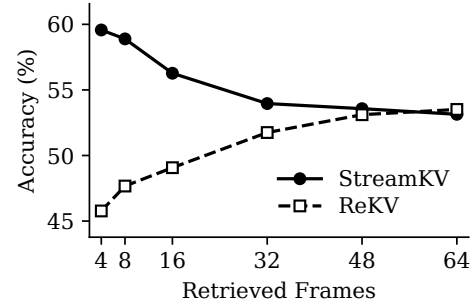


Figure 4: Accuracy versus Retrieved Frames Comparison for StreamKV and ReKV.

sequent question answering process. In contrast, our experiments on ReKV show the opposite trend. ReKV needs to retrieve more frames to ensure that relevant information is included, due to its lower retrieval precision. Moreover, retrieving fewer frames reduces computational overhead during the question answering stage, thereby accelerating inference. This demonstrates the effectiveness of our precise retrieval strategy and efficient inference.

Conclusion

In this paper, we introduced **StreamKV**, a training-free framework designed to address key challenges in streaming video understanding, including context preservation, long-term memory bottlenecks, and precise retrieval. By leveraging semantic partitioning and summary vectors, StreamKV effectively preserves the semantic continuity of video content. For efficient memory management, it employs a guidance prompt to identify and retain only the most informative KV caches within each segment. Furthermore, StreamKV proposes a Unified Layer-Adaptive KV Selection Module to further improve compression and retrieval. Extensive experiments on the StreamingBench benchmark demonstrate that StreamKV significantly outperforms existing online Video-LLMs in accuracy while substantially improving memory efficiency and reducing latency. Our work presents a practical and effective solution for building powerful and efficient Online Video-LLMs, paving the way for more robust real-world applications.

Acknowledgements

This work was supported by the Beijing Natural Science Foundation (L252060), and by the National Natural Science Foundation of China (62476011).

References

- Alaa, T.; Mongy, A.; Bakr, A.; Diab, M.; and Gomaa, W. 2024. Video Summarization Techniques: A Comprehensive Review. *arXiv:2410.04449*.
- Bai, Y.; Wang, J.; Zhu, Z.; et al. 2023. Qwen-72B: A Powerful Language Model at Scale. *arXiv preprint arXiv:2311.15001*.
- Cai, Z.; Zhang, Y.; Gao, B.; Liu, Y.; Li, Y.; Liu, T.; Lu, K.; Xiong, W.; Dong, Y.; Hu, J.; and Xiao, W. 2025. PyramidKV: Dynamic KV Cache Compression based on Pyramidal Information Funneling. *arXiv:2406.02069*.
- Chen, J.; Lv, Z.; Wu, S.; Lin, K. Q.; Song, C.; Gao, D.; Liu, J.-W.; Gao, Z.; Mao, D.; and Shou, M. Z. 2024a. Videollm-online: Online video large language model for streaming video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18407–18418.
- Chen, L.; Zhao, H.; Liu, T.; Bai, S.; Lin, J.; Zhou, C.; and Chang, B. 2024b. An Image is Worth 1/2 Tokens After Layer 2: Plug-and-Play Inference Acceleration for Large Vision-Language Models. *arXiv:2403.06764*.
- Di, S.; Yu, Z.; Zhang, G.; Li, H.; Cheng, H.; Li, B.; He, W.; Shu, F.; Jiang, H.; et al. 2025. Streaming Video Question-Answering with In-context Video KV-Cache Retrieval. In *ICLR*.
- Ding, X.; Wu, H.; Yang, Y.; Jiang, S.; Bai, D.; Chen, Z.; and Cao, T. 2025. StreamMind: Unlocking Full Frame Rate Streaming Video Dialogue through Event-Gated Cognition. *arXiv preprint arXiv:2503.06220*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv:2010.11929*.
- Fu, C.; Dai, Y.; Luo, Y.; Li, L.; Ren, S.; Zhang, R.; Wang, Z.; Zhou, C.; Shen, Y.; Zhang, M.; et al. 2024. Video-MME: The First-Ever Comprehensive Evaluation Benchmark of Multi-modal LLMs in Video Analysis. *arXiv:2405.21075*.
- Fu, Y.; Cai, Z.; Asi, A.; Xiong, W.; Dong, Y.; and Xiao, W. 2025. Not All Heads Matter: A Head-Level KV Cache Compression Method with Integrated Retrieval and Reasoning. In *The Thirteenth International Conference on Learning Representations*.
- Han, C.; Wang, Q.; Xiong, W.; Chen, Y.; Ji, H.; and Wang, S. 2024. LM-Infinite: Simple On-the-Fly Length Generalization for Large Language Models.
- Jin, P.; Takanobu, R.; Zhang, W.; Cao, X.; and Yuan, L. 2024. Chat-univi: Unified visual representation empowers large language models with image and video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13700–13710.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; tau Yih, W.; Rocktäschel, T.; Riedel, S.; and Kiela, D. 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *arXiv:2005.11401*.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; et al. 2024a. Llava-onevision: Easy visual task transfer. *arXiv:2408.03326*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. *arXiv preprint arXiv:2301.12597*.
- Li, W.; Hu, B.; Shao, R.; Shen, L.; and Nie, L. 2025. LION-FS: Fast & Slow Video-Language Thinker as Online Video Assistant. *arXiv:2503.03663*.
- Li, X.; Wang, Y.; Yu, J.; Zeng, X.; Zhu, Y.; Huang, H.; Gao, J.; Li, K.; He, Y.; Wang, C.; et al. 2024b. Videochat-flash: Hierarchical compression for long-context video modeling. *arXiv preprint arXiv:2501.00574*.
- Li, Y.; Huang, Y.; Yang, B.; Venkitesh, B.; Locatelli, A.; Ye, H.; Cai, T.; Lewis, P.; and Chen, D. 2024c. SnapKV: LLM Knows What You are Looking for Before Generation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Lin, B.; Zhu, B.; Ye, Y.; Ning, M.; Jin, P.; and Yuan, L. 2023. Video-llava: Learning united visual representation by alignment before projection. *arXiv:2311.10122*.
- Lin, J.; Fang, Z.; Chen, C.; Wan, Z.; Luo, F.; Li, P.; Liu, Y.; and Sun, M. 2024. StreamingBench: Assessing the Gap for MLLMs to Achieve Streaming Video Understanding. *arXiv:2411.03628*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023a. Visual instruction tuning. *NeurIPS*.
- Liu, J.; Yu, Z.; Lan, S.; Wang, S.; Fang, R.; Kautz, J.; Li, H.; and Alvare, J. M. 2025. StreamChat: Chatting with Streaming Video. *arXiv:2412.08646*.
- Liu, N. F.; Lin, K.; Hewitt, J.; Paranjape, A.; Bevilacqua, M.; Petroni, F.; and Liang, P. 2023b. Lost in the Middle: How Language Models Use Long Contexts. *arXiv:2307.03172*.
- Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2021. CLIP4Clip: An Empirical Study of CLIP for End to End Video Clip Retrieval. *arXiv:2104.08860*.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*.
- Ouyang, L.; Wu, J.; Jiang, X.; et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Qasim, I.; Horsch, A.; and Prasad, D. K. 2023. Dense Video Captioning: A Survey of Techniques, Datasets and Evaluation Protocols. *arXiv:2311.02538*.
- Qian, R.; Ding, S.; Dong, X.; Zhang, P.; Zang, Y.; Cao, Y.; Lin, D.; and Wang, J. 2025. Dispider: Enabling Video LLMs with Active Real-Time Interaction via Disentangled Perception, Decision, and Reaction. *arXiv preprint arXiv:2501.03218*.

- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1986. Learning representations by back-propagating errors. *Nature*, 323(6088): 533–536.
- Su, J.; Lu, Y.; Pan, S.; Murtadha, A.; Wen, B.; and Liu, Y. 2023. RoFormer: Enhanced Transformer with Rotary Position Embedding. *arXiv:2104.09864*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023a. Llama: Open and efficient foundation language models. *arXiv:2302.13971*.
- Touvron, H.; et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Tschannen, M.; Gritsenko, A.; Wang, X.; Naem, M. F.; Alabdulmohsin, I.; Parthasarathy, N.; Evans, T.; Beyer, L.; Xia, Y.; Mustafa, B.; Hénaff, O.; Harmsen, J.; Steiner, A.; and Zhai, X. 2025. SigLIP 2: Multilingual Vision-Language Encoders with Improved Semantic Understanding, Localization, and Dense Features. *arXiv:2502.14786*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Wang, J.; Liu, Z.; Rao, Y.; and Lu, J. 2025. SparseMM: Head Sparsity Emerges from Visual Concept Responses in MLLMs. *arXiv:2506.05344*.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Fan, Y.; Dang, K.; Du, M.; Ren, X.; Men, R.; Liu, D.; Zhou, C.; Zhou, J.; and Lin, J. 2024. Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution. *arXiv preprint arXiv:2409.12191*.
- Wu, S.; Chen, J.; Lin, K. Q.; Wang, Q.; Gao, Y.; Xu, Q.; Xu, T.; Hu, Y.; Chen, E.; and Shou, M. Z. 2024. VideoLLM-MoD: Efficient Video-Language Streaming with Mixture-of-Depths Vision Computation. *arXiv:2408.16730*.
- Xu, M.; Gao, M.; Gan, Z.; Chen, H.-Y.; Lai, Z.; Gang, H.; Kang, K.; and Dehghan, A. 2024. SlowFast-LLaVA: A Strong Training-Free Baseline for Video Large Language Models. *arXiv:2407.15841*.
- Yang, D.; Han, X.; Gao, Y.; Hu, Y.; Zhang, S.; and Zhao, H. 2024. PyramidInfer: Pyramid KV Cache Compression for High-throughput LLM Inference. *arXiv:2405.12532*.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid loss for language image pre-training. In *ICCV*.
- Zhang, B.; Li, K.; Cheng, Z.; Hu, Z.; Yuan, Y.; Chen, G.; Leng, S.; Jiang, Y.; Zhang, H.; Li, X.; et al. 2025a. VideoLLaMA 3: Frontier Multimodal Foundation Models for Image and Video Understanding. *arXiv preprint arXiv:2501.13106*.
- Zhang, H.; Gao, M.; Gan, Z.; Dufter, P.; Wenzel, N.; Huang, F.; Shah, D.; Du, X.; Zhang, B.; Li, Y.; et al. 2025b. Mm1. 5: Methods, analysis & insights from multimodal llm fine-tuning. *ICLR*.
- Zhang, H.; Li, X.; and Bing, L. 2023. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv:2306.02858*.
- Zhang, H.; Wang, Y.; Tang, Y.; Liu, Y.; Feng, J.; Dai, J.; and Jin, X. 2024a. Flash-vstream: Memory-based real-time understanding for long video streams. *arXiv preprint arXiv:2406.08085*.
- Zhang, Y.; Fan, C.-K.; Ma, J.; Zheng, W.; Huang, T.; Cheng, K.; Gudovskiy, D.; Okuno, T.; Nakata, Y.; Keutzer, K.; and Zhang, S. 2024b. SparseVLM: Visual Token Sparsification for Efficient Vision-Language Model Inference. *arXiv:2410.04417*.
- Zhang, Z.; Sheng, Y.; Zhou, T.; Chen, T.; Zheng, L.; Cai, R.; Song, Z.; Tian, Y.; Re, C.; Barrett, C.; Wang, Z.; and Chen, B. 2023. H2O: Heavy-Hitter Oracle for Efficient Generative Inference of Large Language Models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Zhou, J.; Shu, Y.; Zhao, B.; Wu, B.; Xiao, S.; Yang, X.; Xiong, Y.; Zhang, B.; Huang, T.; and Liu, Z. 2024. Mlvu: A comprehensive benchmark for multi-task long video understanding. *arXiv:2406.04264*.
- Zhu, L.; Wang, R.; et al. 2023. A Survey on Multimodal Conversation Modeling. *ACM Computing Surveys (CSUR)*.