

# DyCoke : Dynamic Compression of Tokens for Fast Video Large Language Models

Keda Tao<sup>1,2</sup> Can Qin<sup>3</sup> Haoxuan You<sup>4</sup> Yang Sui<sup>5</sup> Huan Wang<sup>1,\*</sup>

<sup>1</sup>Westlake University <sup>2</sup>Xidian University

<sup>3</sup>Salesforce AI Research <sup>4</sup>Columbia University <sup>5</sup>Rice University

<https://github.com/KD-TAO/DyCoke>

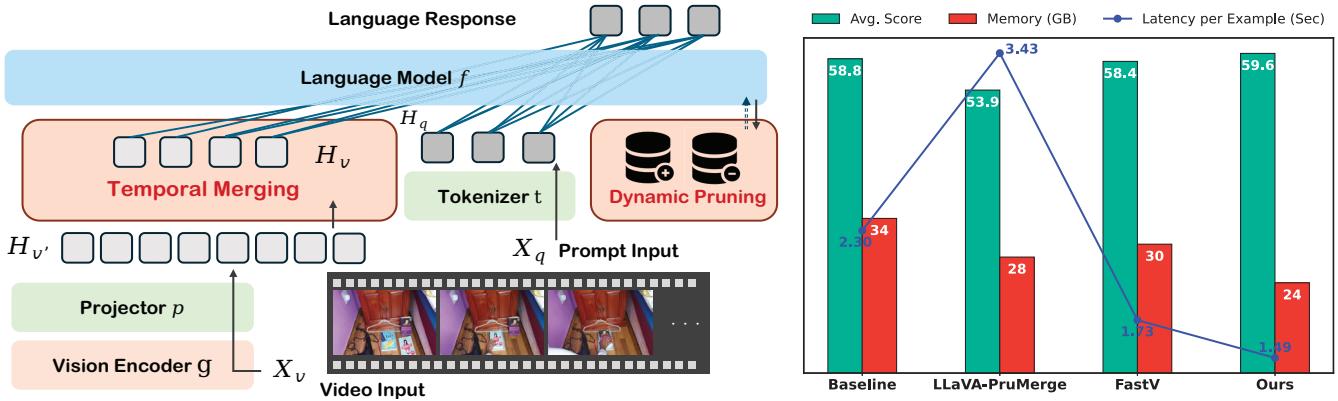


Figure 1. **Left:** We introduce *DyCoke* (dynamic compression of tokens), a training-free token compression method for fast video large language models. The key innovation of *DyCoke* over its predecessors is to *dynamically* remove redundant tokens during the decoding stage, squeezing both the temporal (video frames) and spatial redundancy in visual tokens. **Right:** Efficiency and performance comparison of various training-free token pruning methods on MVBench [23] with LLaVA-OV-7B [18]. *DyCoke* surpasses the SoTA counterparts (PruMerge [39], FastV [3]), with 1.5× inference speedup and a 1.4× reduction in memory usage relative to the baseline, while simultaneously enhancing performance.

## Abstract

*Video large language models (VLLMs) have significantly advanced recently in processing complex video content. Yet, their inference efficiency remains constrained because of the high computational cost stemming from the thousands of visual tokens generated from the video inputs. We empirically observe that, unlike single image inputs, VLLMs typically attend visual tokens from different frames at different decoding iterations. This makes a one-shot pruning strategy prone to removing important tokens by mistake. Motivated by this, we present DyCoke, a training-free token compression method to optimize token representation and accelerate VLLMs. DyCoke incorporates a plug-and-play temporal compression module to minimize temporal redundancy by merging redundant tokens across frames and applying dynamic KV cache reduction to prune spatially redundant tokens selectively. It ensures high-quality inference by dynamically retaining the critical tokens at each decoding step. Extensive experimental results demonstrate that DyCoke can outperform the prior SoTA counterparts, achieving 1.5× inference speedup, and 1.4× memory reduction against the baseline VLLM, while still improving the performance, with no training.*

## 1. Introduction

Video large language models (VLLMs) have advanced significantly in understanding diverse video contexts, primarily because of their enhanced reasoning ability for complex multimodal information [4, 18, 22–25, 43, 49, 52]. Most current VLLMs rely on sequential visual representations. When dozens of video frames are fed into the language model along with a language prompt, the video input is converted into tens of thousands of tokens. Because of the quadratic complexity of the attention mechanism, the inherent visual redundancy across video frames leads to a dramatic surge of computational complexity, resulting in prohibitive training and inference costs. Previous research [5, 6, 51, 54] has largely focused on developing lighter large language models (LLMs) with fewer parameters; however, this often significantly diminishes the complex reasoning capabilities of these models. Consequently, how to reduce the number of video tokens while maintaining model performance emerges

\*Corresponding author: wanghuan@westlake.edu.cn

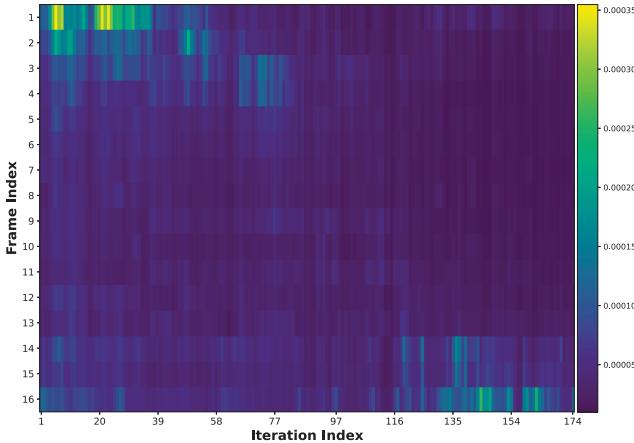


Figure 2. Attention score between the predicted token at different decoding iterations (x-axis) and the input video tokens (y-axis) at the decoding stage of LLaVA-OV-7B [18] (attention score averaged over all attention layers). Note, that some video tokens (*e.g.*, frame #1) become less important as the decoding proceeds, while others may instead become more important (*e.g.*, frame #16). This observation motivates us to develop *DyCoke*, a training-free plug-and-play token compression method that can *dynamically* exploit the token redundancy during decoding.

as a new research direction.

In previous works on token compression for image LLMs, the attention scores of visual tokens serve as the primary metric for assessing token importance, resulting in a single-stage pruning approach. For example, FastV [3] evaluates the distribution of attention between visual tokens and predicted tokens in the language model during the prefilling phase, leveraging the KV cache. In contrast, LLaVA-PruMerge [39] selects key visual tokens using attention scores derived from the CLIP visual encoder [35]. However, the substantial amount of long-term sequential information stored in videos results in considerable temporal and spatial redundancy [11, 36, 41]. We present the distribution of the (averaged) attention scores for each predicted token over the visual tokens in Figure 2. It shows that the overall distribution of attention scores between visual tokens is highly sparse, and the model’s focus shifts to different visual tokens as the decoding proceeds, resembling human attention patterns. Consequently, for video LLMs, unlike image inputs, a single-stage pruning strategy may result in incorrect token filtering, omission of key tokens, and temporal disarray, thereby compromising video comprehension.

Based on the above analysis, we reassess the need for a straightforward and effective token compression method tailored to VLLMs, using the significant temporal and spatial redundancy in video information as a foundational consideration. Thus, in this work, we present the first *temporal-spatial dynamic token compression method* (*DyCoke*) to optimize the token representation tailored for VLLMs, as illustrated in Figure 1 (left). The first phase involves designing

a plug-and-play, lightweight token compression module that addresses temporal redundancy by merging similar tokens across frames. We group consecutive frames by sampling and identifying tokens with overlapping information in adjacent or nearby frames for temporal merging. The second phase maintains a parsimonious KV cache established in the first phase by dynamically pruning less important information, reducing the spatial redundancy of visual tokens, while retaining pruned tokens for secondary activations needed for auxiliary computations. Specifically, our approach enables the model to dynamically select a *distinct* set of tokens at each decoding step, which is essential for preserving performance. Building on this, *DyCoke* maximizes reasoning ability while substantially reducing visual tokens, resulting in a more streamlined and representative visual token set.

Empirically, the proposed *DyCoke* demonstrates excellent performance in video reasoning tasks, simplifying visual tokens as much as possible while maintaining model performance, particularly with long encoded inputs, and does not require fine-tuning or parameter modifications. In the first stage, after temporal merging, redundant visual tokens can be adaptively reduced by 50% - 60%. In the second stage, each iteration can dynamically further reduce visual tokens by an additional 70% - 90% based on the first stage. On average, each video input frame retains 15 tokens for attention matrix calculation, greatly accelerating the inference process. As shown in Figure 1 (right), *DyCoke* achieves a  $1.54 \times$  inference speedup on LLaVA-OV-7B [18], with the lowest memory consumption and highest accuracy. Notably, the method is *training-free*.

Our contributions in this work are summarized as follows:

- We propose a plug-and-play temporal token merging block to effectively reduce visual tokens while preserving the key video content information by leveraging temporal redundancy between frames and merging similar tokens.
- We propose a dynamic KV cache token reduction method that dynamically reduces redundant tokens in the KV cache without relying on additional parameters or training, enabling efficient processing of long input sequences.
- Experimental results on several video inference benchmarks show that *DyCoke* maintains high inference accuracy and speed while compressing visual tokens and enables processing longer video sequences within the same computational budget.

## 2. Related Work

### 2.1. Video Large Language Models

With advances in large language models (LLMs) and their strong multi-modal understanding and reasoning capabilities, many studies have attempted to integrate LLMs with video encoders to leverage these powerful capabilities for video tasks [4, 16, 18, 19, 21–25, 28, 32, 43, 45, 48]. Representa-

tive works, such as VideoChat [22] and VideoLLaMA [24], use video converters to encode video features based on image LLMs, enhancing understanding capabilities by training on extensive video datasets. LLaVA-NeXT-Interleave [19] and LLaVA-OneVision [18] focus on achieving excellent performance across single-image, multi-image, and video scenarios. Although the potential of VLLMs for video understanding and reasoning is being realized, the tens of thousands of visual tokens required for long videos significantly increase both inference time and memory demands. While works such as VILA aim to optimize token usage, substantial hardware resources are still needed for model fine-tuning [16, 24, 27, 43]. Therefore, we desire a token compression method specifically for VLLMs that requires *no* fine-tuning.

## 2.2. Efficient Multi-Modal Large Language Models

While multi-modal large language models (MLLMs) have made significant progress [10, 12, 13, 15, 15, 20, 29, 30, 47, 47, 55], their large-scale training and deployment entail substantial computational costs. LLaVA-1.5 [7, 29, 38] addresses this by using 4-bit rather than 8-bit quantization for compression. MobileVLM [5] and MobileVLM-v2 [6] utilize compact architectures optimized for mobile applications. Additionally, TinyGPT-V [51], LLaVA-Phi [56], and Vary-toy [44] aim to match or exceed the performance of large models by using smaller LLM backbones, such as Phi-2 [14]. MoE-LLaVA [26] employs a mixture of experts to address model sparsity and improve both efficiency and performance. TinyLLaVA [54] explores more lightweight MLLM architectures and training optimizations. However, many of these works have noted that reducing the size of the LMM model backbone often compromises its reasoning capability. Improving LMM efficiency by compressing the number of visual tokens offers a promising alternative.

In previous studies, token pruning has been widely adopted to mitigate token redundancy in vision transformers (ViTs) and large language models (LLMs) [3, 57]. ToMe [1] proposes merging similar tokens within ViTs to consolidate redundant information while preserving task-relevant content across various domains like image, video, and audio processing. TESTA [36] achieves up to a 75% reduction in processed tokens through the use of temporal and spatial aggregation modules. TempMe [40] addresses temporal redundancy by progressively merging tokens across neighboring clips. FastV [2] enhances attention efficiency in MLLMs by pruning redundant image tokens based on attention scores, without requiring additional training. LLaVA-PruMerge [39] introduces adaptive token reduction by selecting key visual tokens based on attention scores derived from the CLIP visual encoder. Furthermore, Look-m [42] applies token merging strategies in the KV cache to decrease computational costs and support extended multimodal contexts. xGen-MM-Vid [37] maps a sequence of tokens across multiple frames

into a compact set of visual tokens, enabling fine-tuning with fewer visual tokens. LazyLLM [9] employs attention maps to progressively prune tokens, reducing time-to-first-token (TTFT). In our study, we present a novel, training-free dynamic token compression strategy specifically designed for VLLMs that fully accounts for the temporal characteristics of visual tokens to preserve model performance as effectively as possible.

## 3. Proposed Method

### 3.1. Background on Video LLM Inference

Video LLM inference typically comprises two stages: *prefilling* and *decoding*.

**(1) Prefilling Stage.** In the prefilling stage, for a video with  $M_v$  frames, the image encoder maps each frame into  $K_v$  embedding vectors,  $\mathbf{z}_i \in \mathbb{R}^{K_v \times d_v}$ , where  $d_v$  is the dimension of each embedding vector. The  $M_v$  frames thus form an embedding sequence  $\mathbf{Z}_v = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{M_v}]$ , which is fed into the projector, which maps visual tokens into the same feature space as text to facilitate information fusion and alignment across modalities. The projector output is processed to generate the set  $H_{v'} \in \mathbb{R}^{M_v N_v \times D}$  of visual tokens, where  $N_v$  is the token length corresponding to one frame of video.

Simultaneously, the model receives a token sequence prompt  $T = \{t_i\}_{i=1}^{N_q}$ , where  $t_i$  represents the  $i$ -th token of  $N_q$  tokens in total. We can get the set of text tokens  $H_q \in \mathbb{R}^{N_q \times D}$  where  $D$  is the dimension of the hidden state. Next, the visual tokens and text tokens are concatenated as LLM input,  $H = \text{concat}[H_{v'}, H_q]$ .

In each transformer layer  $l$  of an LLM, self-attention is applied to  $H$ . This process involves computing the query  $\mathbf{Q}^l$ , key  $\mathbf{K}^l$ , and value  $\mathbf{V}^l$  matrices using linear transformations. Specifically,

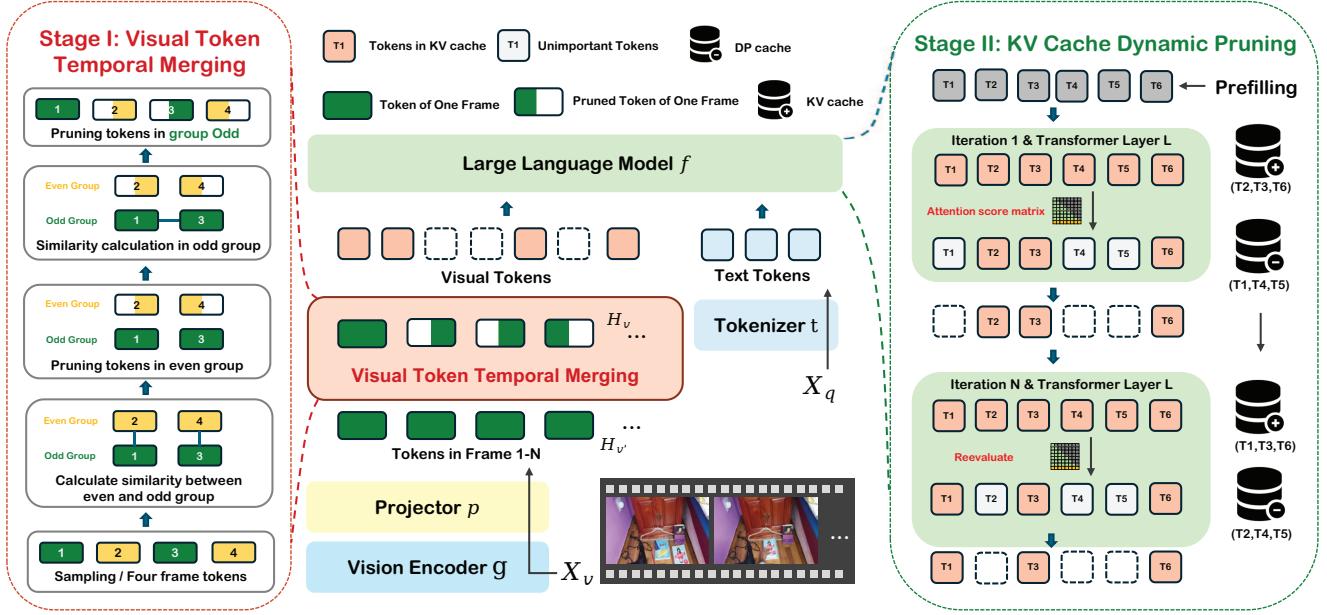
$$\mathbf{Q}^l = H \mathbf{W}_Q^l, \quad \mathbf{K}^l = H \mathbf{W}_K^l, \quad \mathbf{V}^l = H \mathbf{W}_V^l, \quad (1)$$

where  $\mathbf{W}_Q^l, \mathbf{W}_K^l, \mathbf{W}_V^l \in \mathbb{R}^{D \times D}$  are learnable projection matrices. These transformations project the input into a latent space where attention can be efficiently calculated. The  $K$  and  $V$  matrices are computed and subsequently stored in the KV cache to facilitate token generation during decoding.

**(2) Decoding Stage.** In the decoding phase, the model sequentially generates tokens by using and updating the KV cache stored during the prefilling phase. At each time step  $t$ , only the key and value of the *new* token  $h_t$  are computed, without recalculating the attention for the entire sequence. The  $\mathbf{K}$  and  $\mathbf{V}$  values calculated for the new token are updated in the KV cache:

$$\mathbf{K} = [\mathbf{K}, h_t \mathbf{W}_K], \quad \mathbf{V} = [\mathbf{V}, h_t \mathbf{W}_V], \quad (2)$$

which significantly reduces the computational load.



**Figure 3. Detailed overview of our DyCoke method.** DyCoke compresses visual tokens in VLLMs through a two-stage pruning process: *visual token temporal merging* (TTM) and *KV cache dynamic pruning*. Token temporal merging (illustrated in the red dashed box on the left) merges similar tokens in video frames at the prefilling stage, tapping into the temporal redundancy of the video input; KV cache dynamic pruning (illustrated in the blue dashed box on the right) further removes less attended visual tokens in the KV cache dynamically at the decoding stage, exploiting the spatial redundancy in visual tokens. DyCoke is a drop-in *training-free* approach to accelerate VLLMs.

### 3.2. Our Method: DyCoke

DyCoke employs a two-stage token compression strategy. The first stage merges visual tokens that exhibit significant temporal redundancy across frames, and the second stage then dynamically prunes visual tokens for attention weight calculation, building on the first stage. This approach aims to preserve model performance while simplifying the tokens as extensively as possible.

**(1) Visual Token Temporal Merging.** In VLLMs, the video input contains significant temporal redundancy. Some activities often persist across multiple frames with minimal visual change, while backgrounds and stationary objects frequently contain similar information across frames, causing substantial redundancy. Combining these redundant visual tokens at temporal scales can reduce the total token length of the input, which accelerates VLLM inference and decreases memory consumption. We introduce a plug-and-play token temporal merging (TTM) module as a first-stage solution to filter out consecutive, redundant visual tokens.

First, we assume that for an input visual token  $H'_v$ , the goal is to reduce  $k\%$  of tokens through the merge operation. To achieve this, we compute the similarity between all possible token pairs and merge those with the highest similarity. However, this approach significantly increases computational load and processing time, making the cost outweigh the benefits. Therefore, given the high incidence of temporal redundancy between adjacent frames, we employ

continuous sampling of visual tokens corresponding to the input video frames. As shown in Figure 3, TTM initially performs uniform sampling with a sliding window with a length of 4 frames, dividing tokens into groups  $O$  (*Odd*) and  $E$  (*Even*) and calculating token similarity between corresponding positions in adjacent groups. We use cosine similarity to calculate token similarity  $S$ :

$$S = \cos(\theta) = \frac{h_i \cdot h_j}{\|h_i\| \|h_j\|}. \quad (3)$$

We prune tokens in group  $E$  with high similarity to those in group  $O$ . We then calculate the similarity between frames within group  $O$ , retaining the full token of the first frame in the sampling window and pruning the remaining tokens. The pruning rate in TTM is set to  $k\%$ , and this process is repeated for each subsequent sampling window. Finally, the LLM input can be redefined as

$$H = \text{concat}[\text{TTM}(H'_v), H_q]. \quad (4)$$

At this stage, TTM enables visual marker reduction by leveraging temporal dependencies. The TTM module is simple, effective, and plug-and-play, with a negligible processing time of less than  $10^{-3}$  seconds for 32 input frames.

**(2) KV Cache Dynamic Pruning.** To further compress visual tokens, we analyze the distribution of average attention scores for the visual token represented by Figure 2 for each predicted token. Results indicate that the attention score of

the visual token for the next prediction token is *highly sparse*, suggesting substantial redundancy in the input visual token that can be safely pruned without impacting the next prediction. We also observe that each prediction token focuses on a different visual token at various decoding stages. This observation aligns with the human process of understanding long-sequence video information, leading us to consider a *dynamic* pruning scheme for the KV cache during decoding.

At the first decoding iteration, for an LLM with  $N_{lm}$  layers, we compute the cross-attention weights between the predicted token and the visual token at layer  $L$  to calculate the average attention score matrix:

$$\mathbf{A}^{(L)} = \text{Softmax} \left( \frac{\mathbf{Q}^{(L)}(\mathbf{K}^{(L)})^\top}{\sqrt{D}} \right), \quad (5)$$

where  $\mathbf{Q}^{(L)} \in \mathbb{R}^{1 \times d}$ . We then extract the attention scores of visual tokens and predicted tokens to form a subset  $\mathbf{A}_v^{(L)}$ . To obtain the top  $p\%$  attention scores in the cross-attention matrix  $\mathbf{A}_v^{(L)}$ , we calculate a threshold  $\tau$  and define the set  $\mathcal{I}_p^{(L)}$  comprising the indices of these top  $p\%$  attention scores. Then we prune the visual tokens in the KV cache, retain tokens with high attention scores, and update the KV cache:

$$\begin{aligned} \mathbf{K}_v^{(L)} &= \{\mathbf{K}_v^{(L)}[i] \mid i \in \mathcal{I}_p^{(L)}\}, \\ \mathbf{V}_v^{(L)} &= \{\mathbf{V}_v^{(L)}[i] \mid i \in \mathcal{I}_p^{(L)}\}. \end{aligned} \quad (6)$$

where  $\mathbf{K}_v^{(L)}$  and  $\mathbf{V}_v^{(L)}$  denote the set of visual tokens in the KV cache of layer  $L$ .

In the next decoding step, the model may need to *refocus* on the tokens that were previously pruned. If they are discarded directly, the model cannot retrieve their KV cache entries. We also consider that tokens of interest tend to remain consistent across successive iterations. To reduce large-scale indexing requirements, we employ cosine similarity to measure the attention distribution across different decoding iterations. The KV Cache is updated only at iteration  $N$ , where a low similarity is observed. We also introduce a *dynamic pruning cache* (DP cache) to store pruned tokens, which can be denoted as

$$\begin{aligned} K_{DP}^{(L)} &= \{\mathbf{K}^{(L)}[i] \mid i \in \mathcal{J}^{(L)}\}, \\ \mathbf{V}_{DP}^{(L)} &= \{\mathbf{V}^{(L)}[i] \mid i \in \mathcal{J}^{(L)}\}, \end{aligned} \quad (7)$$

where  $\mathcal{J}^{(L)} = \{i \mid i \notin \mathcal{I}_p^{(L)}\}$ . In the iteration  $N$ , the cross-attention matrix is recalculated at layer  $L$ , dynamically adding tokens from the DP cache with increased attention scores into the calculation and storing them in the KV cache. Simultaneously, tokens whose attention scores have decreased since the previous stage and no longer fall within the top  $p\%$  are indexed. These tokens are removed from the KV cache and stored in the DP cache. This process repeats at each decoding stage, with the KV cache and DP cache dynamically updated for optimal token compression.

## 4. Experimental Results

### 4.1. Evaluation Setups and Implementation Details

**Benchmarks.** We evaluate the performance of VLLMs using established video-to-text benchmarks. ActivityNet-QA [50] comprises human-annotated, action-related question-answer pairs derived from the ActivityNet dataset. For evaluation purposes, single-word responses are generated. We assess the model’s accuracy and response quality (scored on a scale from 0 to 5) using GPT-4o-mini [33]. We also utilize the PerceptionTest [34] to assess perception capabilities, alongside VideoMME [8] and NeXTQA [46], which encompass various video domains and durations. VideoDetailCaption [31] assesses the model’s detailed understanding of video content, with scoring also conducted via GPT-4o-mini [33]. Additionally, we further evaluate the model on the MVbench [23] dataset, which includes 20 complex tasks requiring comprehensive video understanding beyond single-frame analysis. Each task contains 200 test samples in a multiple-choice VideoQA format. These samples require the model to choose the correct answer from multiple provided options. For MVBench, we conducted multiple experiments and computed the average results.

**Comparison Methods.** We compare our method with two latest training-free visual token compression methods: *LLaVA-PruMerge* [39] leverages the sparsity of attention scores in CLIP [35] to identify key tokens; *FastV* [3] adopts the attention score between the predicted tokens and the visual tokens during the prefilling stage to identify key tokens. Notably, FastV offers two versions, with and without KV cache. In practice, we found that FastV without KV cache causes a sharp increase in inference time for video input, thereby we mainly compare to the version *with* KV cache. Of note, the above two methods rely on *one-shot* token pruning, while our approach introduces *dynamic* visual token pruning for the first time. We use the official codes<sup>1</sup> of these methods for evaluation under identical hardware conditions.

**Implementation Details.** We implement the proposed DyCoke on the LLaVA-OneVision-0.5B, LLaVA-OneVision-7B, and LLaVA-OneVision-72B models using NVIDIA RTX 4090 (24GB), A6000 (48GB), and A100 (80GB) GPUs, respectively. And we use the PyTorch framework. To set the pruning ratios for all methods, we use total calculated FLOPs to ensure fair comparison. The attention computation layer for FastV is set to layer 5. For video input, we follow the official requirements of the LLaVA-OneVision model, with a default of 32 video input frames and  $N_v = 196$ , except for experiments with specific instructions. In the comparison experiments,  $L$  is set to 3 and  $P$  to 0.7, with the first-stage pruning rate  $K$  serving as the primary experimental variable.

---

1. <https://github.com/pkunlp-icler/FastV>, <https://github.com/42Shawn/LLaVA-PruMerge>

Method	Pruning Settings			ActNet-QA		NextQA		PercepTest		VideoDC		Videomme	
	Retained Ratio (Final)	FLOPs (T)	FLOPs Ratio	Acc.	Sco.	mc	val	test	wo	w-subs			
LLaVA-OV-0.5B													
Full Tokens	100%	3.4	100%	47.93	2.66	57.2	49.1	2.86	44.1	43.5			
FastV	35%	1.4	41%	46.74	<u>2.58</u>	56.5	<u>49.1</u>	2.36	42.0	41.4			
PruMerge	55%	1.5	44%	41.68	2.35	54.0	47.5	2.10	38.8	39.3			
Ours ( $K=0.3, L=3, P=0.7$ )	23.25%	2.4	70%	<b>47.90</b>	<u>2.65</u>	<u>57.2</u>	48.9	<u>2.63</u>	<b>45.4</b>	<b>43.8</b>			
Ours ( $K=0.5, L=3, P=0.7$ )	18.75%	1.8	53%	<u>47.80</u>	<b>2.65</b>	<u>57.2</u>	<b>49.5</b>	<u>2.62</u>	45.1	<u>43.4</u>			
Ours ( $K=0.7, L=3, P=0.7$ )	14.25%	1.2	35%	47.70	<b>2.65</b>	<b>57.7</b>	<b>49.5</b>	2.56	<u>45.2</u>	43.3			
LLaVA-OV-7B													
Full Tokens	100%	41.4	100%	51.93	2.86	79.4	57.1	3.30	58.5	61.3			
FastV	35%	17.9	43%	50.93	2.80	78.2	56.7	3.09	57.3	60.5			
PruMerge	55%	21.1	51%	50.45	2.78	76.0	54.3	2.88	52.9	57.0			
Ours ( $K=0.3, L=3, P=0.7$ )	23.25%	30.8	75%	<u>51.80</u>	<u>2.85</u>	<b>79.1</b>	<u>57.2</u>	3.19	<u>58.8</u>	<u>61.0</u>			
Ours ( $K=0.5, L=3, P=0.7$ )	18.75%	24.1	59%	<b>52.08</b>	<b>2.88</b>	<u>78.5</u>	<b>57.6</b>	<u>3.29</u>	<b>59.5</b>	<b>61.4</b>			
Ours ( $K=0.7, L=3, P=0.7$ )	14.25%	17.9	43%	<u>51.80</u>	<u>2.85</u>	78.2	<b>57.6</b>	<u>3.20</u>	58.3	60.7			
LLaVA-OV-72B													
Full Tokens	100%	436.1	100%	52.96	2.92	80.2	66.9	3.34	66.2	69.5			
FastV	45%	202.4	46%	<u>52.63</u>	2.82	77.2	58.5	3.01	62.1	66.7			
PruMerge	55%	229.8	53%	50.91	2.80	75.2	55.6	2.81	61.9	63.8			
Ours ( $K=0.5, L=3, P=0.7$ )	18.75%	262.5	60%	<u>52.81</u>	<u>2.92</u>	<b>79.1</b>	<u>60.2</u>	<u>3.35</u>	<b>66.3</b>	<b>69.7</b>			
Ours ( $K=0.7, L=3, P=0.7$ )	14.25%	195.1	44%	52.38	<u>2.88</u>	<u>78.8</u>	<u>59.6</u>	<u>3.27</u>	64.9	68.7			

Table 1. Comparison of different methods on video QA and description benchmarks. For all the values, the higher is better. In this context,  $K$  represents the pruning rate in the first stage of DyCoke;  $L$  denotes the attention evaluation layer; and  $P$  indicates the pruning rate in the second stage of our method. The **best** result among token pruning methods of each metric is in bold, second best underlined.

Method	FR	AS	AP	AA	FA	UA	OE	OI	OS	MD	AL	ST	AC	MC	MA	SC	FP	CO	EN	ER	CI	Avg.	
LLaVA-OV-7B																							
Full Tokens	100%	70.7	71.5	84.6	45.0	79.0	57.1	81.0	38.0	24.5	46.0	91.5	44.0	46.5	72.0	51.0	51.0	68.0	36.0	71.5	51.0	58.0	
PruMerge	51%	63.3	63.0	84.6	39.5	76.5	51.0	61.0	36.0	<b>32.5</b>	<b>49.0</b>	88.0	41.0	39.5	62.5	46.5	45.5	54.0	<u>35.5</u>	67.5	38.5	52.6	
FastV	43%	<b>73.5</b>	<u>73.5</u>	76.9	44.0	76.5	<u>56.1</u>	76.0	<b>42.5</b>	19.5	40.0	<u>92.0</u>	43.5	<u>43.0</u>	68.5	48.0	<b>50.0</b>	<u>67.0</u>	33.5	69.0	43.5	56.1	
Ours ( $K=0.5$ )	59%	69.7	<u>73.5</u>	84.6	<b>47.5</b>	<u>79.0</u>	<b>57.1</b>	77.5	39.0	22.5	<u>47.0</u>	<u>93.0</u>	44.0	<b>45.0</b>	<b>74.0</b>	<u>50.5</u>	<b>50.0</b>	66.0	<u>36.5</u>	<u>71.5</u>	<b>51.0</b>	<b>58.0</b>	
Ours ( $K=0.7$ )	43%	67.6	<b>74.0</b>	<u>92.3</u>	<u>47.0</u>	<b>80.0</b>	55.6	<u>78.0</u>	<u>39.0</u>	<u>24.0</u>	45.0	<u>93.0</u>	<b>45.0</b>	<u>45.5</u>	<u>71.0</u>	50.0	49.0	<b>67.5</b>	34.5	70.0	49.0	<u>57.5</u>	
LLaVA-OV-0.5B																							
Full Tokens	100%	53.7	59.5	30.8	37.5	60.5	46.5	68.5	35.0	21.0	32.0	87.0	44.5	29.5	54.5	37.0	45.0	46.0	29.0	70.0	39.5	47.1	
PruMerge	46%	37.8	48.5	23.1	30.5	53.5	45.5	45.5	31.0	20.0	<b>39.5</b>	83.0	38.5	<b>30.5</b>	47.0	35.5	<u>43.5</u>	38.0	28.0	69.0	<b>44.0</b>	42.5	
FastV	46%	52.1	<b>60.5</b>	<b>38.5</b>	<u>35.0</u>	<b>62.0</b>	<u>46.0</u>	63.0	<u>35.5</u>	<u>23.5</u>	28.5	84.5	42.0	<u>29.0</u>	50.5	37.0	41.0	<b>47.5</b>	<u>28.5</u>	67.5	<b>44.0</b>	46.1	
Ours ( $K=0.5$ )	60%	<u>53.7</u>	<u>59.5</u>	<u>30.8</u>	<u>39.0</u>	60.0	<u>46.0</u>	<b>68.5</b>	32.5	21.0	<u>30.0</u>	87.0	<u>42.0</u>	<u>29.0</u>	<b>55.5</b>	<u>39.0</u>	<b>44.0</b>	<u>46.0</u>	<u>28.5</u>	<b>70.0</b>	<u>42.5</u>	47.0	
Ours ( $K=0.7$ )	44%	<b>54.8</b>	<u>59.5</u>	<u>30.8</u>	<u>39.0</u>	61.0	<b>47.5</b>	68.0	34.0	21.5	<u>30.0</u>	<b>88.0</b>	<u>44.5</u>	29.0	<u>54.5</u>	<u>38.0</u>	41.0	45.5	<b>29.5</b>	69.0	41.5	<b>47.1</b>	

Table 2. Comparison of different methods on the MVBench [23] dataset (number of input frames: 16). FR refers to FLOPs ratio; FR = 100% indicates no tokens are removed, which is the original baseline. From AS to CI are the different sub-tasks in MVBench. For all the values, the higher is better. The **best** result among token pruning methods of each metric is in bold, second best underlined.

For benchmarks such as PerceptionTest [34], VideoMME [8], NeXTQA, VideoDetailCaption [31], and ActivityNet-QA [50], we use the LMMs-Eval [17, 53] for evaluation, while MVBench [23] is evaluated using the official code.

## 4.2. Main Results

**Video QA.** (1) As shown in Tab. 1, our method *significantly* outperforms the counterpart methods FastV and PruMerge at similar or lower computational cost, on different benchmarks on 0.5B, 7B, and 72B VLLMs.

(2) Notably, with proper pruning, we achieve superior average performance compared to the original model on the PercepTest and Videomme benchmarks, suggesting that effectively reducing temporal redundancy can enhance the model’s ability to understand and reason about video information. Specifically, LLaVA-PruMerge filters out most tokens not relevant to the task, while FastV suffers from performance degradation due to the inherent limitations of

one-time pruning, which prevents fine-grained, accurate filtering of important tokens.

(3) We further conduct evaluations on the more challenging Multi-Choice VideoQA task in MVBench. Results in Tab. 2 show that DyCoke also achieves the *best* quantitative results. The merits of our method can be further confirmed in Figure 4, where DyCoke is shown to improve the model’s ability to focus on finer details by reducing the token redundancy. These results collectively demonstrate DyCoke’s effectiveness in accurately retaining temporal information and can accurately prune and merge non-essential tokens while preserving model performance.

**Video Description.** The video description task requires the model to summarize and describe the video by understanding detailed actions and events, which involves generating long paragraphs of text. A representative benchmark for this task is VideoDetailCaption (VideoDC). As shown in Tab. 2, DyCoke achieves minimal performance degradation

## Challenging Video Understanding



User                    *What color is the object that is stationary? | What direction is the yellow sphere moving in?*

LLaVA-OV            The object that is stationary is **gray**. ☺ | From the left to the right side of the frame. ☺

LLaVA-OV w/ FastV    The object that is stationary is blue. ☺ | Towards the left side of the frame. ☺

LLaVA-OV w/ DyCoke: The object that is stationary is **gray**. ☺ | From the left to the right side of the frame. ☺



User                    *The person uses multiple similar objects to play an occlusion game. Where is the hidden object at the end of the game from the person's point of view?*

LLaVA-OV            The hidden object is under the middle purple cup. ☺

LLaVA-OV w/ FastV    The hidden object is under the middle purple cup. ☺

LLaVA-OV w/ DyCoke: The hidden object is under the **third purple cup from the left**. ☺

**Figure 4. Showcases of our DyCoke compared to FastV with LLaVA-OV 7B on MVBench.** The first row shows that after token compression by FastV, the model generates the *wrong* answer while our method still retains the correct answer. The second row demonstrates a case that our token compression method can calibrate the mistake from attending full tokens, suggesting that retaining *less but key* information can enhance the model’s capability for correct video understanding.

Method	Total Latency ↓	GPU Mem. ↓	Accuracy ↑	Latency per Example ↓
LLaVA-OV-7B				
Full Tokens	1:19:27	27G	57.58 ( $\pm 0.11$ )	1.19s (1.00×)
PruMerge	2:02:17	20G	52.59 ( $\pm 0.08$ )	1.83s (0.65×)
FastV	1:03:50	24G	56.05 ( $\pm 0.13$ )	0.96s (1.24×)
Ours ( $K = 0.5$ )	59:02	21G	<b>57.88</b> ( $\pm 0.10$ )	0.88s (1.35×)
Ours ( $K = 0.7$ )	<b>57:13</b>	<b>19G</b>	57.45 ( $\pm 0.18$ )	<b>0.85s (1.40×)</b>
LLaVA-OV-0.5B				
Full Tokens	34:05	19G	47.09 ( $\pm 0.09$ )	0.56s (1.00×)
PruMerge	1:35:08	<b>8.5G</b>	42.53 ( $\pm 0.30$ )	1.42s (0.41×)
FastV	32:30	10G	46.14 ( $\pm 0.03$ )	0.49s (1.14×)
Ours ( $K = 0.5$ )	32:17	8.9G	46.76 ( $\pm 0.13$ )	0.48s (1.16×)
Ours ( $K = 0.7$ )	<b>31:45</b>	<b>7.4G</b>	<b>47.09</b> ( $\pm 0.11$ )	<b>0.47s (1.19×)</b>

(a) The number of video input frames is 16

Method	Total Latency ↓	GPU Mem. ↓	Accuracy ↑	Latency per Example ↓
LLaVA-OV-7B				
Full Tokens	2:33:30	34G	58.81 ( $\pm 0.14$ )	2.30s (1.00×)
PruMerge	3:48:59	28G	53.93 ( $\pm 0.06$ )	3.43s (0.64×)
FastV	1:55:20	30G	58.36 ( $\pm 0.09$ )	1.73s (1.32×)
Ours ( $K = 0.5$ )	1:53:17	28G	59.09 ( $\pm 0.23$ )	1.69s (1.36×)
Ours ( $K = 0.7$ )	<b>1:39:49</b>	<b>24G</b>	<b>59.56</b> ( $\pm 0.19$ )	<b>1.49s (1.54×)</b>
LLaVA-OV-0.5B				
Full Tokens	1:19:06	21G	48.23 ( $\pm 0.15$ )	1.18s (1.00×)
PruMerge	3:06:06	13G	42.55 ( $\pm 0.18$ )	2.79s (0.42×)
FastV	1:13:25	15G	47.04 ( $\pm 0.05$ )	1.10s (1.07×)
Ours ( $K = 0.5$ )	1:10:32	12G	<b>48.13</b> ( $\pm 0.09$ )	1.06s (1.11×)
Ours ( $K = 0.7$ )	<b>1:02:15</b>	<b>10G</b>	47.77 ( $\pm 0.10$ )	<b>0.93s (1.27×)</b>

(b) The number of video input frames is 32

**Table 3. Actual inference efficiency comparison on MVBench.** MVBench dataset is used here to eliminate the impact of output sequence length on decoding time, where the model only outputs *one* token. Experiments with 7B and 0.5B models are conducted on a single A6000 GPU and a single 4090 GPU, respectively.

relative to other limiting methods due to the implementation of dynamic token pruning during the decoding phase. FastV incorrectly prunes important tokens overlooked by the LLM during the prefilling phase of pruning, resulting in severe

Methods	#Params	Decoding Latency	VideoDC Acc.
Full Tokens	7B	42 ms/token	3.30
Ours ( $K = 0.5$ )	7B	35 ms/token	3.29
Ours ( $K = 0.7$ )	7B	31 ms/token	3.20

**Table 4. Actual inference efficiency evaluation on VideoDC.** VideoDC is a video description benchmark (32 input frames are used here). Unlike the MVBench results in Tab. 3, here the model outputs *multiple* tokens.

performance degradation.

### 4.3. Efficiency Analysis

**Latency and Memory Comparison.** We first compare the speed and memory consumption of model inference with different numbers of sampling frames (16 and 32 frames). To ensure result robustness, we test on the MVBench dataset [23] to minimize the influence of output length. Results in Tab. 3 show that the model with compressed tokens by our method runs significantly faster than its full-token counterpart, with a speedup of  $1.4\times$  on 7B models. The advantage is even more pronounced for longer visual sequences, with a speedup of  $1.54\times$ . The speedup comes along with lower memory consumption than other baseline methods.

For long text generation, we evaluate the time required to predict each new token during the decoding stage on the VideoDC [31]. As shown in Tab. 4, our method preserves model performance comparable to the original model while significantly reducing latency compared to full tokens.

Method	Pruning Settings				ActNet-QA		NextQA	PercepTest	VideoDC	Videommme	
	K	L	P	Retained Ratio	Acc.	Sco.	mc	val	test	wo	w-subs
LLaVA-OV-7B											
Full Tokens	-	-	-	100%	51.93	2.86	79.4	57.1	3.30	58.5	61.3
w/o DP	0.7	3	0.7	14.25%	51.06	2.82	77.2	56.6	3.01	58.1	60.2
Random Pruning	0.7	3	0.7	14.25%	50.90	2.79	77.9	56.4	2.98	55.8	59.3
DyCoke	0.7	3	0.7	14.25%	51.80	2.85	78.2	57.6	3.20	58.3	60.7
DyCoke	0.7	10	0.7	14.25%	51.81	2.85	78.2	57.5	3.20	58.4	60.7
DyCoke	0.7	3	0.9	4.75%	51.48	2.83	78.2	57.5	2.86	58.3	60.7
DyCoke	0.9	0	0.9	3.25%	40.21	2.24	79.1	57.4	2.76	57.8	60.1
LLaVA-OV-0.5B											
Full Tokens	-	-	-	100%	47.93	2.66	57.2	49.1	2.86	44.1	43.5
w/o DP	0.7	3	0.7	14.25%	46.44	2.53	56.0	48.6	2.36	42.0	41.4
Random Pruning	0.7	3	0.7	14.25%	41.68	2.35	54.0	47.5	2.10	38.8	39.3
DyCoke	0.7	3	0.7	14.25%	47.70	2.65	57.7	49.5	2.56	45.2	43.3
DyCoke	0.9	0	0.9	3.25%	47.10	2.61	57.5	49.1	1.75	43.5	42.8

**Table 5. Ablation study of our DyCoke method.**  $K$ ,  $L$ ,  $P$  are three hyper-parameters of our method, through which we can control the retained token ratio (*Retained Ratio*):  $K$  represents the pruning rate in the first stage of DyCoke;  $L$  denotes the attention evaluation layer; and  $P$  indicates the pruning rate in the second stage of DyCoke. The key innovation of our method is dynamic pruning (*DP*) of tokens. The rows of *w/o DP* and *Random Pruning* are to see the effect of not using dynamic pruning or using a naive alternative.

Method	FLOPs	Input Frames	Accuracy			
			S Video	M Video	L Video	Avg.
Full Token	18.99T	16	67.9	52.8	47.9	56.2
DyCoke	17.91T	32	71.0	55.6	48.3	58.3
Full Token	41.40T	32	71.0	55.0	49.7	58.5

**Table 6. Cost-effectiveness analysis.** By using our DyCoke, the model can process *more* video frames under the same computational budget, leading to improved video understanding performance. The performance is accessible to the full-token model yet at a dramatically increased cost (from 17.91T FLOPs to 41.40T).

**Cost-Effectiveness.** With visual tokens reduced, our method allows for longer video frames as input while maintaining the same computational budget. Experiments on the VideoMME benchmark, which includes short, medium-length, and long videos, show that our method, as illustrated in Tab. 6, notably improves performance on short and medium-length videos. Long videos, due to accumulated content, convey meaning with fewer frames, whereas short videos depend on dense frame sequences. Thus, DyCoke achieves a superior cost-performance balance.

#### 4.4. Ablation Study

**Token Selection Strategy.** In our proposed TTM (token temporal merging) module, the token similarity between adjacent needles and preceding frames is used to evaluate temporal redundancy for pruning and merging. In Tab. 5 (marked in blue background), the effect of random token selection for pruning in TTM on model performance is analyzed. The effectiveness of the visibility model in understanding video content drops substantially, highlighting the superiority of our token filtering strategy. Furthermore, we investigate the effects of over-pruning in the TTM module (highlighted in gray) and observe a sharp performance drop, further demonstrating the TTM module’s effectiveness in

mitigating the temporal redundancy of visual tokens.

**Dynamic Pruning.** As mentioned in Sec. 3.2, the attention of each visual token in VLLMs varies at each decoding iteration stage, inspiring us to design a dynamic token pruning strategy (DP). As shown in Tab. 5, when dynamic pruning is replaced by one-shot pruning, model performance declines on various tasks, especially on the VideoDC benchmark (marked in green). This demonstrates the effectiveness of dynamic pruning in preventing the undesired pruning of important tokens. In addition, we also explore the impact of attention evaluation layer  $L$  on the overall performance. However, we observe that when  $L > 0$ , dynamic pruning does not significantly affect model performance, indicating the stability of dynamic token pruning.

## 5. Conclusion

This paper presents *DyCoke*, a new *training-free* method to dynamically reduce the visual tokens for faster video large language models (VLLMs). We develop a two-stage token compression strategy that leverages the temporal and spatial redundancy in video information: In the first stage, highly similar temporal tokens among frames are merged; the second stage further reduces the visual tokens used for attention computation during the decoding stage. To the best of our knowledge, this is the first dynamic token pruning method specifically tailored for VLLMs. Extensive benchmark and analysis results on a wide range of video QA and reasoning tasks with three VLLMs (0.5B, 7B, 70B parameters) show our method consistently surpasses the prior SoTA counterparts. Using our method on VLLMs, we can achieve up to  $1.4\times$  memory reduction, and  $1.5\times$  speedup, with performance still improved.

## References

- [1] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. *arXiv preprint arXiv:2210.09461*, 2022. 3
- [2] 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. *arXiv preprint arXiv:2403.06764*, 2024. 3
- [3] 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. *arXiv preprint arXiv:2403.06764*, 2024. 1, 2, 3, 5
- [4] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024. 1, 2
- [5] 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. 1, 3
- [6] Xiangxiang Chu, Limeng Qiao, Xinyu Zhang, Shuang Xu, Fei Wei, Yang Yang, Xiaofei Sun, Yiming Hu, Xinyang Lin, Bo Zhang, et al. Mobilevlm v2: Faster and stronger baseline for vision language model. *arXiv preprint arXiv:2402.03766*, 2024. 1, 3
- [7] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35:30318–30332, 2022. 3
- [8] Chaoyou Fu, Yuhan Dai, Yondong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. *arXiv preprint arXiv:2405.21075*, 2024. 5, 6
- [9] Qichen Fu, Minsik Cho, Thomas Merth, Sachin Mehta, Mohammad Rastegari, and Mahyar Najibi. Lazyllm: Dynamic token pruning for efficient long context llm inference. *arXiv preprint arXiv:2407.14057*, 2024. 3
- [10] Zhangwei Gao, Zhe Chen, Erfei Cui, Yiming Ren, Weiyun Wang, Jinguo Zhu, Hao Tian, Shenglong Ye, Junjun He, Xizhou Zhu, et al. Mini-internvl: a flexible-transfer pocket multi-modal model with 5% parameters and 90% performance. *Visual Intelligence*, 2(1):1–17, 2024. 3
- [11] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 2
- [12] De-An Huang, Shijia Liao, Subhashree Radhakrishnan, Hongxu Yin, Pavlo Molchanov, Zhiding Yu, and Jan Kautz. Lita: Language instructed temporal-localization assistant. In *ECCV*, 2024. 3
- [13] Wei Huang, Xingyu Zheng, Xudong Ma, Haotong Qin, Chengtao Lv, Hong Chen, Jie Luo, Xiaojuan Qi, Xianglong Liu, and Michele Magno. An empirical study of llama3 quantization: From llms to mllms. *Visual Intelligence*, 2(1):36, 2024. 3
- [14] Mojan Javaheripi, Sébastien Bubeck, Marah Abdin, Jyoti Aneja, Sébastien Bubeck, Caio César Teodoro Mendes, Weizhu Chen, Allie Del Giorno, Ronen Eldan, Sivakanth Gopi, et al. Phi-2: The surprising power of small language models. *Microsoft Research Blog*, 1:3, 2023. 3
- [15] Yao Jiang, Xinyu Yan, Ge-Peng Ji, Keren Fu, Meijun Sun, Huan Xiong, Deng-Ping Fan, and Fahad Shahbaz Khan. Effectiveness assessment of recent large vision-language models. *Visual Intelligence*, 2(1):17, 2024. 3
- [16] Peng Jin, Ryuichi Takanobu, Wancai Zhang, Xiaochun Cao, and Li Yuan. Chat-univ: Unified visual representation empowers large language models with image and video understanding. In *CVPR*, 2024. 2, 3
- [17] Bo Li, Peiyuan Zhang, Kaichen Zhang, Fanyi Pu, Xinrun Du, Yuhao Dong, Haotian Liu, Yuanhan Zhang, Ge Zhang, Chunyuan Li, and Ziwei Liu. Lmms-eval: Accelerating the development of large multimodal models, 2024. 6
- [18] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024. 1, 2, 3, 11
- [19] Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*, 2024. 2, 3
- [20] 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 *ICML*, 2023. 3
- [21] Jun Li, Zedong Zhang, and Jian Yang. Tp2o: Creative text pair-to-object generation using balance swap-sampling. In *ECCV*, 2024. 2
- [22] KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhui Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. 1, 3
- [23] Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In *CVPR*, 2024. 1, 5, 6, 7, 11
- [24] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models. In *ECCV*. Springer, 2024. 3
- [25] Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023. 1, 2
- [26] Bin Lin, Zhenyu Tang, Yang Ye, Jiaxi Cui, Bin Zhu, Peng Jin, Jinfang Huang, Junwu Zhang, Yatian Pang, Munan Ning, et al. Moe-llava: Mixture of experts for large vision-language models. *arXiv preprint arXiv:2401.15947*, 2024. 3
- [27] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. Vila: On pre-training for visual language models. In *CVPR*, 2024. 3

- [28] Haogeng Liu, Qihang Fan, Tingkai Liu, Linjie Yang, Yunzhe Tao, Huabo Huang, Ran He, and Hongxia Yang. Videoteller: Enhancing cross-modal generation with fusion and decoupling. *arXiv preprint arXiv:2310.04991*, 2023. 2
- [29] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *CVPR*, 2024. 3
- [30] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *NeurIPS*, 2024. 3
- [31] LMMs-Lab. Video detail caption, 2024. Accessed: 2024-11-5, 6, 7
- [32] Ruipu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Da Li, Pengcheng Lu, Tao Wang, Linmei Hu, Minghui Qiu, and Zhongyu Wei. Valley: Video assistant with large language model enhanced ability. *arXiv preprint arXiv:2306.07207*, 2023. 2
- [33] OpenAI. Gpt-4 model, 2023. Accessed: 2023-11-08. 5
- [34] Viorica Patraucean, Lucas Smaira, Ankush Gupta, Adria Recasens, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Mateusz Malinowski, Yi Yang, Carl Doersch, et al. Perception test: A diagnostic benchmark for multimodal video models. *Advances in Neural Information Processing Systems*, 36, 2024. 5, 6
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 5
- [36] Shuhuai Ren, Sishuo Chen, Shicheng Li, Xu Sun, and Lu Hou. Testa: Temporal-spatial token aggregation for long-form video-language understanding. *arXiv preprint arXiv:2310.19060*, 2023. 2, 3
- [37] Michael S Ryoo, Honglu Zhou, Shrikant Kendre, Can Qin, Le Xue, Manli Shu, Silvio Savarese, Ran Xu, Caiming Xiong, and Juan Carlos Niebles. xgen-mm-vid (blip-3-video): You only need 32 tokens to represent a video even in vlms. *arXiv preprint arXiv:2410.16267*, 2024. 3
- [38] Yuzhang Shang, Zhihang Yuan, Qiang Wu, and Zhen Dong. Pb-llm: Partially binarized large language models. *arXiv preprint arXiv:2310.00034*, 2023. 3
- [39] Yuzhang Shang, Mu Cai, Bingxin Xu, Yong Jae Lee, and Yan Yan. Llava-prumerge: Adaptive token reduction for efficient large multimodal models. *arXiv preprint arXiv:2403.15388*, 2024. 1, 2, 3, 5
- [40] Leqi Shen, Tianxiang Hao, Sicheng Zhao, Yifeng Zhang, Pengzhang Liu, Yongjun Bao, and Guiguang Ding. Tempme: Video temporal token merging for efficient text-video retrieval. *arXiv preprint arXiv:2409.01156*, 2024. 3
- [41] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022. 2
- [42] Zhongwei Wan, Ziang Wu, Che Liu, Jinfa Huang, Zhihong Zhu, Peng Jin, Longyue Wang, and Li Yuan. Look-m: Look-once optimization in kv cache for efficient multimodal long-context inference. *arXiv preprint arXiv:2406.18139*, 2024. 3
- [43] Jiawei Wang, Liping Yuan, and Yuchen Zhang. Tarsier: Recipes for training and evaluating large video description models. *arXiv preprint arXiv:2407.00634*, 2024. 1, 2, 3
- [44] Haoran Wei, Lingyu Kong, Jinyue Chen, Liang Zhao, Zheng Ge, En Yu, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. Small language model meets with reinforced vision vocabulary. *arXiv preprint arXiv:2401.12503*, 2024. 3
- [45] Jiangwei Weng, Zhiqiang Yan, Ying Tai, Jianjun Qian, Jian Yang, and Jun Li. Mamballie: Implicit retinex-aware low light enhancement with global-then-local state space. In *NeurIPS*, 2024. 2
- [46] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *CVPR*, 2021. 5
- [47] Yisong Xiao, Aishan Liu, Tianyuan Zhang, Haotong Qin, Jinyang Guo, and Xianglong Liu. Robustmq: benchmarking robustness of quantized models. *Visual Intelligence*, 1(1):30, 2023. 3
- [48] Zeren Xiong, Zedong Zhang, Zikun Chen, Shuo Chen, Xiang Li, Gan Sun, Jian Yang, and Jun Li. Novel object synthesis via adaptive text-image harmony. In *NeurIPS*, 2024. 2
- [49] Lin Xu, Yilin Zhao, Daquan Zhou, Zhijie Lin, See Kiong Ng, and Jiashi Feng. Pllava: Parameter-free llava extension from images to videos for video dense captioning. *arXiv preprint arXiv:2404.16994*, 2024. 1
- [50] Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueling Zhuang, and Dacheng Tao. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *AAAI*, 2019. 5, 6
- [51] Zhengqing Yuan, Zhaoxu Li, Weiran Huang, Yanfang Ye, and Lichao Sun. Tinygpt-v: Efficient multimodal large language model via small backbones. *arXiv preprint arXiv:2312.16862*, 2023. 1, 3
- [52] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023. 1
- [53] Kaichen Zhang, Bo Li, Peiyuan Zhang, Fanyi Pu, Joshua Adrian Cahyono, Kairui Hu, Shuai Liu, Yuanhan Zhang, Jingkang Yang, Chunyuan Li, and Ziwei Liu. Lmms-eval: Reality check on the evaluation of large multimodal models, 2024. 6
- [54] Baichuan Zhou, Ying Hu, Xi Weng, Junlong Jia, Jie Luo, Xien Liu, Ji Wu, and Lei Huang. Tinyllava: A framework of small-scale large multimodal models. *arXiv preprint arXiv:2402.14289*, 2024. 1, 3
- [55] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. 3
- [56] Yichen Zhu, Minjie Zhu, Ning Liu, Zhiyuan Xu, and Yaxin Peng. Llava-phi: Efficient multi-modal assistant with small language model. In *Proceedings of the 1st International Workshop on Efficient Multimedia Computing under Limited*, pages 18–22, 2024. 3
- [57] Jiedong Zhuang, Lu Lu, Ming Dai, Rui Hu, Jian Chen, Qiang Liu, and Haoji Hu. St<sup>3</sup>: Accelerating multimodal large language model by spatial-temporal visual token trimming. In *AAAI*, 2025. 3