

# Contribution-aware Token Compression for Efficient Video Understanding via Reinforcement Learning

Yinchao Ma, Qiang Zhou, Zhibin Wang, Xianing Chen, Hanqing Yang, Jun Song\*, Bo Zheng

Taobao & Tmall Group of Alibaba  
{mayinchao.myc, jianchong.zq, jsong.sj}@taobao.com

## Abstract

Video large language models have demonstrated remarkable capabilities in video understanding tasks. However, the redundancy of video tokens introduces significant computational overhead during inference, limiting their practical deployment. Many compression algorithms are proposed to prioritize retaining features with the highest attention scores to minimize perturbations in attention computations. However, the correlation between attention scores and their actual contribution to correct answers remains ambiguous. To address the above limitation, we propose a novel Contribution-aware token Compression algorithm for VIDEO understanding (CaCoVID) that explicitly optimizes the token selection policy based on the contribution of tokens to correct predictions. First, we introduce a reinforcement learning-based framework that optimizes a policy network to select video token combinations with the greatest contribution to correct predictions. This paradigm shifts the focus from passive token preservation to active discovery of optimal compressed token combinations. Secondly, we propose a combinatorial policy optimization algorithm with online combination space sampling, which dramatically reduces the exploration space for video token combinations and accelerates the convergence speed of policy optimization. Extensive experiments on diverse video understanding benchmarks demonstrate the effectiveness of CaCoVID. Codes will be released.

## Introduction

Video large language models (Bai et al. 2025; Li et al. 2025; Lin et al. 2024) have achieved remarkable advancements in video understanding tasks. However, the integration of video data into LLMs introduces substantial computational challenges, primarily due to the numerous processing demands caused by densely encoded video tokens and the inherent quadratic complexity of attention mechanisms. Thus, compressing video tokens with minimal performance loss has attracted increasing attention for video understanding.

Numerous studies (Chen et al. 2024; Yang et al. 2025b) have explored token compression for multimodal large language models (MLLMs). The design motivations behind these methods can generally be categorized into two

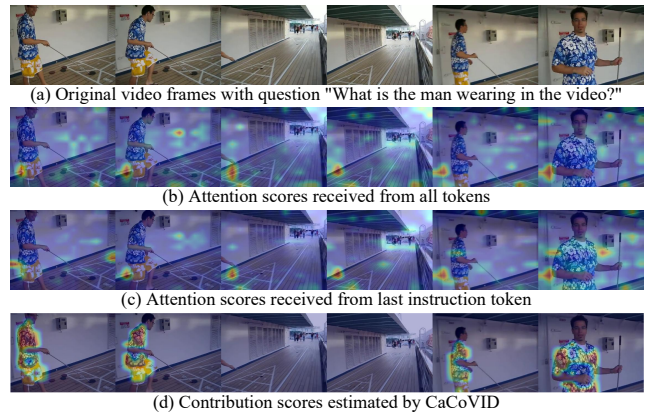


Figure 1: As observed, the attention scores of LLMs demonstrate ambiguity in their correlation with token contributions to correct question answering. Higher attention scores are not allocated to critical tokens such as “the clothing of the man”, which may stem from the visual attention sink phenomenon (Kang et al. 2025). In contrast, our contribution scores learned through LLM prediction feedback effectively focus on the regions most critical to answering the question correctly - specifically, the “the clothing of the man” area.

paradigms: content-based token compression and model-based token compression. **Content-based token compression** aims to preserve content diversity or spatio-temporal structure as design objectives, employing handcrafted content metrics to guide token compression. Typically, DivPrune (Alvar et al. 2025) prunes features by maximizing the diversity of visual tokens based on the solution of max-min diversity problem (MMDP). TopV (Yang et al. 2025a) comprehensively considers feature similarity, spatial relationships, and central distance to evaluate token importance for pruning. Although these approaches preserve rich visual contents, their compression strategy is query-agnostic and may prune tokens critical to the query. **Model-based token compression** aims to minimize perturbations in LLM inference as design objectives, typically pruning tokens with low attention scores to compress the token sequence. Typically, FastV (Chen et al. 2024) proposes to filter out unimportant tokens by using the average attention score that each

\*Corresponding Author

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

token receives from all other tokens as a metric. Pyramid-Drop (Xing et al. 2024) designs a multi-stage pruning strategy to progressively prune visual tokens at different layers based on the attention scores between the last-instruction token and visual tokens. However, the correlation between attention scores and their actual contribution to correct answers remains ambiguous, which may lead to suboptimal token compression, as shown in Figure 1(b-c).

To address the above limitations, we shift the focus from passive token preservation to active discovery of optimal compressed token combinations for correct prediction. To achieve this, there are two critical issues that need to be considered for active exploration of optimal token combinations. 1) **How to empower active token selection.** In existing token compression algorithms (Chen et al. 2024; Xing et al. 2024), LLMs fail to actively participate in the token selection process, with tokens instead being passively selected based on handcrafted metrics or pretrained attention scores. Therefore, the selected tokens may be suboptimal for correct predictions. To address this, we seek to train a small compression policy network that actively learns to select optimal token combinations with feedback of LLM predictions. 2) **How to explore optimal token combinations.** There currently exists no large-scale video training data annotated with key tokens for token selection optimization. Although reinforcement learning algorithms (Schulman et al. 2017; Shao et al. 2024) provide solutions to optimize networks via sampling-based exploration with environmental rewards in unsupervised settings, the combinatorial exploration space of  $n$  video tokens reaches  $2^n$  where  $n$  typically exceeds 1000. Such astronomically large exploration spaces make native sampling strategies impractical and prone to divergent learning dynamics. Therefore, it is necessary to develop a new reinforcement learning framework for token combination exploration and optimization.

Based on the above discussions, we propose a novel **Contribution-aware token Compression** algorithm for **VIDEo** understanding (**CaCoVID**) that explicitly optimizes a video token compression policy by actively exploring the contributions of different token combinations to the correct prediction. Specifically, we design a token compression policy network to estimate the contribution of video tokens and frames to correct prediction by interacting with the query. The most contributed tokens are retained for LLM inference as shown in Figure 1(d). Further, to achieve effective combinatorial exploration and policy optimization, we develop a novel combinatorial policy optimization algorithm with online combinatorial space sampling (OCSS). By partitioning the combinatorial sub-space, OCSS constrains sampling within tokens exhibiting similar contribution scores, thereby dramatically reducing ineffective combinatorial exploration and accelerating the convergence of the policy optimization.

To summarize, the main contributions of this work are: (1) To the best of our knowledge, we are the first to propose a reinforcement learning-based token compression algorithm (CaCoVID) that ranks and prunes video tokens by directly estimating the contribution to the correct prediction. (2) We develop a novel combinatorial policy optimization algorithm with online combination space sampling, which dramati-

cally reduces the exploration space for video token combinations and accelerates the convergence of policy optimization. (3) Extensive experimental results on diverse video understanding benchmarks demonstrate that CaCoVID achieves state-of-the-art performance with lower latency.

## Related Work

In this section, we briefly overview related works on video large language models and visual token compression.

### Video Large Language Model

The evolution of video large language models (Ataallah et al. 2024; Chen et al. 2025; Wang et al. 2024) has sparked significant research efforts in developing effective frameworks for spatiotemporal understanding. Due to the redundancy of video tokens, many studies propose video-specific feature learning frameworks to compress video tokens during LLM training. Video-ChatGPT (Maaz et al. 2024) employs temporal and spatial pooling to compress video tokens. PLLaVA (Xu et al. 2024) reduces video tokens to the desired feature dimension via adaptive average structure pooling. Video-XL (Shu et al. 2025) utilizes KV sparsification mechanisms in LLMs to compress the visual information into visual summarization tokens. Further, with the development of unified multimodal frameworks (image, multi-image, video), unified MLLMs attract more attention, *e.g.* LLaVA-OneVision (Li et al. 2025) and Qwen2.5-VL (Bai et al. 2025), which commonly maintain original frame-level tokens to ensure compatibility across modalities. However, these approaches still face challenges in video processing due to high token consumption during inference. Thus, we propose a contribution-aware token compression algorithm for video understanding, a framework-agnostic and compatible solution. We optimize a small compression policy network through active token exploration, effectively unlocking the reasoning potential of pretrained models under limited video token budgets without requiring LLMs retraining.

### Visual Token Compression

Visual feature compression (Bolya et al. 2023; Shen et al. 2025) has gained increasing attention due to the computational burdens caused by redundant visual tokens in multimodal large language models (MLLMs). The design motivations behind these methods can generally be categorized into two paradigms: content-based token compression and model-based token compression.

**Content-based token compression** aims to preserve content diversity or spatio-temporal structure as design objectives, employing handcrafted content metrics to guide token compression. TopV (Yang et al. 2025a) optimizes token pruning through solving a cost matrix based on token distance and similarity. DivPrune (Alvar et al. 2025) enhances diversity preservation by solving a max-min diversity problem to retain structurally distinct tokens. VisionZip (Yang et al. 2025b) reduces spatial redundancy through attention-driven feature analysis in vision encoders. PruneVID (Huang et al. 2025) merges static tokens in

video segments and reduces spatial redundancy via DPC-KNN (Rodriguez and Laio 2014) merging. Although preserving diverse visual contents, their compression strategies are query-agnostic and may prune tokens critical to queries.

**Model-based token compression** aims to minimize perturbations in model inference as design objectives, typically pruning tokens with low attention scores to compress the token sequence. FastV (Chen et al. 2024) prunes visual tokens in early LLM layers through average attention scores received from other tokens. PyramidDrop (Xing et al. 2024) implements multi-stage pruning by progressively eliminating low-attention tokens at different network depths. Dy-Coke (Tao et al. 2025) dynamically prunes key-value cache via prioritizing tokens with high attention scores to the predicted token. SparseVLM (Zhang et al. 2025) filters key-visual tokens through attention scores with question tokens. However, the correlation between attention scores and their actual contribution to correct answers remains ambiguous, which may lead to suboptimal token compression.

To address the above limitations, we shift the focus from passive token preservation to active discovery of optimal compressed token combinations for correct prediction. We achieve this by exploring the contribution of different token combinations to the correct prediction and optimizing a policy network that ranks and prunes video tokens without requiring retraining of the LLM.

## Method

In this section, we first present the preliminaries of video large language models (Video LLMs), elaborating on the critical importance of video token compression for Video LLMs. Subsequently, we formally introduce our proposed contribution-aware token compression for video understanding, which incorporates a learnable compression policy network and a combinatorial policy optimization algorithm.

### Preliminaries

**Video Large Language Models.** Video LLMs typically encode densely sampled video frame features through a visual encoder, then map these visual features into the LLM’s embedding space via a modality projector, enabling alignment between video tokens and language tokens. The concatenation of video tokens  $X_{vid} \in \mathbb{R}^{n_{vid} \times d}$ , and question tokens  $X_{qst} \in \mathbb{R}^{n_{qst} \times d}$  serves as the input to LLM for the prefilling and decoding.

**Computation Complexity.** The encoder of LLM typically consists of the self-attention mechanism and the feed-forward network (FFN). The total computational costs of LLM can be expressed as  $FLOPs = T \times (4nd^2 + 2n^2d + 2ndm)$ . Here,  $T$  is the number of transformer layers,  $n = n_{vis} + n_{qst}$  is the input sequence length,  $d$  is the hidden dimension size, and  $m$  denotes the intermediate size of the FFN. The computational complexity quadratically depends on the sequence length  $n$ . In video LLMs, video token length  $n_{vis}$  dominates the total (e.g.,  $32 \times 196 = 6,272$  video tokens in LLaVA-OneVision), while the question token length  $n_{qst}$  typically remain below 512. Consequently, over 90% of FLOPs arise from attention interactions with video tokens.

This highlights urgent demand for video token compression to eliminate redundant spatial-temporal information.

### Contribution-aware Token Compression

In existing token compression algorithms (Chen et al. 2024; Xing et al. 2024), LLMs fail to actively participate in the token selection process, with tokens instead being passively selected based on handcrafted metrics or pretrained attention scores. The compression objectives may not align with the goals of correct prediction. To address this, we seek to actively explore critical tokens through feedback from LLM and optimize a small network to estimate the contribution of video tokens to correct predictions for token compression.

**Compression Policy Network.** The compression policy network is composed of self-attention mechanisms and two multi-layer perceptrons ( $MLP_t$  and  $MLP_f$ ), which output two-dimensional logits, as shown in Figure 2(b). Given encoded video tokens  $X_{vid} \in \mathbb{R}^{n_{vid} \times d}$  and question tokens  $X_{qst} \in \mathbb{R}^{n_{qst} \times d}$ , the self-attention mechanism first establishes cross-modal interactions between video tokens and text tokens to generate question-aware video tokens  $X'_{vid}$ . Here,  $n_{vid} = t \times h \times w$ , and  $t, h, w$  are the number of frames, height and width of encoded video tokens. Subsequently, these video tokens are fed into two MLPs to estimate the contribution of each video token and frame to correctly answering the question. Formally,

$$[X'_{vid}, X'_{qst}] = \text{SelfAttn}([X_{vid}, X_{qst}]), \quad (1)$$

$$S_f = \text{MLP}_f(\text{Pool}(X'_{vid})) \in \mathbb{R}^{t \times 2}, \quad (2)$$

$$S_t = \text{MLP}_t(X'_{vid}) \in \mathbb{R}^{n_{vid} \times 2}, \quad (3)$$

$$\hat{S}_t = S_t[:, 1] - S_t[:, 0], \hat{S}_f = S_f[:, 1] - S_f[:, 0], \quad (4)$$

where  $\text{Pool}(\cdot)$  means pooling along the spatial dimension, the two output channels of MLPs indicate whether the corresponding token or frame should be selected for answering the question. The difference between the two channels can be regarded as the potential contribution of tokens  $\hat{S}_t$  or frames  $\hat{S}_f$  to correctly answering the question. During inference,  $\hat{S}_t$  is used to select the most contributed token in each frame and  $\hat{S}_f$  is used to allocate token budget across frames.

**Combinatorial Policy Optimization.** There currently exists no large-scale video training data annotated with key tokens for token selection optimization, which precludes the direct application of supervised learning in policy network training. Traditional reinforcement learning algorithms (Schulman et al. 2017; Shao et al. 2024) for large language models typically perform exploration in an autoregressive manner over the vocabulary-sized dimension, and are usually built upon pretrained models. As a result, the exploration space is relatively constrained, which facilitates model convergence. However, for thousands of video token selections, autoregressive estimation of individual token contributions is highly inefficient. Meanwhile, parallel prediction of all token contributions entails an exponentially large combinatorial exploration space  $O(2^n)$ , rendering native sampling-based exploration strategies computationally infeasible and prone to divergent learning dynamics.

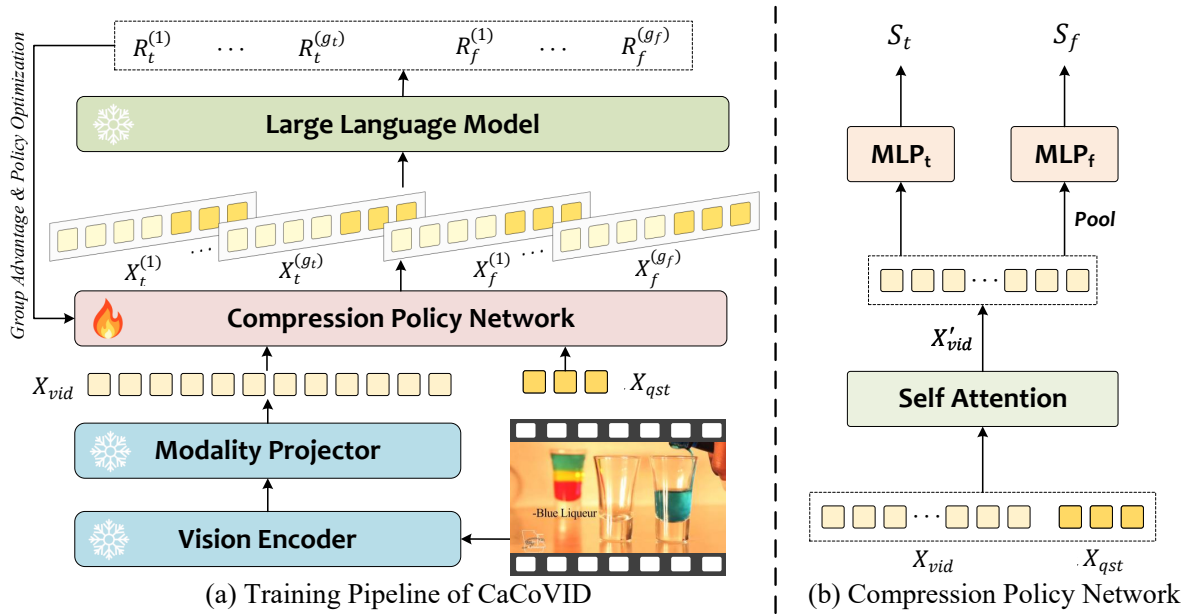


Figure 2: (a) illustrates the training pipeline of CaCoVID, which optimizes the compression policy network by exploring the contributions of different token combinations to correct predictions. (b) presents the architecture of the compression policy network, which builds interactions between video tokens and question tokens in a self-attention mechanism and estimates the contributions of video tokens or frames to the correct prediction by MLPs.

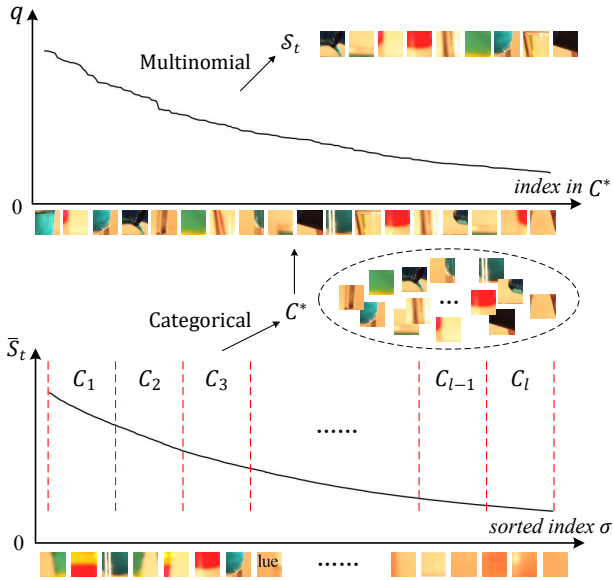


Figure 3: Illustration of the online combinatorial space sampling strategy. Tokens are first sorted by their contribution scores estimated by the policy network and partitioned into combinatorial sub-spaces. Then, a categorical distribution is applied to sample the subspace according to the sum of contribution scores within each subspace. Finally, tokens within the selected subspace are sampled according to a multinomial distribution to acquire the final token combination.

How to narrow down the exploration space and leverage online learning experiences to minimize ineffective exploration constitutes critical issues in policy optimization. To address the above issues, we propose a combinatorial policy optimization algorithm (CPO), integrating a novel online combinatorial space sampling strategy. The goal of the compression policy network is to discover the video token combinations that contribute most to correct predictions, which drives our focus toward systematically exploring token combinations with the highest potential to guide the LLM toward correct answers for given questions, rather than indiscriminately exploring arbitrary combinations. To this end, we develop the online combinatorial space sampling strategy (OCSS), which divides the token combinatorial space into multiple combinatorial sub-spaces according to the estimated contribution scores, such that tokens with similar contribution levels are more likely to be sampled together, as shown in Figure 3. Given a token sampling ratio  $r$ , we first sort all tokens based on their estimated contribution scores of the policy network, and then divide them into  $l = 1/(\lambda \times r)$  combinatorial sub-spaces, each containing  $m = \lambda \times r \times n_{vid}$  tokens,  $\lambda = 2$  by default. Formally,

$$\sigma = \text{argsort}(\hat{S}_t), \quad (5)$$

$$C_i = \{\sigma_j \mid j \in [(i-1)m+1, im]\}, i = 1, 2, \dots, l \quad (6)$$

where  $\sigma$  denotes the sorted index arrangement of video token,  $C_i \subset \mathcal{Z}_t$  is the token index set of  $i^{th}$  combinatorial sub-space,  $\mathcal{Z}_t = \{1, 2, \dots, n_{vid}\}$  is the universal set of the video token indexes. Then, we employ probabilistic sampling based on the total contribution scores of tokens in each

combinatorial sub-space. Formally,

$$\bar{S}_t = \text{Softmax}(\hat{S}_t) \in \mathbb{R}^{n_{vid}}, \quad (7)$$

$$p_i = \sum_{k \in C_i} \bar{S}_t[k], \quad (8)$$

$$C^* \sim \text{Categorical}([p_1, p_2, \dots, p_l]), \quad (9)$$

where  $C^*$  is the sampled combinatorial sub-space,  $\text{Categorical}(\cdot)$  denotes the categorical distribution. Subsequently, a probabilistic sampling is conducted within the selected combinatorial sub-space to further explore its internal token combinations. Formally,

$$q_k = \frac{\exp(\hat{S}_t[k])}{\sum_{u \in C^*} \exp(\hat{S}_t[u])}, \quad k \in C^*, \quad (10)$$

$$\mathcal{S}_t \sim \text{Multinomial}(r \cdot n_{vid}, [q_k]_{k \in C^*}), \quad (11)$$

$$\hat{X}_t = X_{vid}[\mathcal{S}_t, :] \in \mathbb{R}^{(r \cdot n_{vid}) \times d}, \quad (12)$$

where  $\mathcal{S}_t \subset C^*$  is the final sampled  $r \cdot n_{vid}$  video token index,  $\text{Multinomial}(\cdot, \cdot)$  denotes the multinomial distribution,  $\hat{X}_t$  denotes the sampled video tokens.

During each training iteration, we sample  $g_t$  groups of video tokens  $\{\hat{X}_t^{(i)}\}_{i=1}^{g_t}$  and feed them into the LLM along with question tokens  $X_{qst}$  for reasoning, as shown in Figure 2(a). We evaluate the correctness of the predictions by comparing them with the answer  $\mathcal{A}$  as the reward for the video token combinations. Formally,

$$X_t^{(i)} = [\hat{X}_t^{(i)}, X_{qst}], \quad i = 1, 2, \dots, g_t \quad (13)$$

$$Y_t^{(i)} = \text{LLM}(X_t^{(i)}), R_t^{(i)} = \text{Reward}(Y_t^{(i)}, \mathcal{A}), \quad (14)$$

where the reward functions  $\text{Reward}(\cdot, \cdot)$  are consistent with Video-R1 (Feng et al. 2025). Then, we compute the group advantage for policy optimization. Formally,

$$A_t^{(i)} = \frac{R_t^{(i)} - \text{mean}(\{R_t^{(i)}\}_{i=1}^{g_t})}{\text{std}(\{R_t^{(i)}\}_{i=1}^{g_t})}. \quad (15)$$

Finally, we optimize the policy model for token contribution estimation by maximizing the following objective,

$$\mathcal{J}_{CPO}^t(\theta) = \mathbb{E}_{j \in \mathcal{Z}_t} \left[ \frac{1}{g_t} \sum_{i=1}^{g_t} r_{i,j}^{clip}(\theta) \right], \quad (16)$$

$$r_{i,j}^{clip}(\theta) = \min \left[ r_{i,j}(\theta) A_t^{(i)}, \hat{r}_{i,j}(\theta) A_t^{(i)} \right], \quad (17)$$

$$\hat{r}_{i,j}(\theta) = \text{clip}(r_{i,j}(\theta), 1 - \epsilon_{low}, 1 + \epsilon_{high}) \quad (18)$$

$$r_{i,j}(\theta) = \frac{\pi_{\theta}^t(j|[X_{vid}, X_{qst}], \mathcal{S}_t^{(i)})}{\pi_{\theta_{old}}^t(j|[X_{vid}, X_{qst}], \mathcal{S}_t^{(i)})}, \quad j \in \mathcal{Z}_t \quad (19)$$

where  $\epsilon_{low}, \epsilon_{high}$  is a clipping hyper-parameter (Yu et al. 2025) for stabilizing training,  $\pi_{\theta}^t(\cdot|\cdot, \cdot)$  and  $\pi_{\theta_{old}}^t(\cdot|\cdot, \cdot)$  denote the current and old compression policy network with parameters  $\theta$  and  $\theta_{old}$ , whose output can be expressed as,

$$\pi_{\theta}^t(j|[X_{vid}, X_{qst}], \mathcal{S}_t^{(i)}) = \begin{cases} \sigma(\mathcal{S}_t[j, 1]), & \text{if } j \in \mathcal{S}_t^{(i)} \\ \sigma(\mathcal{S}_t[j, 0]), & \text{if } j \in \mathcal{Z}_t \setminus \mathcal{S}_t^{(i)} \end{cases}$$

where  $\sigma(\cdot)$  means softmax along the channel dimension. Similarly, we sample  $g_f$  groups of video frames through the online combinatorial space sampling strategy and calculate the group advantage function through the prediction of LLM to optimize the compression policy network for video frames, as detailed in the supplementary materials.

**Data Exploration Efficiency.** The effective utilization of training samples during model training is another critical aspect for policy network optimization, which involves three key considerations. 1) *Ineffective sample filtering.* Certain video questions can be answered correctly through blind testing (without video input), rendering these samples non-informative for policy network learning. We address this by filtering out simple samples through blind testing. 2) *Experience replay.* To fully explore the video token combinations, we iterate each training sample multiple times combined with experience replay mechanisms to generate more exploration experience. 3) *Dynamic sample ratio.* Specific video-question pairs may consistently produce correct/incorrect responses under the sample ratio  $r$ , which hinders effective exploration. We introduce a dynamic sample ratio strategy during training. For multiple iterations of a training sample, if the average reward from the previous iteration exceeds  $\alpha_{high}$ , we decrease the video sample ratio by half in the next iteration; if the average reward falls below  $\alpha_{low}$ , we increase the video sample ratio by two times. The detailed training algorithm is shown in supplementary materials.

## Experiment

### Experimental Setting

**Training Data.** Our training data is sourced from Video-R1, which aggregates over 116,000 video data. To ensure reliable reward estimation, we specifically selected multiple-choice QA pairs as our training corpus, comprising 104,000 instances. Further, as detailed in Section 3.2, we conducted blind evaluations on these multiple-choice QA data using LLaVA-OneVision-7B (Li et al. 2025) and Qwen2.5-VL-3B (Bai et al. 2025) to filter out trivial questions that could be correctly answered without video context, ultimately retaining a refined dataset of 31,000 high-quality video-relevant questions for policy network training.

**Training Details.** We implement the contribution-aware token compression algorithm for video understanding (CaCoVID) on LLaVA-OneVision-7B and Qwen2.5-VL-3B. The parameters of the vision encoder, projector, and large language model are frozen during training, with only the compression policy network being optimized. The attention mechanism of the compression policy network is initialized using pretrained parameters from the first layer of the LLM, while the MLPs are initialized via Xavier init (Glorot et al. 2010). We employ distinct learning rates of  $1 \times 10^{-7}$  for the attention mechanism and  $1 \times 10^{-6}$  for MLPs. The token sample ratio is set to  $r = 0.02$  and group number is set to  $g_t = 24, g_f = 8$  during training, with each sample undergoing 5 iterative optimization steps with experience replay, as discussed in Section 3.2. We set  $\alpha_{high} = 0.875, \alpha_{low} = 0.125, \epsilon_{high} = 0.28, \epsilon_{low} = 0.2$ . The policy network parameters are optimized through our CPO, completing full

Method	Equivalent Retention Ratio	LongVideoBench 1~60min	MLVU		VideoMME				Avg. Acc.	
			Dev	Test	Short 1~3min	Medium 3~30min	Long 30~60min	Overall 1~60min	Score	%
LLaVA-OneVision-7B	100%	56.4	63.0	47.7	70.3	56.8	48.6	58.6	56.4	100
FastV (Chen et al. 2024)	25%	53.5	59.0	41.7	64.1	53.1	48.1	55.1	52.3	92.7
VisionZip (Yang et al. 2025b)	25%	56.0	61.5	43.5	69.0	55.0	48.3	57.4	54.6	96.8
DivPrune (Alvar et al. 2025)	25%	<b>56.4</b>	<b>62.4</b>	44.1	69.0	54.6	47.9	57.1	55.0	97.5
PruneVID (Huang et al. 2025)	25%	56.1	60.9	45.4	68.6	<b>56.7</b>	<b>48.9</b>	58.0	55.1	97.7
FrameFusion (Fu et al. 2025b)	25%	55.7	61.1	44.2	68.7	55.2	48.2	57.4	54.6	96.7
CaCoVID	25%	56.2	62.3	<b>46.3</b>	<b>70.7</b>	56.3	48.4	<b>58.5</b>	<b>55.8</b>	<b>98.9</b>
FastV (Chen et al. 2024)	20%	53.0	58.3	42.0	62.1	52.2	46.9	53.7	51.8	91.8
VisionZip (Yang et al. 2025b)	20%	55.7	61.5	44.7	68.8	<b>55.9</b>	47.8	<b>57.5</b>	54.9	97.2
DivPrune (Alvar et al. 2025)	20%	55.6	<b>62.2</b>	43.9	68.4	54.4	48.1	57.1	54.7	96.9
PruneVID (Huang et al. 2025)	20%	54.9	61.4	<b>45.7</b>	67.4	54.2	47.8	56.8	54.7	96.9
FrameFusion (Fu et al. 2025b)	20%	55.1	60.5	44.2	67.7	54.4	47.2	56.4	54.1	95.8
CaCoVID	20%	<b>56.5</b>	62.0	45.3	<b>69.7</b>	54.3	<b>48.6</b>	<b>57.5</b>	<b>55.3</b>	<b>98.1</b>
FastV (Chen et al. 2024)	15%	49.8	55.5	41.7	58.0	52.7	46.0	52.2	49.8	88.3
VisionZip (Yang et al. 2025b)	15%	54.0	61.2	41.9	65.7	53.4	<b>48.2</b>	55.8	53.2	94.3
DivPrune (Alvar et al. 2025)	15%	55.1	<b>61.9</b>	43.9	66.6	<b>54.9</b>	48.0	56.5	54.3	96.3
PruneVID (Huang et al. 2025)	15%	55.1	59.9	44.3	66.6	54.4	47.9	56.3	53.9	95.5
FrameFusion (Fu et al. 2025b)	15%	52.9	58.4	41.8	66.2	53.0	46.6	55.3	52.1	92.3
CaCoVID	15%	<b>56.8</b>	61.2	<b>44.9</b>	<b>69.3</b>	54.4	47.4	<b>57.1</b>	<b>55.0</b>	<b>97.5</b>
FastV (Chen et al. 2024)	10%	47.6	52.9	35.2	53.3	48.0	44.2	48.5	46.1	81.6
VisionZip (Yang et al. 2025b)	10%	48.9	57.6	41.5	58.6	52.3	46.2	52.4	50.1	88.8
DivPrune (Alvar et al. 2025)	10%	53.3	<b>60.7</b>	<b>44.9</b>	63.9	52.9	46.2	54.3	53.3	94.5
PruneVID (Huang et al. 2025)	10%	54.4	59.7	42.7	66.0	52.4	46.3	54.9	52.9	93.8
FrameFusion (Fu et al. 2025b)	10%	49.1	56.5	39.8	59.6	49.8	45.1	51.5	49.2	87.2
CaCoVID	10%	<b>55.2</b>	60.6	44.6	<b>67.6</b>	<b>54.9</b>	<b>46.6</b>	<b>56.3</b>	<b>54.2</b>	<b>96.1</b>
Qwen2.5-VL-3B	100%	54.3	62.1	45.6	71.4	60.1	49.6	60.4	55.6	100
FastV (Chen et al. 2024)	25%	49.5	57.0	43.0	66.0	56.0	48.2	56.7	51.6	92.7
VisionZip (Yang et al. 2025b)	25%	51.7	58.5	43.8	67.9	55.6	48.9	57.4	52.9	95.1
DivPrune (Alvar et al. 2025)	25%	<b>53.4</b>	58.8	<b>46.4</b>	68.4	55.2	47.9	57.2	53.9	97.0
PruneVID (Huang et al. 2025)	25%	53.0	59.4	43.0	64.8	53.1	47.7	55.2	52.6	94.6
FrameFusion (Fu et al. 2025b)	25%	52.4	58.5	44.8	<b>69.1</b>	<b>57.8</b>	<b>50.0</b>	<b>59.0</b>	53.7	96.5
CaCoVID	25%	53.3	<b>59.9</b>	46.2	67.7	56.1	48.8	57.5	<b>54.2</b>	<b>97.5</b>
FastV (Chen et al. 2024)	20%	48.2	55.6	41.3	64.1	54.4	47.3	55.3	50.1	90.1
VisionZip (Yang et al. 2025b)	20%	51.9	57.9	43.5	66.3	54.6	48.2	56.4	52.4	94.2
DivPrune (Alvar et al. 2025)	20%	52.3	58.4	<b>45.6</b>	67.6	53.3	47.2	56.0	53.1	95.5
PruneVID (Huang et al. 2025)	20%	52.1	58.1	43.7	64.1	53.4	46.9	54.8	52.2	93.8
FrameFusion (Fu et al. 2025b)	20%	50.7	57.3	43.4	<b>68.0</b>	<b>56.6</b>	<b>49.6</b>	<b>58.0</b>	52.3	94.1
CaCoVID	20%	<b>52.8</b>	<b>60.2</b>	44.8	67.9	<b>56.6</b>	48.4	57.6	<b>53.9</b>	<b>96.9</b>
FastV (Chen et al. 2024)	15%	47.8	53.7	37.3	61.3	52.6	48.2	54.0	48.2	86.7
VisionZip (Yang et al. 2025b)	15%	52.0	57.7	<b>43.8</b>	66.6	53.8	48.6	56.3	52.4	94.3
DivPrune (Alvar et al. 2025)	15%	51.5	57.7	42.0	66.3	52.4	48.0	55.6	51.7	92.9
PruneVID (Huang et al. 2025)	15%	52.1	58.0	42.1	61.8	51.7	47.4	53.6	51.5	92.5
FrameFusion (Fu et al. 2025b)	15%	49.9	55.4	41.5	<b>67.0</b>	<b>55.0</b>	<b>49.2</b>	<b>57.1</b>	51.0	91.6
CaCoVID	15%	<b>52.4</b>	<b>58.2</b>	43.2	66.2	54.4	47.8	56.1	<b>52.5</b>	<b>94.4</b>
FastV (Chen et al. 2024)	10%	45.0	50.8	32.0	55.2	48.1	46.1	49.8	44.4	79.9
VisionZip (Yang et al. 2025b)	10%	<b>51.5</b>	57.2	42.4	63.8	52.1	46.6	54.2	51.3	92.3
DivPrune (Alvar et al. 2025)	10%	50.5	56.4	40.8	<b>64.3</b>	51.4	47.2	<b>54.3</b>	50.5	90.9
PruneVID (Huang et al. 2025)	10%	51.3	56.6	40.4	59.7	51.1	46.6	52.4	50.2	90.2
CaCoVID	10%	<b>51.5</b>	<b>57.9</b>	<b>44.2</b>	63.1	<b>52.2</b>	<b>47.6</b>	<b>54.3</b>	<b>52.0</b>	<b>93.4</b>

Table 1: Comparison of state-of-the-art methods on LLaVA-OneVision-7B and Qwen2.5-VL-3B. For compression algorithms such as FastV and FrameFusion applied in the LLM prefilling phase, we report the results with layer-averaged retention ratio termed equivalent retention ratio as FrameFusion does. Bold highlights optimal performance with similar retention ratios.

training in a single epoch. The training process takes about 10 hours on  $8 \times H100$  GPUs with one batch per GPU. The random seed for training is set as 42.

**Benchmarks.** We evaluate CaCoVID on diverse video understanding benchmarks including LongVideoBench (Wu et al. 2024), MLVU (Zhou et al. 2025), and VideoMME (Fu et al. 2025a). These benchmarks consist of videos with varying temporal scales, task complexities, and domain diversities, providing a comprehensive and generalized metric framework for evaluating performance.

**Inference Details.** We reserve the top contribution tokens of each frame estimated by the compression policy network to answer the given question. The number of tokens  $n_j$  reserved in  $j^{th}$  frame is determined by the estimated contribution of the frame  $n_j = \text{Softmax}(\hat{S}_f)[j] \cdot r \cdot n_{vid}$ . To preserve the spatio-temporal structure of videos, we sample tokens uniformly in the video as complementary to the top-contributed tokens, with these spatio-temporal tokens accounting for 50% of the retention token allocation.

Method	Compression Time (ms)↓	LLM Prefilling Time (ms)↓	Avg. Acc.↑
Vanilla	-	200.4	56.4 (100%)
DivPrune	134.3	53.2	55.0 (97.5%)
PruneVID	34.1	53.2	55.1 (97.7%)
CaCoVID	11.2	53.2	55.8 (98.9%)

Table 2: Comparison of compression efficiency on LLaVA-OneVision-7B with 25% retention ratio.

### State-of-the-art Comparisons

We comprehensively evaluate our CaCoVID under various retention ratios  $r \in \{10\%, 15\%, 20\%, 25\%\}$  and compare it with state-of-the-art methods. The implementation details of these algorithms being compared are provided in the supplementary materials. Accuracy is the main evaluation metric.

**Results on LLaVA-OneVision and Qwen2.5-VL.** During evaluation on LLaVA-OneVision-7B, we uniformly sample 32 frames from the video, and each frame is encoded into 196 tokens. During evaluation on Qwen2.5-VL-3B, we uniformly sample 64 frames from the video, and each frame is encoded into no more than 256 tokens. As shown in Table 1, we compare the performance of different compression algorithms under identical hardware and software environments. Our method consistently outperforms the previous state-of-the-art approach across various retention ratios. Compared to content-based compression algorithms like VisionZip, DivPrune, and PruneVID, CaCoVID achieves question-aware token contribution estimation, effectively retaining tokens critical for answering specific questions. Compared to model-based compression algorithms like FastV and FrameFusion, CaCoVID aligns the optimization objectives of compression algorithms with the prediction accuracy of LLMs, thereby enhancing the overall performance.

**Comparison of compression efficiency.** As shown in Table 2, CaCoVID achieves superior compression speed compared to previous compression algorithms, while demonstrating significantly improved performance. On one hand, the policy network in CaCoVID can estimate the contribution of each token to accurate predictions in parallel, thereby reducing compression latency. On the other hand, our contribution-aware token compression for video understanding explicitly aligns the objective of the compression policy with the correct prediction of LLM, thereby achieving superior performance.

### Ablation Study

We perform ablation study on LLaVA-OneVision-7B with retention ratio 25%.

**Analysis of the online combinatorial space sampling.** Firstly, given the token sampling ratio of  $r\%$ , we explored two training strategies: concatenating  $r\%$  tokens uniformly sampled from each frame by OCSS (Frame) versus directly sampling  $r\%$  tokens from the entire video token set by OCSS (Video). As shown in Table 3, we find that intra-frame token sampling followed by concatenation yields superior performance. The underlying reason could be that the

Sample	Range	$r\%$	MLVU <sub>dev</sub>	VideoMME
OCSS	Video	2%	61.7	57.2
OCSS	Frame	2%	62.3	58.5
Random	Frame	2%	57.9	54.4
Multinomial	Frame	2%	59.2	55.1
OCSS	Frame	1%	62.2	58.3
OCSS	Frame	2%	62.3	58.5
OCSS	Frame	5%	61.6	57.3

Table 3: Performance with different online sampling strategy during training.

ISF	ER	DSR	MLVU <sub>dev</sub>	VideoMME
			60.9	57.1
✓			61.2	57.3
✓	✓		62.1	58.1
✓	✓	✓	62.3	58.5

Table 4: Performance with different approaches for data exploration efficiency.

combinatorial space of intra-frame sampling reduces sampling complexity while preserving the complete temporal structure of the video, which facilitates better understanding of video content by the LLM. Furthermore, we compare different sampling strategies including random sampling, multinomial probability sampling, and our proposed online combinatorial space sampling (OCSS). Experimental results demonstrate that OCSS significantly outperforms both random sampling and multinomial distribution sampling. This improvement stems from OCSS’s ability to narrow down the exploration space and fully leverage online learning experiences to minimize inefficient exploration, thereby enabling the compression policy network to efficiently discover optimal token combinations for question answering. Furthermore, we can observe that training with smaller sampling ratios tends to yield higher performance. This is because numerous token combinations can provide correct answers to the question under high sampling ratios, leading to ineffective policy optimization. Meanwhile, excessively small sampling ratios could also hinder the sampling of effective token combinations required for answering questions, thereby slightly reducing performance.

**Analysis of the data exploration efficiency.** We investigate the impact of data exploration efficiency on model performance, as shown in Table 4. As discussed in Section 3.2, we employ three approaches to enhance data exploration efficiency: Ineffective Sample Filtering (ISF), Experience Replay (ER), and Dynamic Sample Ratio (DSR). We find that filtering out ineffective samples improves model performance, indicating that samples that are guessed correctly in blind tests are detrimental to the learning of the policy network. When implementing experience replay, the performance shows significant improvement as this allows each sample more exploration opportunities, leading to more stable training. The dynamic sample ratio mechanism also enhances model performance by enabling differentiated exploration for samples with varying difficulty levels.

Strategy	MLVU <sub>dev</sub>	VideoMME
FrameAvg	61.2	57.4
FrameAda	62.0	58.1
FrameAda+ST	62.3	58.5

Table 5: Performance with different token retention strategy.

**Analysis of the token retention strategy.** As shown in Table 5, we study token retention strategies during inference, proposing three approaches: 1) retaining equal tokens across all frames, FrameAvg. 2) adaptively allocating token budgets per frame based on contribution scores of frames, FrameAda. 3) further enhancing spatio-temporal structural information by supplementing FrameAda with spatio-temporal tokens, FrameAda+ST. Comparing 1 and 2 reveals that allocating more tokens to frames with higher contributions improves model performance. Comparing 2 and 3 further shows that sacrificing a portion of sub-contributed tokens to maintain the video’s spatiotemporal structure enhances the video understanding capabilities.

**Visualization** analysis is shown in supplementary materials.

## Conclusion

In this paper, we present CaCoVID, the first reinforcement learning-based token compression framework for video understanding that optimizes a policy network to rank and prune video tokens by actively exploring optimal token combinations to correct predictions. We propose a novel combinatorial policy optimization algorithm with online combination space sampling, which effectively narrows exploration complexity while accelerating policy convergence. Extensive experiments on diverse video understanding benchmarks demonstrate the effectiveness of CaCoVID.

## References

Alvar, S. R.; Singh, G.; Akbari, M.; and Zhang, Y. 2025. DivPrune: Diversity-based Visual Token Pruning for Large Multimodal Models. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 9392–9401.

Ataallah, K.; Shen, X.; Abdelrahman, E.; Sleiman, E.; Zhu, D.; Ding, J.; and Elhoseiny, M. 2024. MiniGPT4-Video: Advancing Multimodal LLMs for Video Understanding with Interleaved Visual-Textual Tokens. In *Workshops of Proceedings of the Computer Vision and Pattern Recognition Conference*.

Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; Zhong, H.; Zhu, Y.; Yang, M.; Li, Z.; Wan, J.; Wang, P.; Ding, W.; Fu, Z.; Xu, Y.; Ye, J.; Zhang, X.; Xie, T.; Cheng, Z.; Zhang, H.; Yang, Z.; Xu, H.; and Lin, J. 2025. Qwen2.5-VL Technical Report. arXiv:2502.13923.

Bolya, D.; Fu, C.-Y.; Dai, X.; Zhang, P.; Feichtenhofer, C.; and Hoffman, J. 2023. Token Merging: Your ViT but Faster. In *International Conference on Learning Representations*.

Chen, L.; Zhao, H.; Liu, T.; Bai, S.; Lin, J.; Zhou, C.; and Chang, B. 2024. An Image is Worth 1/2 Tokens After Layer

2: Plug-and-Play Inference Acceleration for Large Vision-Language Models. In *European Conference on Computer Vision*.

Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Cui, E.; Zhu, J.; Ye, S.; Tian, H.; Liu, Z.; Gu, L.; Wang, X.; Li, Q.; Ren, Y.; Chen, Z.; Luo, J.; Wang, J.; Jiang, T.; Wang, B.; He, C.; Shi, B.; Zhang, X.; Lv, H.; Wang, Y.; Shao, W.; Chu, P.; Tu, Z.; He, T.; Wu, Z.; Deng, H.; Ge, J.; Chen, K.; Zhang, K.; Wang, L.; Dou, M.; Lu, L.; Zhu, X.; Lu, T.; Lin, D.; Qiao, Y.; Dai, J.; and Wang, W. 2025. Expanding Performance Boundaries of Open-Source Multimodal Models with Model, Data, and Test-Time Scaling.

Feng, K.; Gong, K.; Li, B.; Guo, Z.; Wang, Y.; Peng, T.; Wu, J.; Zhang, X.; Wang, B.; and Yue, X. 2025. Video-R1: Reinforcing Video Reasoning in MLLMs. arXiv:2503.21776.

Fu, C.; Dai, Y.; Luo, Y.; Li, L.; Ren, S.; Zhang, R.; Wang, Z.; Zhou, C.; Shen, Y.; Zhang, M.; Chen, P.; Li, Y.; Lin, S.; Zhao, S.; Li, K.; Xu, T.; Zheng, X.; Chen, E.; Shan, C.; He, R.; and Sun, X. 2025a. Video-MME: The First-Ever Comprehensive Evaluation Benchmark of Multi-modal LLMs in Video Analysis.

Fu, T.; Liu, T.; Han, Q.; Dai, G.; Yan, S.; Yang, H.; Ning, X.; and Wang, Y. 2025b. FrameFusion: Combining Similarity and Importance for Video Token Reduction on Large Visual Language Models. In *International Conference on Computer Vision*.

Glorot, X.; et al. 2010. Understanding the difficulty of training deep feedforward neural networks.

Huang, X.; Zhou, H.; Han, K.; et al. 2025. PruneVid: Visual Token Pruning for Efficient Video Large Language Models. In *Annual Meeting of the Association for Computational Linguistics*.

Kang, S.; Kim, J.; Kim, J.; and Hwang, S. J. 2025. See What You Are Told: Visual Attention Sink in Large Multimodal Models. In *International Conference on Learning Representations*.

Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; and Li, C. 2025. LLaVA-OneVision: Easy Visual Task Transfer. *Transactions on Machine Learning Research*.

Lin, B.; Ye, Y.; Zhu, B.; Cui, J.; Ning, M.; Jin, P.; and Yuan, L. 2024. Video-LLaVA: Learning United Visual Representation by Alignment Before Projection. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 5971–5984. Miami, Florida, USA: Association for Computational Linguistics.

Maaz, M.; Rasheed, H.; Khan, S.; and Khan, F. S. 2024. Video-ChatGPT: Towards Detailed Video Understanding via Large Vision and Language Models. In *Annual Meeting of the Association for Computational Linguistics*.

Rodriguez, A.; and Laio, A. 2014. Clustering by fast search and find of density peaks. *Science*, 344(6191): 1492–1496.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms.

- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y. K.; Wu, Y.; and Guo, D. 2024. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models.
- Shen, L.; Hao, T.; He, T.; Zhao, S.; Zhang, Y.; Liu, P.; Bao, Y.; and Ding, G. 2025. TempMe: Video Temporal Token Merging for Efficient Text-Video Retrieval. In *International Conference on Learning Representations*.
- Shu, Y.; Liu, Z.; Zhang, P.; Qin, M.; Zhou, J.; Liang, Z.; Huang, T.; and Zhao, B. 2025. Video-XL: Extra-Long Vision Language Model for Hour-Scale Video Understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*.
- Tao, K.; Qin, C.; You, H.; Sui, Y.; and Wang, H. 2025. DyCoke: Dynamic Compression of Tokens for Fast Video Large Language Models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*.
- 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.
- Wu, H.; Li, D.; Chen, B.; and Li, J. 2024. LongVideoBench: A Benchmark for Long-context Interleaved Video-Language Understanding. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 28828–28857. Curran Associates, Inc.
- Xing, L.; Huang, Q.; Dong, X.; Lu, J.; Zhang, P.; Zang, Y.; Cao, Y.; He, C.; Wang, J.; Wu, F.; et al. 2024. Pyramidrop: Accelerating your large vision-language models via pyramid visual redundancy reduction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*.
- Xu, L.; Zhao, Y.; Zhou, D.; Lin, Z.; Ng, S. K.; and Feng, J. 2024. PLLaVA : Parameter-free LLaVA Extension from Images to Videos for Video Dense Captioning.
- Yang, C.; Sui, Y.; Xiao, J.; Huang, L.; Gong, Y.; Li, C.; Yan, J.; Bai, Y.; Sadayappan, P.; Hu, X.; et al. 2025a. Topv: Compatible token pruning with inference time optimization for fast and low-memory multimodal vision language model. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 19803–19813.
- Yang, S.; Chen, Y.; Tian, Z.; Wang, C.; Li, J.; Yu, B.; and Jia, J. 2025b. VisionZip: Longer is Better but Not Necessary in Vision Language Models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*.
- Yu, Q.; Zhang, Z.; Zhu, R.; Yuan, Y.; Zuo, X.; Yue, Y.; Dai, W.; Fan, T.; Liu, G.; Liu, L.; Liu, X.; Lin, H.; Lin, Z.; Ma, B.; Sheng, G.; Tong, Y.; Zhang, C.; Zhang, M.; Zhang, W.; Zhu, H.; Zhu, J.; Chen, J.; Chen, J.; Wang, C.; Yu, H.; Song, Y.; Wei, X.; Zhou, H.; Liu, J.; Ma, W.-Y.; Zhang, Y.-Q.; Yan, L.; Qiao, M.; Wu, Y.; and Wang, M. 2025. DAPO: An Open-Source LLM Reinforcement Learning System at Scale.
- Zhang, Y.; Fan, C.-K.; Ma, J.; Zheng, W.; Huang, T.; Cheng, K.; Gudovskiy, D.; Okuno, T.; Nakata, Y.; Keutzer, K.; et al. 2025. SparseVLM: Visual Token Sparsification for Efficient Vision-Language Model Inference. In *International Conference on Machine Learning*.
- Zhou, J.; Shu, Y.; Zhao, B.; Wu, B.; Liang, Z.; Xiao, S.; Qin, M.; Yang, X.; Xiong, Y.; Zhang, B.; Huang, T.; and Liu, Z. 2025. MLVU: Benchmarking Multi-task Long Video Understanding.