

Beyond Training: Dynamic Token Merging for Zero-Shot Video Understanding

Yiming Zhang^{1,2}, Zhuokai Zhao³, Zhaorun Chen³, Zenghui Ding^{1*}, Xianjun Yang^{1,2}, Yining Sun^{1,2}
¹HFIPS, Chinese Academy of Sciences ²University of Science and Technology of China
³University of Chicago

Abstract

Recent advancements in multimodal large language models (MLLMs) have opened new avenues for video understanding. However, achieving high fidelity in zero-shot video tasks remains challenging. Traditional video processing methods rely heavily on fine-tuning to capture nuanced spatial-temporal details, which incurs significant data and computation costs. In contrast, training-free approaches, though efficient, often lack robustness in preserving context-rich features across complex video content. To this end, we propose DYTO, a novel dynamic token merging framework for zero-shot video understanding that adaptively optimizes token efficiency while preserving crucial scene details. DYTO integrates a hierarchical frame selection and a bipartite token merging strategy to dynamically cluster key frames and selectively compress token sequences, striking a balance between computational efficiency with semantic richness. Extensive experiments across multiple benchmarks demonstrate the effectiveness of DYTO, achieving superior performance compared to both fine-tuned and training-free methods and setting a new state-of-the-art for zero-shot video understanding.

1. Introduction

In recent years, video understanding has seen substantial progress, largely thanks to the rapid advancements in multimodal large language models (MLLMs) [27, 33, 49]. Traditional video understanding methods often rely on specific training to align video frames with natural language, using spatial-temporal cues to construct coherent narratives across video sequences [1, 21, 22]. In contrast, MLLM-based approaches provide a more flexible and generalized framework, incorporating diverse open-world knowledge across multiple data modalities in the pre-training phase [24, 29]. By utilizing this pre-trained knowledge, MLLMs can dynamically adapt to various tasks such as captioning [44], question answering [10], retrieval [8], and zero-shot or few-shot reasoning over various video content [4]. MLLM-based video understanding approaches generally fall into

two categories, where one requires domain-specific finetuning [6, 9, 15, 18, 20, 23, 26, 41, 43, 47, 48], and the other is completely training-free [5, 11, 35, 39].

Often times, trade-offs arise between accuracy and efficiency in these approaches. While finetuned models typically achieve higher accuracy by tailoring their capabilities to specific tasks, this comes at the cost of extensive labeled data and increased computational resources. In contrast, training-free methods are more efficient and flexible, leveraging generalization capabilities of pre-trained MLLMs for zero-shot inference. However, they may fall short in some settings, where specific reasoning tasks or varied video types require more nuanced understanding. Thus, balancing these trade-offs remains a pivotal challenge in advancing MLLM-based video understanding.

While training-free methods offer the significant advantage of zero-shot adaptability, a central challenge lies in retaining crucial spatial-temporal details across frames without specific finetuning or frame-by-frame annotation. Existing work has often encountered trade-offs between retaining semantic richness and maintaining computational efficiency. For instance, many approaches reduce frame counts and perform aggressive token pooling [11, 35, 39], which risk losing contextually significant frames or visual tokens that capture nuanced actions. Consequently, existing methods struggle to balance semantic fidelity with token efficiency in a manner that adapts to varying temporal complexities across video content.

To this end, we propose **DYTO**, a novel **DY**namic **TO**ken merging approach that leverages hierarchical frame selection and a fine-grained bipartite token compression method. DYTO effectively preserves essential spatial-temporal information while significantly reducing token redundancy. Specifically, our method dynamically clusters key frames across hierarchical temporal scales, allowing it to adaptively capture critical events without rigid sampling strategies. Furthermore, by implementing a bipartite token merging mechanism, we optimize token counts based on the semantic content of each frame, which supports richer video representations under computational constraints.

By employing this dynamic, adaptive framework, DYTO

*Correspondence to Zenghui Ding: dingzenghui@iim.ac.cn.

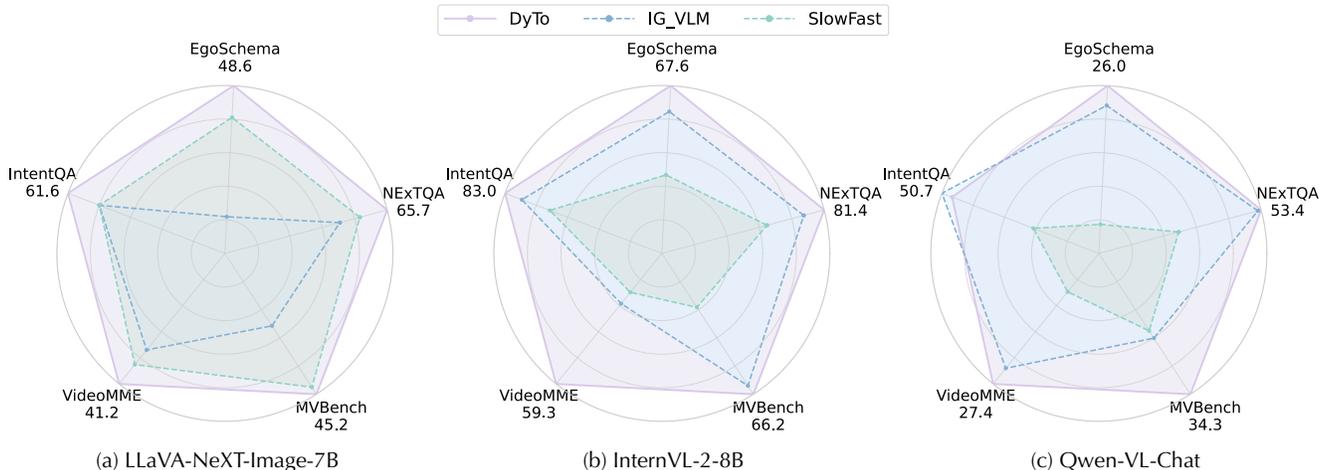


Figure 1. Comparison with two SOTA training-free video-based LLM decoding methods over three different model backbones across five video benchmarks. DYTO and other baselines are marked using solid (—) and dashed (---) lines, respectively. DYTO outperforms existing training-free SOTA methods on almost all the benchmarks and achieves even better performance than most SFT-based methods.

achieves superior performance in zero-shot video understanding, outperforming both finetuned models and other training-free methods. This paper’s contributions include:

- A novel hierarchical bipartite merging strategy that dynamically selects key frames and performs adaptive token merging to optimize spatial-temporal fidelity and achieve more fine-grained feature retention across extended frame sequences.
- Comprehensive evaluation across multiple benchmarks, demonstrating DYTO’s superior performance in both understanding capabilities and computational efficiency over existing methods, including both finetuned and training-free approaches.

2. Related Work

2.1. Video Large Language Models

Video language model has witnessed significant advancements over the past years [6, 20, 32, 38, 46]. VideoChatGPT [26] retrieves features from each frame and subsequently combines them using two operations for spatial and temporal pooling prior to feeding them into a large language model. LLaMA-VID [18] utilized a dual token approach to effectively compress the video token by differentiating between context and content. Vista-LLaMA [23] presented EDVT-Attention along with a sequential vision projector that emphasizes visual tokens while decreasing temporal tokens by sequentially merging them using a Q-former. Video-LLaVA [20] aligns the encoders for images and videos beforehand, enabling shared projections and joint training across both image and video tasks, thus mapping them into the language space. VideoChat [15] utilized cross-attention to condense video tokens alongside user inquiries and conversational context. Video-LLaMA [43] in-

troduces a Video Q-Former and an Audio Q-Former, allowing for the integration of multiple modalities in video understanding. In contrast, Video-LLaMA2 [6] designs a spatial-temporal convolution connector to replace the Q-Former for spatial-temporal representation learning. Chat-UniVi [9] developed a unified model for images and videos that uses dynamic token merging with k-NN to simplify spatial and temporal tokens. SeViLA [41] focused on extracting key frames that are relevant to the inquiries and analyzed the video by transforming these keyframes into video tokens.

2.2. Training-free Video LLMs

Recent research explored and demonstrated that Image LLMs require no additional fine-tuning to apply for video understanding scenarios. FreeVA [35] explores the different spatial-temporal pooling strategies and versions of close-sourced GPT evaluation that influence the video understanding performance. IG-VLM [11] design image grid format and assemble multiple video frames as an image before sending them to an Image LLM. SlowFast-LLaVA [39] (SF-LLaVA for short) introduced a new fusion technique for short-long sampling and various pooling strategies. These methods demonstrate promising results across various video benchmarks, but they have two main limitations. Initially, they all sampled video frames uniformly to a fixed length as the representation of the video. This approach inevitably loses important event information within the videos. Although SF-LLaVA samples longer sequences than other methods, its design of using a small number of tokens fails to capture abundant spatial information present in each frame. Furthermore, the average or maximum pooling method employed does not adequately preserve the significant changes in action over the temporal dimension. In this paper, we present a new method to dynamically select

frames and merge visual tokens, enabling us to comprehensively and efficiently capture the complete semantic information of every video. We also expand the sample video frame sequence to 100 or more frames to enhance performance on longer video understanding task.

3. Method

To reduce the loss of crucial information during token compression, we introduce a hierarchical bipartite graph merging mechanism. Initially, For a video that uniformly sampling with N Frames $X = \{I_1, I_2, I_3, \dots, I_N\}$, the frames of the video are encoded separately by visual encoders $f : Visual_{enc}(I_i)$, producing visual tokens $V \in \mathbb{R}^{N \times L \times D}$. These are subsequently processed by the token aggregator module, transitioning from coarse-grained to fine-grained feature.

3.1. Coarse-grained Hierarchical Clustering

For N frames tokens sequence $V = \{v_1, v_2, v_3, \dots, v_N\}$ of a video, $v_i \in \mathbb{R}^{L \times D}$, we define a directed graph $G = (C, E)$ where C is the set of nodes(the points to be clustered) and the edges E connect the nodes in the feature space. Previous studies demonstrate the deep layers tend to aggregate global semantics in the image [14, 30], therefore we collect the [CLS] tokens of each v_i and concat them as the compressed representation $\mathbf{v} \in \mathbb{R}^{N \times D}$ of the video. The time-stamps are defined as $T = 1, 2, \dots, N$. We build W_t by calculating the feature space distances between frames and adjusting them based on their temporal positions.

$$W_t(i, j) = \begin{cases} (1 - \langle \mathbf{v}_i, \mathbf{v}_j \rangle) \cdot |t_i - t_j|/N & \text{if } i \neq j \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

v_i and v_j indicate the [CLS] token embedding. The inner product is calculated on L2-normalized feature vectors to keep the distance within the $[0, 1]$ range. The term $|t_i - t_j|/N$ serves as a weighting mechanism based on the sequence length. Finally, from this we construct a 1-NN graph by keeping only the closest node to each node and setting all other edges to zero.

$$G(i, j) = \begin{cases} 0 & \text{if } W_t(i, j) > \min_{\forall j} W_t(i, j) \\ 1 & \text{otherwise} \end{cases} \quad (2)$$

The 1-NN temporal graph G defines an adjacency matrix where each node is connected to its nearest neighbor based on the temporally weighted distances W_t . We make the links symmetric by setting $G(j, i) = 1$ to encode space-temporal distances and form connected components to clusters conveniently. The connected components of Graph in Equ.2 automatically partition the data into discovered clusters. We then recursively merge these clusters according to the similarity of their feature averages. Throughout this

merging process, we identify multiple clustering results and select the second-largest one as the final segmentation of the video. We denote by n the index set of frames belonging to the n_{th} cluster. Consequently, the set of frames in the n_{th} cluster can be written as:

$$V_{all} = \{V_n | n \in \{1, 2, \dots, K\}\} \quad (3)$$

We uniformly sample frames from every cluster and combine them as the keyframe sequence $V_s = \{v_1, v_2, \dots, v_K\} \in \mathbb{R}^{K \times (L-1) \times D}$ from the segmentations. The segmentation can yield the distribution of events contained in the current video sequence, thereby guiding the fine-grained merging in subsequent steps.

3.2. Fine-grained Dynamic Bipartite Merging

Previous method [11, 35, 39] either simply concatenates frames and designs prompts for the VLM or uses pooling method on each image token to construct the video representation. DYTO propose a novel method to effectively construct video representations from selected frames. After acquiring the segmentations of the frame sequence through the previously described method, we implement a dynamic bipartite token merging approach inspired by ToMe [2] to minimize the number of visual tokens. For each individual frame in v_k , we sequentially divide the R_i tokens into two non-overlap tokens sets \mathbb{P} with r_i tokens and \mathbb{Q} with $R_i - r_i$ tokens at the i^{th} steps, where initial $R_0 = L - 1$ (exclude CLS token). We set the $r_i = R_0 - \frac{Z}{K}$ to dynamically merge the image to preserve more tokens under the fixed visual token length Z . To obtain the similarity scores, each visual token is split into H heads along channel dimensions, each with D/H channels. The similarity score for each token pair is obtained by averaging the cosine similarity scores over all heads following Eq. (4)

$$a^{p_i q_i} = \frac{1}{H} \left[\sum_{h=1}^H \cos \left(\mathbf{p}_h^{(p_i)}, \mathbf{p}_h^{(q_i)} \right) \right] \quad (4)$$

where $p_i \in \{1, \dots, r_i\}$ and $q_i \in \{1, \dots, (R_i - r_i)\}$ are the indexes of patch feature \mathbf{p} in set \mathbb{P}_i and \mathbb{Q}_i , respectively. We choose the top- r_i token pairs that have the highest similarity scores and combine the matched tokens through pooling. Finally, the remaining tokens from both sets are joined back together, resulting in $R_i - r_i$ after the i^{th} step. In the end, we efficiently encode the video features from $N \times L \times D$ to $K \times (L - \frac{Z}{K} - 1) \times D$ to reduce redundancy in the original visual tokens sequence.

4. Experiment

To evaluate DYTO, we employ over 10 benchmarks that test both structured and open-ended video question answering (VQA) capabilities of the method, specifically using the multiple-choice and GPT-assisted benchmarks.

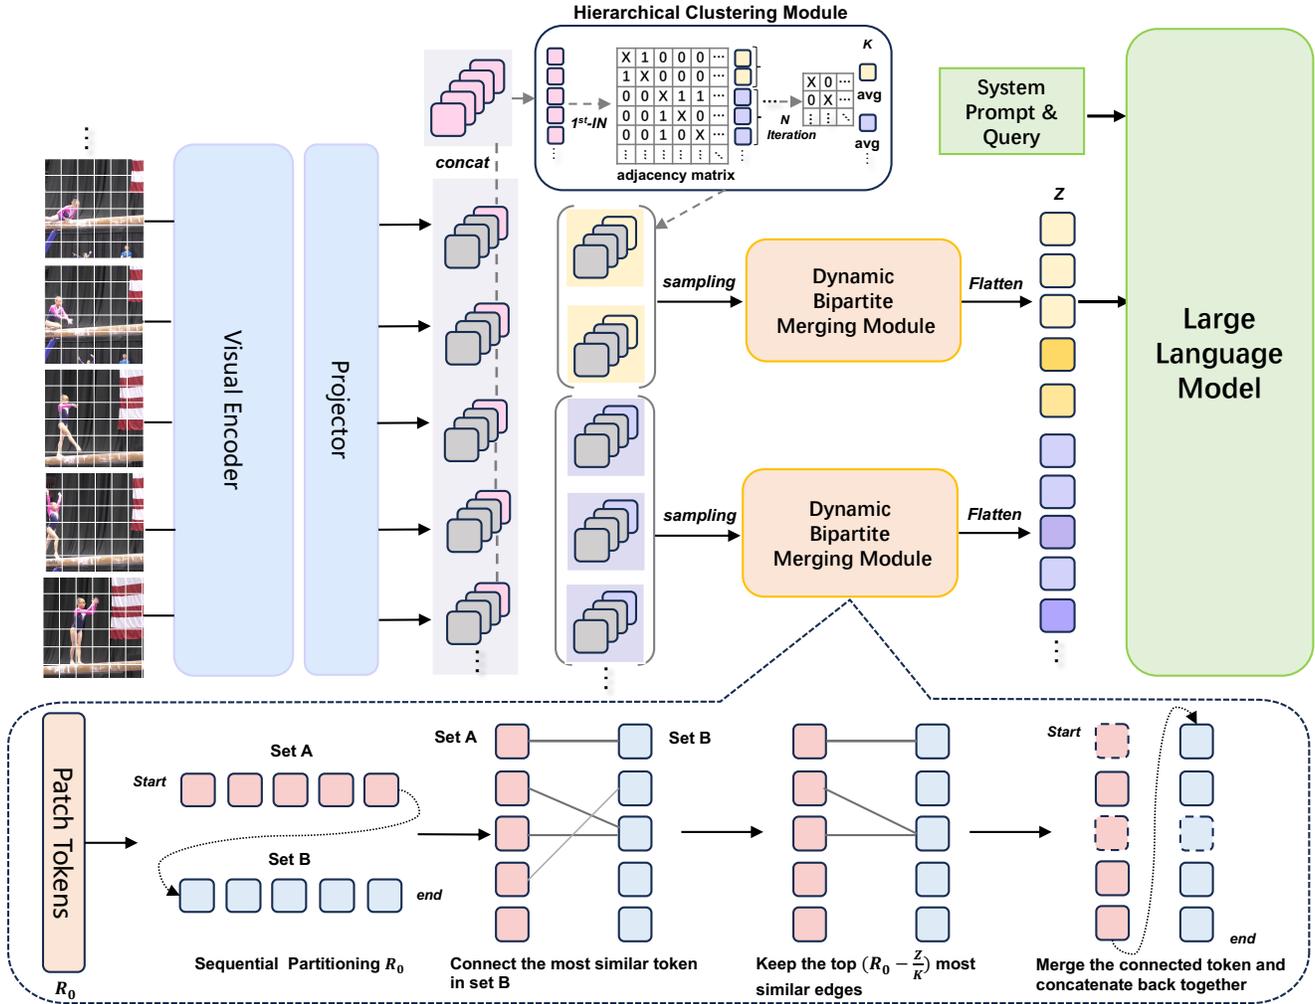


Figure 2. The overview of DYTO, a training-free model built upon image-based MLLM without any fine-tuning. Specifically, DYTO first divides the video into K clusters using the [CLS] token (pink block). Then the dynamic bipartite merging module samples frames from each cluster and controls the final output length as Z , resulting in better balance between computational efficiency and semantic richness.

4.1. Evaluation Benchmarks

Structured VQA benchmarks. For structured VQA, we evaluate DYTO on a diverse set of multiple-choice benchmarks, including NExTQA [36], EgoSchema [28], IntentQA [13], VideoMME [7], and MVBench [16], all of which are designed to quantify model’s capability on video understanding by selecting the correct answer among predefined options. Collectively, these benchmarks provide a comprehensive evaluation of DYTO to interpret complex, multimodal data and to select accurate, contextually rich responses across varying levels of task structures. Notably, we perform experiments on the VideoMME [7] benchmark under the “w/o subs” configuration, which restricts access to subtitles, thereby isolating the model’s reliance on visual and temporal cues in video understanding.

Open-ended VQA benchmarks. We also assess DYTO on the open-ended VQA tasks, specifically, we evaluate the zero-shot performance of DYTO on MSVD-QA [3], MSRVT-QA [37], TGIF-QA [17], ANet-QA [42], and Video-ChatGPT Generation (VCG) [24] benchmarks. These benchmarks require the model to autonomously generate free-form responses, simulating real-world question-answering contexts that demand nuanced understanding of video content. Specifically, for VCG benchmark, we assess across five key dimensions, including Correctness of Information (CI), Detail Orientation (DO), Contextual Understanding (CU), Temporal Understanding (TU), and Consistency (CO). Following Wu [35], we utilize *GPT-3.5-Turbo-0125* to ensure fair comparisons with other methods.

Method	LVLMM Size	Base Model	Frame Length	NExTQA acc. (%)	EgoSchema acc. (%)	IntentQA acc. (%)	VideoMME acc. (%)	MVBench acc. (%)
<i>Models fine-tuned with video data</i>								
SeViLA [40]	7B	CLIP-L + FlanT5	8	63.6	25.7	60.9	44.6	-
Video-LLaVA [20]	7B	ViT-L + Vicuna	8	-	38.4	-	40.4	41.0
LLaMA-VID [19]	7B	EVA-G + Vicuna	1fps	-	38.5	-	-	41.9
InternVideo2 [34]	7B	VideoMAEv2-g+InternVL	16	59.1	32.1	-	41.6	-
LLaVA-NeXT-Video [45]	7B	CLIP-L + Vicuna	32	-	43.9	-	46.5	33.7
VideoLLaMA2 [6]	7B	CLIP-L+ Mistral-Instruct	32	-	51.7	-	46.6	54.6
VideoChat2 [16]	7B	UMT-L+Vicuna	16	61.7	-	59.0	39.5	51.1
LLaVA-OneVision [12]	7B	SigLIP+Qwen-2	32	79.4	60.1	-	58.2	56.7
VideoLLaMA2 [6]	46.7B	CLIP-L+ Mistral-Instruct	32	-	53.3	-	47.9	53.9
LLaVA-NeXT-Video [45]	32B	CLIP-L + Qwen1.5	32	77.3	60.9	-	60.2	-
<i>Training-free approaches</i>								
IG-VLM [11]	7B		6	63.1	35.8	60.1	39.8	41.3
SlowFast [39]	7B	LLaVA-NeXT-image	50	64.2	45.5	60.1	40.4	44.8
DYTO	7B		dynamic	65.7	48.6	61.6	41.2	45.2
IG-VLM [11]	8B		6	79.9	65.6	81.8	49.4	65.3
SlowFast [39]	8B	InternVL2	50	77.2	60.7	79.8	48.0	57.0
DYTO	8B		dynamic	81.4	67.6	83.0	59.3	66.2
IG-VLM [11]	9B		6	53.2	24.8	51.4	26.5	32.0
SlowFast [39]	9B	Qwen-VL-Chat	50	48.4	17.6	45.1	22.1	31.7
DYTO	9B		dynamic	53.4	26.0	51.7	27.4	34.3
IG-VLM [11]	26B		6	80.6	56.0	83.1	50.8	66.4
SlowFast [39]	26B	InternVL2	50	79.2	54.8	82.8	49.4	61.8
DYTO	26B		dynamic	81.1	59.2	83.6	53.0	68.1
IG-VLM [11]	34B		6	70.9	53.6	65.3	52.0	48.4
SlowFast [39]	34B	LLaVA-NeXT-image	50	71.9	55.8	66.2	53.2	51.2
DYTO	34B		dynamic	72.9	56.8	67.5	53.4	52.9

Table 1. Structured VQA benchmarks results comparing DYTO with SOTA training-free approaches as well as models that have been fine-tuned with additional video data from various LVLMM architectures.

4.2. Experimental Setup

Input video and model setting. In our approach, we uniformly sample each video to $N = 100$ frames. Each frame is resized to match the input dimensions of different visual encoder, which then outputs visual tokens plus a [CLS] token. Following Sarfraz et al. [31], we collect the [CLS] token as the coarse-grained feature for each frame and subsequently partition the sequence into K clusters, using recursive iterations of *hierarchical clustering* (§3.1) to group semantically similar frames.

Due to hardware limitations, we set the visual token sequence length Z to either 3680 or 7200, corresponding to model sizes of 7B and 34B, respectively. To optimize feature representation and mitigate token redundancy, we dynamically adjust the merge ratio based on the formulation $r = Z/K$.

4.3. Main Results

Structured VQA benchmarks. As shown in Table 1, DYTO outperforms all the training-free approaches as well

as all the fine-tuned models across all benchmarks by clear margins. Notably, DYTO demonstrated its adaptability to diverse video understanding contexts, achieving superior accuracy even compared to models that employ extensive fine-tuning. For example, on NExTQA [36], DYTO sets a new state-of-the-art performance by achieving 81.4% correctness when pairing with InternVL2-8B base model, which is significantly higher than the 59.1% accuracy achieved by InternVideo2 [34]. Similarly, performance on EgoSchema [28], IntentQA [13], VideoMME [7], and MVBench [16] illustrate the effectiveness of DYTO in handling task-specific reasoning within video content.

Open-ended VQA benchmarks. In open-ended settings, as shown in Table 2, DYTO also demonstrates competitive zero-shot capabilities. It consistently outperforms existing methods on benchmarks including MSVD-QA [3], MSRVT-QA [37], TGIF-QA [17], and ANet-QA [42]. Notably, DYTO performs exceptionally well on VideoChatGPT [24], surpassing even those methods that require

Method	LVLM Size	Base Model	Frame Length	MSVD-QA	MSRVTT-QA	TGIF-QA	ANet-QA	Video-ChatGPT (VCG Benchmark)					
				acc./score	acc./score	acc./score	acc./score	CI	DO	CU	TU	CO	Average
<i>Models fine-tuned with video data</i>													
Video-ChatGPT[25]	7B	CLIP-L+Vicuna	6	64.9/3.5	49.3/2.9	51.4/3.0	35.2/2.7	2.50	2.57	2.69	2.16	2.20	2.42
VideoGPT+[26]	3.8B	CLIP-L+Phi-3-mini	16	72.4/3.9	60.6/3.6	74.6/4.1	50.6/3.6	3.27	3.18	3.74	2.83	3.39	3.28
Video-LLaVA[20]	7B	ViT-L + Vicuna	8	70.7/3.9	59.2/3.5	70.0/4.0	45.3/3.3	-	-	-	-	-	-
MovieChat[32]	7B	CLIP-G+Vicuna	2048	75.2/3.8	52.7/2.6	-	45.7/3.4	2.76	2.93	3.01	2.24	2.42	2.67
LLama-VID[19]	13B	EVA-G+Vicuna	1fps	69.7/3.7	57.7/3.2	-	47.4/3.3	2.96	3.00	3.53	2.46	2.51	2.89
VideoChat2[16]	7B	UMT-L+Vicuna	16	70.0/3.9	54.1/3.3	-	49.1/3.3	3.02	2.88	3.51	2.26	2.81	2.98
Vista-LLaMA[23]	7B	CLIP-L+Vicuna	16	65.3/3.6	60.5/3.3	-	48.3/3.3	2.44	2.64	3.18	2.26	2.31	2.57
Video-LLaMA2[6]	7B	CLIP-L+Mistral-Instruct	32	70.9/3.8	-	-	50.2/3.3	3.16	3.08	3.69	2.56	3.14	3.13
PLLaVA[38]	7B	CLIP-L+Vicuna	16	76.6/4.1	62.0/3.5	77.5/4.1	56.3/3.5	-	-	-	-	-	-
LLaVA-NeXT-Video[45]	7B	CLIP-L + Vicuna	32	-	-	-	53.5/3.5	-	-	-	-	-	-
Video-LLaMA2[6]	46.7B	CLIP-L+Mistral-Instruct	32	70.5/3.8	-	-	50.3/3.4	3.08	3.11	3.64	2.67	3.26	3.15
LLaVA-NeXT-Video[45]	34B	CLIP-L + Qwen1.5	32	-	-	-	58.8/3.4	-	-	-	-	-	-
<i>Training-free approaches</i>													
IG-VLM-7B[11]	7B	LLaVA-NeXT-image	6	78.3/3.9	63.7/3.4	72.7/4.0	53.8/3.2	3.11	2.78	3.48	2.44	3.29	3.03
IG-VLM-34B[11]	34B	LLaVA-NeXT-image	6	79.6/4.1	62.4/3.5	79.1/4.2	58.4/3.5	3.21	2.87	3.54	2.51	3.34	3.09
SlowFast-7B[39]	7B	LLaVA-NeXT-image	50	78.7/3.9	66.2/3.4	77.5/4.0	53.9/3.1	3.09	2.70	3.57	2.52	3.35	3.04
SlowFast-34B[39]	34B	LLaVA-NeXT-image	50	78.7/4.1	67.1/3.7	80.6/4.3	58.8/3.5	3.48	2.96	3.84	2.70	3.54	3.30
DyTO	7B	LLaVA-NeXT-image	dynamic	77.6/3.9	64.1/3.4	78.0/4.0	54.3/3.2	3.12	2.93	3.72	2.52	3.52	3.16
DyTO	34B	LLaVA-NeXT-image	dynamic	79.6/4.1	<u>66.2/3.6</u>	<u>80.7/4.2</u>	59.0/3.5	<u>3.45</u>	2.95	3.82	<u>2.72</u>	<u>3.63</u>	<u>3.32</u>

Table 2. Open-ended VQA benchmark results for DyTO, demonstrating competitive performance across accuracy and detailed video understanding metrics. Scores indicate the robustness of DyTO in generating contextually rich responses without finetuning.

fine-tuning, highlighting its robustness in zero-shot settings. DyTO’s ability to maintain high accuracy without fine-tuning underscores the strength of its adaptive framework in handling open-ended, real-world VQA tasks, achieving high scores across dimensions like correctness, contextual understanding, and temporal understanding.

5. Analysis

5.1. Scalability on Base Model Size

As shown in Table 1, scaling up model sizes significantly enhances the performance of DyTO across structured VQA tasks, demonstrating notable gains over performance with 7B base models. Specifically, with a 34B model, DyTO achieves a 7.2% accuracy increase on NExTQA [36], an 8.2% boost on EgoSchema [28], and a 7.7% improvement on MVBench [16]. These more substantial increments reflect the method’s enhanced capacity to capture and reason over complex spatial-temporal interactions, particularly in tasks requiring nuanced contextual comprehension.

Compared to other training-free approaches at 34B, DyTO consistently outperforms IG-VLM [11] and SlowFast-LLaVA [39] across benchmarks, with clear accuracy advantages on tasks like VideoMME [7]. This suggests that the proposed hierarchical clustering and bipartite token merging mechanisms are particularly effective in leveraging the additional model capacity to retain critical semantic information, even in the absence of fine-tuning.

In the open-ended VQA tasks shown in Table 2, our per-

formance with a 34B model also surpasses its 7B counterpart, reinforcing the scalability of our approach. For instance, DyTO achieves a 2.3-point average increase across correctness and contextual understanding dimensions of VCG benchmark, as well as improvements in detail orientation and temporal understanding. These gains further highlight the robustness of DyTO’s adaptive framework, which scales efficiently with model size to provide more contextually enriched and temporally accurate responses across various VQA benchmarks.

5.2. Performance w.r.t. Video Lengths

We analyze the performance of DyTO across varying video lengths to better demonstrate its robustness, particularly in maintaining accuracy with longer videos. As shown in Fig. 3 (Top), while the performance of DyTO does show some degradation as video length increases, the decline is significantly less pronounced compared to other methods like IG-VLM [11] and SlowFast-LLaVA [39]. This stability is largely attributed to DyTO’s adaptive, video-dependent dynamic token merging, which effectively prioritizes essential frames and contextual tokens, preserving key information even as the video sequence lengths.

In Fig. 3 (Bottom), we further examine the impact of increasing the number of sampled frames on performance with extended videos. As the number of sampled frames increases, DyTO’s accuracy on longer videos notably improves, demonstrating its capacity to capture detailed temporal and contextual information more effectively than

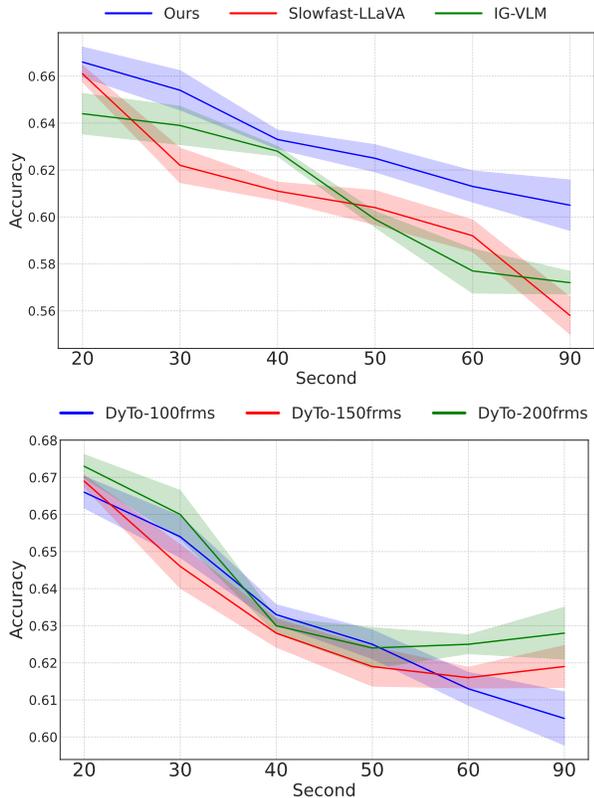


Figure 3. *Top*: Performance comparison of baseline method under various video lengths. *Bottom*: Effect of different input sampling lengths under various video lengths.

competing methods. This improvement highlights the advantage of DYTO’s hierarchical clustering and bipartite token merging, which dynamically adjust to sample frames that represent critical moments, ensuring more comprehensive and accurate video understanding.

This indicates a robust generalization across varying video lengths, reinforcing DYTO’s strength in managing diverse content scales without substantial accuracy loss or the need for fine-tuning.

5.3. Ablation of Clustering and Merging Module

The experimental results from the table demonstrate significant differences in performance when DYTO adopts different strategies. Overall, DyTo consistently outperforms both the baseline (SlowFast) and its ablated variants. Notably, on NEXTQA, DyTo w/o Clustering slightly surpasses the full DyTo model. A closer inspection reveals that, due to the relatively short clips in NEXTQA, the benefits of clustering become less pronounced, thus minimizing DyTo’s advantage. Nevertheless, in all other datasets, DyTo exhibits clear improvements over the baseline and ablation models, highlighting its stronger capability in handling more complex or longer video tasks. Furthermore, this performance edge persists regardless of the language model size (7B or

Method	LLM NEXTQA		EgoSchema IntentQA		VideoMME MV Bench	
	Size	acc.	acc.	acc.	acc.	acc.
SlowFast	7B	64.2	45.5	60.1	40.4	44.8
DyTo w/o Clustering	7B	65.6	47.8	61.4	42.3	44.8
DyTo w/o Token Merge	7B	64.9	45.6	60.5	41.2	44.6
DyTo	7B	65.7	48.6	61.6	42.7	45.2
SlowFast	34B	71.9	55.8	66.2	53.2	51.2
DyTo w/o Clustering	34B	73.2	55.8	66.8	52.5	52.7
DyTo w/o Token Merge	34B	72.2	56.0	66.3	51.7	51.8
DyTo	34B	72.9	56.8	67.3	53.4	52.9

Table 3. Ablation Study on clustering and token merging modules with parameter sizes ranging from 7B to 34B.

34B), demonstrating DyTo’s effectiveness and robustness in diverse settings.

5.4. Visualizations of Hierarchical Clustering

To further illustrate DYTO’s improved sampling and segmentation, Fig. 4 visualizes the clustering output from our method alongside other approaches, including IG-VLM [11] and SlowFast-LLaVA [39]. By mapping the video events using Hungarian matching, we demonstrate that DYTO has a clear advantage in accurately capturing and segmenting all critical events without omissions. Unlike IG-VLM and SlowFast-LLaVA, which may overlook certain events by failing to sample any key-frame from them—such as IG-VLM missing the first event of the video completely—DYTO reliably identifies each segment’s temporal boundaries, allowing for precise key-frame selection. Specifically, although SlowFast-LLaVA aims to mitigate this limitation by sampling 50 frames with coarse 4×4 tokens as the fast part input, shown as smaller red dots in Fig. 4, its main performance heavily depends on the 10 frames (the slow part) that carry the main temporal load [39]. Furthermore, Fig. 5, where each color represents different events segmented over time, highlights how DYTO’s clustering naturally organizes video content in a temporally coherent manner. This helps DYTO to achieve a higher fidelity in key-frame selection, making it better suited for scenarios requiring comprehensive, action-specific video understanding.

5.5. Qualitative Case Study

Fig. 9 provides a qualitative illustration of DYTO’s robust performance in a complex zero-shot video understanding scenario from the NEXTQA benchmark. Based models is LLaVA-NeXT-Image-7B. In this instance, DYTO accurately identifies and preserves critical visual details, such as the red toy on the stroller, the child’s attire, and the architectural background. Due to the simple uniform sampling and pooling method, IG-VLM and SlowFast-LLaVA both output the hallucination content (*carrying a backpack* and *holding a cup of cereal* respectively). Unlike these methods, DYTO captures nuanced visual cues and contextual elements. DYTO’s hierarchical clustering, which segments

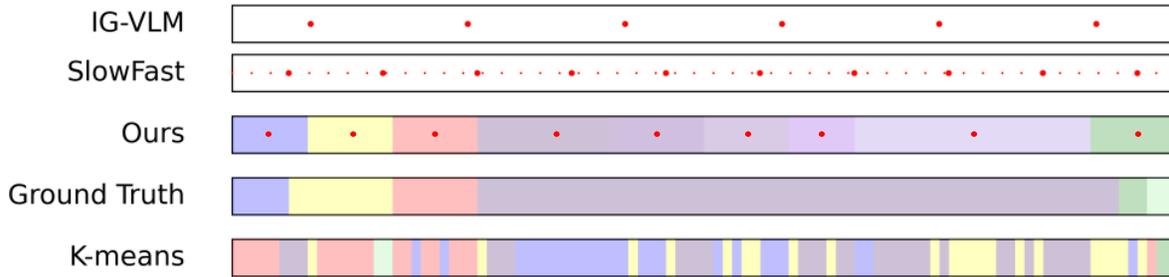


Figure 4. The sampling method and clustering module output visualization on a video. Our method offers more comprehensive video representation frames compared to other methods.



Figure 5. Clustering module output example from videos. Colors indicate different events in temporal order. The differences are clearly visible in the video clips.

Question: Describe what happens in this video in detail.

IG-VLM-7B: A woman and a young child are seen walking together through a city square. They **pass by a large stone archway with a plaque** on it. The woman is **carrying a backpack**. The child is wearing a **green jacket** and is **holding onto a pole**. They are surrounded by other people and buildings, indicating an urban setting. The video captures a moment of everyday life, with the woman and child being the main focus.

SlowFast-LLaVA-7B: In this video, a young child is seen walking on a street, **holding a cup of cereal**. The child is wearing a **green jacket** and **pink pants**. The setting appears to be a European city with historical architecture. There are other people and a horse-drawn carriage visible in the background. The child seems to be enjoying a casual, everyday moment, possibly on a family outing or a leisurely walk.

DyTo-7B: In this video, a young child is seen sitting in a **stroller**, which is located in a public space with a classical architectural style, possibly a **square** or a pedestrian area. The child is wearing a **green jacket** and **pink pants**. The stroller is positioned in front of a building with columns, which suggests a historical or significant location. The child appears to be looking upwards, possibly at the building or the sky. There are other people in the background, indicating a bustling environment. The overall content of the video is a candid moment of a child exploring their surroundings.

Figure 6. Example from NEXtQA benchmark. **red** part denotes the incorrect or hallucination content. **green** part denotes the correct object content. DyTo demonstrates significantly better performance compared to the other two methods.

video frames dynamically, alongside its bipartite merging strategy that avoids excessive compression that might lead to semantic loss, result in a more detailed and accurate scene understanding.

6. Conclusion and Future Work

In this paper, we introduced DyTo, a training-free framework designed for dynamic token merging in zero-shot video understanding. Through a novel combination of hierarchical frame selection and bipartite token merging, DyTo addresses the challenges of maintaining semantic richness while enhancing computational efficiency. Experimental results demonstrate that DyTo achieves state-of-the-art performance across multiple structured and open-ended VQA benchmarks, outperforming both fine-tuned and other training-free models. By dynamically adjusting token granularity based on frame content, DyTo successfully captures critical spatial-temporal details, offering a scalable solution that adapts to varying video lengths and complexities. DyTo not only sets a new standard in zero-shot video tasks but also paves the way for more efficient and contextually aware video understanding. Future work may explore extending DyTo to enhancing token adaptability for real-time applications, further pushing the boundaries of training-free video comprehension.

References

- [1] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, page 4, 2021. 1
- [2] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your ViT but faster. In *International Conference on Learning Representations*, 2023. 3
- [3] David Chen and William Dolan. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 190–200, Portland, Oregon, USA, 2011. Association for Computational Linguistics. 4, 5
- [4] Shimin Chen, Yitian Yuan, Shaoxiang Chen, Zequn Jie, and Lin Ma. Fewer tokens and fewer videos: Extending video understanding abilities in large vision-language models. *arXiv preprint arXiv:2406.08024*, 2024. 1
- [5] Zhaorun Chen, Zhuokai Zhao, Hongyin Luo, Huaxiu Yao, Bo Li, and Jiawei Zhou. Halc: Object hallucination reduction via adaptive focal-contrast decoding. *arXiv preprint arXiv:2403.00425*, 2024. 1
- [6] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, and Lidong Bing. VideoLLaMA 2: Advancing Spatial-Temporal Modeling and Audio Understanding in Video-LLMs, 2024. 1, 2, 5, 6
- [7] Chaoyou Fu, Yuhan Dai, Yondong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, Peixian Chen, Yanwei Li, Shaohui Lin, Sirui Zhao, Ke Li, Tong Xu, Xiawu Zheng, Enhong Chen, Rongrong Ji, and Xing Sun. Video-MME: The First-Ever Comprehensive Evaluation Benchmark of Multi-modal LLMs in Video Analysis, 2024. 4, 5, 6
- [8] Ziyang Jiang, Rui Meng, Xinyi Yang, Semih Yavuz, Yingbo Zhou, and Wenhui Chen. Vlm2vec: Training vision-language models for massive multimodal embedding tasks. *arXiv preprint arXiv:2410.05160*, 2024. 1
- [9] Peng Jin, Ryuichi Takano, Wancai Zhang, Xiaochun Cao, and Li Yuan. Chat-UniVi: Unified Visual Representation Empowers Large Language Models with Image and Video Understanding, 2024. 1, 2
- [10] Ehsan Kamaloo, Nouha Dziri, Charles LA Clarke, and Davood Rafiei. Evaluating open-domain question answering in the era of large language models. *arXiv preprint arXiv:2305.06984*, 2023. 1
- [11] Wonkyun Kim, Changin Choi, Wonseok Lee, and Wonjong Rhee. An Image Grid Can Be Worth a Video: Zero-shot Video Question Answering Using a VLM, 2024. 1, 2, 3, 5, 6, 7
- [12] 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. 5
- [13] Jiapeng Li, Ping Wei, Wenjuan Han, and Lifeng Fan. Intentqa: Context-aware video intent reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11963–11974, 2023. 4, 5
- [14] Kunchang Li, Yali Wang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Qiao. Uniformer: Unified transformer for efficient spatiotemporal representation learning, 2022. 3
- [15] KunChang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhui Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. VideoChat: Chat-Centric Video Understanding, 2024. 1, 2
- [16] Kunchang Li, Yali Wang, Yanan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. MVBench: A Comprehensive Multi-modal Video Understanding Benchmark, 2024. 4, 5, 6
- [17] Yuncheng Li, Yale Song, Liangliang Cao, Joel Tetreault, Larry Goldberg, Alejandro Jaimes, and Jiebo Luo. Tgif: A new dataset and benchmark on animated gif description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4641–4650, 2016. 4, 5
- [18] Yanwei Li, Chengyao Wang, and Jiaya Jia. LLaMA-VID: An Image is Worth 2 Tokens in Large Language Models, 2023. 1, 2
- [19] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models. 2024. 5, 6
- [20] Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-LLaVA: Learning Unified Visual Representation by Alignment Before Projection, 2023. 1, 2, 5, 6
- [21] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7083–7093, 2019. 1
- [22] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3202–3211, 2022. 1
- [23] Fan Ma, Xiaojie Jin, Heng Wang, Yuchen Xian, Jiashi Feng, and Yi Yang. Vista-LLaMA: Reliable Video Narrator via Equal Distance to Visual Tokens, 2023. 1, 2, 6
- [24] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023. 1, 4, 5
- [25] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-ChatGPT: Towards Detailed Video Understanding via Large Vision and Language Models, 2023. 6
- [26] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*, 2024. 1, 2, 6
- [27] Neelu Madan, Andreas Møgelmoose, Rajat Modi, Yogesh S Rawat, and Thomas B Moeslund. Foundation models for video understanding: A survey. *arXiv preprint arXiv:2405.03770*, 2024. 1
- [28] Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. Egoschema: A diagnostic benchmark for very long-

- form video language understanding. *Advances in Neural Information Processing Systems*, 36:46212–46244, 2023. 4, 5, 6
- [29] Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconnier, Sam Dodge, Bowen Zhang, Philipp Dufter, Dhruvi Shah, Xi-anzhi Du, Futang Peng, Floris Weers, et al. Mm1: Methods, analysis & insights from multimodal llm pre-training. *arXiv preprint arXiv:2403.09611*, 2024. 1
- [30] Zizheng Pan, Bohan Zhuang, Haoyu He, Jing Liu, and Jian-fei Cai. Less is more: Pay less attention in vision transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2035–2043, 2022. 3
- [31] Saquib Sarfraz, Vivek Sharma, and Rainer Stiefelhausen. Efficient parameter-free clustering using first neighbor relations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8934–8943, 2019. 5
- [32] Enxin Song, Wenhao Chai, Guanhong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, Yan Lu, Jenq-Neng Hwang, and Gaoang Wang. MovieChat: From Dense Token to Sparse Memory for Long Video Understanding, 2024. 2, 6
- [33] Yunlong Tang, Jing Bi, Siting Xu, Luchuan Song, Susan Liang, Teng Wang, Daoan Zhang, Jie An, Jingyang Lin, Rongyi Zhu, et al. Video understanding with large language models: A survey. *arXiv preprint arXiv:2312.17432*, 2023. 1
- [34] Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yanan He, Guo Chen, Baoqi Pei, Rongkun Zheng, Zun Wang, Yansong Shi, et al. Internvideo2: Scaling foundation models for multimodal video understanding. In *European Conference on Computer Vision*, pages 396–416. Springer, 2024. 5
- [35] Wenhao Wu. Freeva: Offline mllm as training-free video assistant. 2024. 1, 2, 3, 4
- [36] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9777–9786, 2021. 4, 5, 6
- [37] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5288–5296, 2016. 4, 5
- [38] 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, 2024. 2, 6
- [39] Mingze Xu, Mingfei Gao, Zhe Gan, Hong-You Chen, Zhengfeng Lai, Haiming Gang, Kai Kang, and Afshin Dehghan. SlowFast-LLaVA: A Strong Training-Free Baseline for Video Large Language Models, 2024. 1, 2, 3, 5, 6, 7
- [40] Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. Self-chained image-language model for video localization and question answering. *Advances in Neural Information Processing Systems*, 36:76749–76771, 2023. 5
- [41] Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. Self-chained image-language model for video localization and question answering. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 2
- [42] Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yuet-ing Zhuang, and Dacheng Tao. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9127–9134, 2019. 4, 5
- [43] Hang Zhang, Xin Li, and Lidong Bing. Video-LLaMA: An Instruction-tuned Audio-Visual Language Model for Video Understanding, 2023. 1, 2
- [44] Haotian Zhang, Haoxuan You, Philipp Dufter, Bowen Zhang, Chen Chen, Hong-You Chen, Tsu-Jui Fu, William Yang Wang, Shih-Fu Chang, Zhe Gan, et al. Ferret-v2: An improved baseline for referring and grounding with large language models. *arXiv preprint arXiv:2404.07973*, 2024. 1
- [45] Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. Llavavnext: A strong zero-shot video understanding model, 2024. 5, 6
- [46] Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. Llavavnext: A strong zero-shot video understanding model, 2024. 2
- [47] Yiming Zhang, Zhuokai Zhao, Zhaorun Chen, Zhili Feng, Zenghui Ding, and Yining Sun. Rankclip: Ranking-consistent language-image pretraining. *arXiv preprint arXiv:2404.09387*, 2024. 1
- [48] Zhuokai Zhao, Harish Palani, Tianyi Liu, Lena Evans, and Ruth Toner. Multimodal guidance network for missing-modality inference in content moderation. In *2024 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*, pages 1–4. IEEE, 2024. 1
- [49] Pengyuan Zhou, Lin Wang, Zhi Liu, Yanbin Hao, Pan Hui, Sasu Tarkoma, and Jussi Kangasharju. A survey on generative ai and llm for video generation, understanding, and streaming. *arXiv preprint arXiv:2404.16038*, 2024. 1

A. Time Consumption Experiment

we conducted the experiments using the same hardware specifications. The table below shows the time consumption for inference with 500 samples from EgoSchema using a single NVIDIA A100 GPU.

	SlowFast-LLaVA	DYTo
Dataset	Egoschema	
Model	LLaVA-NeXT-34B	
Input	50 + 10 frames	100 frames
Merge Strategy	Pooling	Dynamic token merging
Device	1 Nvidia A100 GPU	
Time Consumption	5.74 s/item	6.22 s/item

Table 4. time consumption

As shown in the table, the difference in time consumption is negligible. Although our method is slightly slower than SlowFast, we think it may be attributed to hardware optimizations or variance.

B. Implementation details

All evaluations could be conducted on a single Nvidia A100 80G graphics card. To accelerate inference, we use a Linux server equipped with 8 Nvidia A100 80G cards. We carry out our evaluation across three model series and five model size. The weights for these models are available on Huggingface¹²³. We implement rotary position embedding (RoPE) and apply a scaling factor of 2, extending the original context length from 4096 to 8192 tokens.

¹<https://huggingface.co/collections/liuhaotian/llava-16-65b9e40155f60fd046a5ccf2>

²<https://huggingface.co/OpenGVLab/InternVL2>

³<https://huggingface.co/Qwen/Qwen-VL-Chat>

C. Visualizations of Dynamic Bipartite Merging

To help understand dynamic token merging effectively, we provide the visualization comparing our method with the pooling method. As shown in Fig. 7, the proposed approach effectively maintains the object’s actions while making every effort to prevent the disruption of the original spatial information. We set the constant merge ratio of $r=288$ to enable a convenient and fair comparison with the pooling method, while r is a dynamic integer that varies based on the number of clusters in DYTO. It is important to emphasize that our proposed token merging method operates without the need for any labels. To create the visualizations in Figure 7, we follow each final merged token back to its original input patches. For each token, we color its corresponding input patches, referred to as "Patchfied," using the average color of that region. To ensure that different tokens are distinguishable, we assign each token a random border color. It’s important to note that tokens do not necessarily correspond to contiguous input regions. The only spatial information comes from the position encodings.



Figure 7. Results of merging on video frames. From top to bottom is original images, after pooling, patchfied, and the proposed bipartite merging. Patchfied means

D. More Qualitative Cases



Question: Describe what happens in this video in detail.

IG-VLM-7B: A man is seen sitting in front of a laptop, watching a movie. The laptop is placed on a wooden desk. The man is wearing glasses and a blue shirt. The movie he is watching appears to be a romantic drama, as suggested by the content displayed on the laptop screen. The overall context of the video seems to be a casual, indoor setting, possibly a home or office environment, where the man is enjoying some leisure time by watching a movie.

SlowFast-LLaVA-7B: In this video, a man is sitting at a desk with a laptop in front of him. He appears to be engaged in a video call with another person, who is not visible in the image. The man is wearing a green shirt and plaid pants. On the desk, there is a glass of water and a sandwich. The setting suggests a casual, indoor environment, possibly a home or office. The man seems to be in a relaxed posture, indicating a comfortable and informal interaction. The context of the video suggests a personal or professional video call, possibly for work or personal reasons.

DyTo-7B: In this video, a man is seen sitting at a desk with a laptop, engaging in an activity that involves a teddy bear. The scene appears to be set in a room with a window, and there are other objects on the desk, such as a cup and a bottle. The man seems to be focused on the laptop screen, possibly working or browsing the internet. The presence of the teddy bear suggests a personal or playful element to the scene. The overall content of the video seems to be a casual, everyday moment captured in a home setting.

Figure 8. Example from MVBench benchmark. red part denotes the incorrect or hallucination content. green part denotes the correct object content.



Question: Describe what happens in this video in detail.

IG-VLM-7B: A person is shown sewing a garment, specifically a pair of **black pants**. The sequence of images captures the process of sewing, from the initial preparation of the fabric and thread, to the actual stitching, and finally the completion of the garment. The person is using a sewing machine and various sewing tools, such as **scissors** and a **needle**. The video provides a detailed look at the craftsmanship involved in creating clothing.

SlowFast-LLaVA-7B: A person is seen sewing a garment, specifically a pair of **pants**. The sequence of key frames captures the process of sewing, starting with the person **cutting the fabric**, then stitching the pieces together, and finally adjusting the fit of the pants. The video provides a detailed look at the craftsmanship involved in creating clothing, showcasing the skill and precision required in sewing.

🐛DyTo-7B: A person is shown in the process of sewing a garment. The sequence of key frames captures the following events:

1. The person is sitting at a table with **fabric**.
2. They are holding a piece of fabric in their hands, preparing to sew it.
3. The person is using a sewing machine, which is visible in the image.
4. They are sewing the fabric onto another piece of fabric, which is also visible.
5. The person is holding a pair of **scissors**, likely for **cutting the fabric**.
6. The person is using a **measuring tape**, which is also visible in the image.
7. The person is holding a **cup**, possibly containing a beverage.
8. The person is wearing a **tattoo on their arm**.

Throughout the video, the focus is on the hands-on process of sewing, with the person demonstrating various tools and techniques. The setting appears to be a workspace or a crafting area, with the person engaging in a creative activity.

Figure 9. Example from Egoschema benchmark. **green** part denotes the correct object content. **blue** part denotes the missing content in IG-VLM and SlowFast-LLaVA answers. DYTO can effectively capture the events in the video.