

Learning Compact Video Representations for Efficient Long-form Video Understanding in Large Multimodal Models

Yuxiao Chen, Jue Wang, Zhikang Zhang, Jingru Yi, Xu Zhang, Yang Zou
Zhaowei Cai, Jianbo Yuan, Xinyu Li, Hao Yang, Davide Modolo
Amazon AGI

Abstract

With recent advancements in video backbone architectures, combined with the remarkable achievements of large language models (LLMs), the analysis of long-form videos spanning tens of minutes has become both feasible and increasingly prevalent. However, the inherently redundant nature of video sequences poses significant challenges for contemporary state-of-the-art models. These challenges stem from two primary aspects: 1) efficiently incorporating a larger number of frames within memory constraints, and 2) extracting discriminative information from the vast volume of input data. In this paper, we introduce a novel end-to-end schema for long-form video understanding, which includes an information-density-based adaptive video sampler (AVS) and an autoencoder-based spatiotemporal video compressor (SVC) integrated with a multimodal large language model (MLLM). Our proposed system offers two major advantages: it adaptively and effectively captures essential information from video sequences of varying durations, and it achieves high compression rates while preserving crucial discriminative information. The proposed framework demonstrates promising performance across various benchmarks, excelling in both long-form video understanding tasks and standard video understanding benchmarks. These results underscore the versatility and efficacy of our approach, particularly in managing the complexities of prolonged video sequences.

1. Introduction

"The whole is greater than the sum of its parts."

— Aristotle

Long-form video understanding has emerged as a hot topic in the research community, presenting significant challenges due to its complexity and dynamic nature, and remains an unsolved problem despite ongoing advancements. Conventional approaches [11, 15, 40, 43, 47] often focus

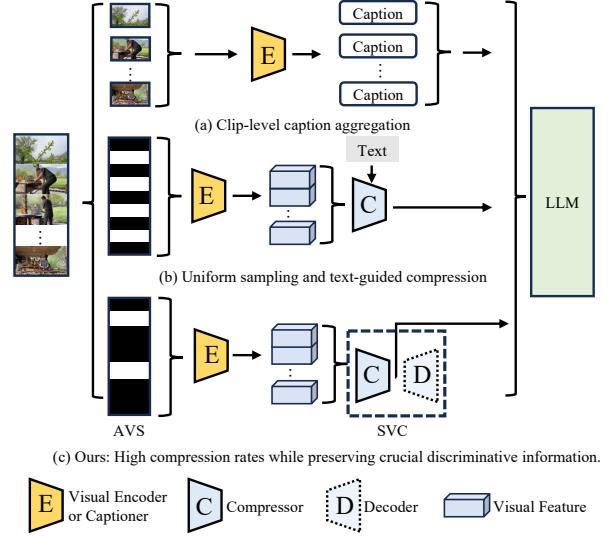


Figure 1. An overview of our method and previous works, showing the different way of modeling long-form video with LLM. Specifically, (a) interpret clip into clip-level caption and aggregate via LLM in the linguistic space, (b) uniformly sample the video frame and leverage paired text to compress video tokens, and (c) our proposed method that leverage adaptive video sampler (AVS) and autoencoder based spatiotemporal video compressor (SVC).

on classification tasks, where the goal is to map an entire video to one or a few predefined ground truth labels. These methods typically treat the video as a single, unified entity, aiming to assign a label that best represents the overall content. Although effective in some contexts [15, 40, 47], these approaches are insufficient to understand long-form video, as they do not capture the complexity and contextual relationships unfolded across multiple scenes over time.

Recently, multimodal large language models (MLLMs) [18, 26, 26, 29, 37, 38, 49] have made a significant impact on the field of video understanding, showcasing remarkable emergent capabilities across a variety of video-related tasks, such as perception, common

sense reasoning, and open dialogue [20, 30, 33, 48, 54]. However, many of these models face challenges in effectively learning and modeling long-form video. This difficulty arises from the redundant nature of video sequences, which often results in an overwhelming number of visual tokens, leading to substantial computational overhead in LLMs. RoPe [39] provides a solution to handle the extended context during inference. However, the complexity of a long-form video sequence requires a more effective encoding solution during training.

To address these challenges, one successful attempt [12, 25, 45, 56], illustrated in Figure 1 (a), is to represent the long-video hierarchy with the natural language, which segments the long video sequences into several short clips and produces video segment captions. These short captions are subsequently fed into the LLMs for inference. However, these models interpret video clips into natural languages in the early stage, leading to the loss of much low-level visual information. Additionally, aggregated hallucinations from segment-level captions may result in poor generalizability and performance across various tasks and models. Another direction [23, 38] focuses on reducing the large number of tokens in long-form videos by introducing a learnable compression or token merging module, or a memory bank linked to the LLM, as shown in Figure 1 (b). However, scaling these methods is challenging, as they require a large volume of video-text pairs, with the linguistic information serving as guidance for training the compressor. Finally, parameter-free pooling operations (such as average pooling) have been proven effective [18, 26, 26, 29, 37, 38, 49], but information loss remains unavoidable and is even more severe in long-form videos, where the diversity within video sequences is greater than in short videos.

In this paper, we introduce a comprehensive schema for long-form video understanding, covering the entire pipeline, from sampling to compression to high-level interpretation. This system includes: 1) an Adaptive Video Sampler (AVS) that selects video frames based on information density, 2) an Auto-Encoder based [36] Spatiotemporal Video Compressor (SVC) that can be effectively trained with video-only data, and 3) a MLLM that integrates seamlessly with both the compressor and sampler. The AVS works jointly with SVC delivering 64-time compression to save the token budget for MLLMs. An overview of our proposed system is illustrated in Figure 1 (c). Our approach significantly reduces the token budget by 64 times, enabling MLLMs to process hours-long video sequences in their entirety. The proposed AVS and SVC modules ensure a sufficient capture of discriminative information while considerably expanding the perception field of the MLLM in long-form videos.

We conducted extensive experiments in a diverse range of tasks and problem settings to evaluate the proposed

method. Our approach significantly outperforms state-of-the-art methods on various video benchmark downstream tasks while demonstrating superior efficiency by using substantially fewer visual tokens and reducing computational overhead. Notably, our method surpasses Llava-OV [18] by 2.6% and 3.3% on the EgoSchema and PercepTest tasks, respectively, while utilizing 80% fewer visual tokens. Our contributions are as follows.

- We proposed a novel long-form video understanding schema to equip MLLM, which includes two novel components: Adaptive Video Sampler (AVS) which samples frames based on the information density, and Spatio-temporal Video Compressor (SVC) that is an AE-based video compressor.
- The proposed AVS and SVC work jointly to deliver 64-time compression ratio, which significantly save the token budget in MLLM while preserving the discriminative information of the long-form videos.
- Compared to baseline and competitive counterparts, we demonstrate the effectiveness of the proposed elements and achieve performance on par with the SoTA methods in various video understanding benchmarks, while only utilizing 20% video tokens compared to the previous SoTA method [18].

2. Related Work

2.1. MLLMs in Long-form videos

In contrast to conventional long-form video understanding tasks that classify entire long video as one or a few ground truth labels [11, 15, 40, 43, 47], Multi-modal Large Language Models (MLLMs) interpret long-form video within a well-defined linguistic space, thus expanding the scope of video understanding to include tasks such as contextual reasoning, open-ended dialogue about video content, and complex question-answering across long-video sequences. However, the high computational cost of transformer-based LLMs poses a significant barrier to effectively processing entire sequences of long-form videos within memory constraints. To address this issue, one solution is to divide the entire long-form video into short clips and generate clip-level captions, which will be hierarchically aggregated as long-form video representation via LLM [12, 25, 45, 56]. However, these models often lose significant low-level visual information, and the accumulated hallucinations from segment-level captions can lead to poor generalizability and reduced performance across various tasks and models. To this end, an efficient and effective video encoding scheme is required in long-form video understanding tasks.

2.2. Long-form Video Encoding

Token Reduction Token reduction is one of the most effective methods to relieve computational burden. It can be

further categorized as 1) *Token Selection*, 2) *Token Merging* and 3) *Token Compression*.

Token Selection: Token selection is a common strategy for improving model efficiency, utilizing lightweight modules to retain only the most "useful" tokens while discarding those considered less essential. Typical works, such as STTS [42], AdaViT [31] and similar methods [24, 31, 34, 42] introduce a selection mechanism to assign importance scores to each token and selects the top-K important ones. However, it is impractical for a lightweight selection module to select discriminative information from a large volume of long-form video data without risking a loss of contextual information.

Token Merging: Originally proposed by [2], token merging is designed to enhance throughput in vision transformers without the need for further training. The merging technique splits visual tokens into two equal groups; for each token on the edge of one group, it identifies the closest matching token in the other group, merges them using average pooling, and then recombines the groups. Although this method is effective and straightforward, its use has been mainly limited to the image and short video domains [1, 21, 35]. This is due to the high diversity in long-form videos, which often leads to mismatches in a similarity-based merging system. More practically, a sliding window-based averaging pooling cubic is widely adapted in recent long-form video understanding works [6, 18, 26, 29, 49]. In particular, Video-ChatGPT [26] and SlowFast-LLaVA [49] employ two distinct pooling strategies across spatial and temporal dimensions to capture object and motion information, respectively. However, these pooling strategies treat each video frame equally, which may result in information distortion when two dissimilar candidates are combined. Effectively reducing redundant visual tokens while preserving key information remains a challenge and further exploration is needed.

Token Compression: Recent works [23, 37, 38, 46, 53] have successfully leveraged a compressor to reduce the number of visual tokens while preserving useful information in various applications, such as classification, generation, and general understanding with LLMs. For example, Wiles et al. [46] leverage the VQ-VAE encoder-decoder model [41] to apply augmentations directly on the latent codes of compressed video and use them for classification tasks. Yu et al. propose MAGViT-v2 [53], which generates concise and expressive tokens for both videos and images utilizing a shared VQ-VAE codebook. In the field of long-form video understanding, MovieChat [38] introduces a recurrent memory bank to consolidate excessive video tokens, while LLaMA-VID [23] uses cross attention to represent each video frame as two tokens. Furthermore, LongVU [37] utilizes DINOv2 [32] features to eliminate redundant frames that exhibit high similarity. These works either rely on ex-

ternal prior or well-structured video and text pairs, which may easily introduce bias or prevent the algorithm scaling up. In this work, we argue that the token reduction mechanism should be implemented as a holistic pipeline that jointly considers video sampling and compression problems as a whole. To this end, we propose a simple yet effective adaptive video sampler as well as a video compressor that delivers a 64-time compression ratio over long-form video while preserving discriminative information.

Video Backbone Recently, Mamba [10] and its huge family [4, 28, 43, 51] received promising results in various long context modeling tasks. Specifically, it utilizes a state-space sequence model with a structured convolutional kernel to efficiently capture long-context information, achieving linear complexity. However, stringent initialization conditions make it difficult for Mamba to be trained in large scale with MLLMs [28]. As a result, we follow previous works [18, 26, 29, 37, 38, 49] by using ViT [7] as the video backbone. The exploration of Mamba with MLLMs will be left for future work.

3. Methodology

3.1. Preliminaries

MLLM For Long-form Video Understanding

Multimodal Large Language Models (MLLMs) process vision data using LLMs to interpret visual information extracted from inputs alongside task-specific text prompts (e.g., visual questions about visual content), ultimately generating the corresponding responses [6, 23, 27].

To be specific, given an input video V as an example, a set of frames, denoted as X , are first *uniformly* sampled from it. A pretrained visual encoder \mathcal{E} then extracts visual token embeddings, denoted as f , from these sampled frames:

$$f = \mathcal{E}(X), \quad f \in \mathcal{R}^{T \times H \times W \times C} \quad (1)$$

where T represents the frame number, H , and W are height and width of the extracted feature tensor, respectively.

The LLM subsequently processes the visual tokens with the text question (Q) to generate the answer A .

$$A = LLM(P(f), Q) \quad (2)$$

where P is a projector, such as the MLP module, that maps visual token features into the input space of the LLM.

However, in the context of long-form video, the volume of f increases significantly, which impacts the capacity of MLLM¹ in terms of both efficiency and effectiveness. Additionally, uniform sampling combined with standard average pooling is suboptimal for long-form video, as redundant and irrelevant frames may consume the limited token

¹The complexity of transformer-based LLMs is quadratic in relation to input length.

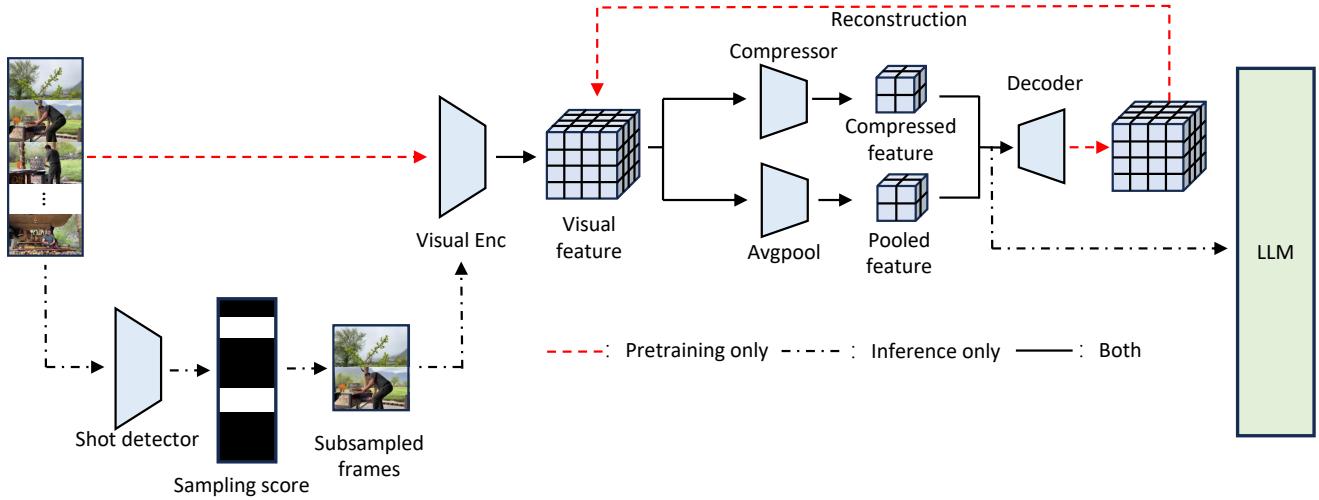


Figure 2. Overview of the proposed method.

budget of LLMs, ultimately reducing their ability to capture complex long-term dependencies. Motivated by these, we propose a long-form video understanding framework in this work, offering solutions and best practices for the research community.

3.2. Proposed Framework

In this paper, we propose a novel long-form video understanding framework, addressing the aforementioned challenges, from frame loading strategy to video token compression technique. Our framework consists of an Adaptive Video Sampler (AVS), a ViT-based video encoder, an autoencoder-based Spatiotemporal Video Compressor (SVC) and the recent popular LLM, QWen2 [50]. An illustration is provided in Figure 2.

3.2.1. Adaptive Video Sampling

In long-form videos, the discriminative information is not evenly distributed. Given the tight token budget of LLMs, it is important to carefully pick visual tokens from the entire sequence. To this end, we propose an information-density-based frame sampler that can adaptively select informative video frames from a long video sequence. We follow the decomposition rule of a movie, borrowing the concept of chapter, scene, and shot in the movie. It is assumed that each long-video sequence consists of several continuous information tubelets, with intra-tubelet information being generally homogeneous, while inter-tubelet data distributions vary significantly. As a result, the gradient of information shifts is crucial for identifying where discriminative information lies in the long-form video. In this work, we employ a shot boundary detection module as a core component of our adaptive frame sampler to capture dynamic moments that contain key information.

Specifically, we first feed the entire video V to a shot boundary detector S , which outputs a confidence score that indicates the likelihood of a content change with respect to each frame:

$$S = \mathcal{S}_D(V), \quad S = \{s_1, s_2, \dots, s_N\} \quad (3)$$

where s_i is the confidence score for the i -th frame, and N is the number of frames.

We then apply the non-maximum suppression (NMS) scheme to filter out redundant detections:

$$S' = \text{NMS}(S) \quad (4)$$

Next, we sample frames with the top- k highest confidence scores and sort them temporally for video understanding:

$$K = \text{Sort}(\text{Top}_k(S')) \quad (5)$$

where K represent the index of sampled frames

3.2.2. Autoencoder-based Video Compressor

With the smart sampling strategy, the perception field of LLMs is still constrained by the massive number of visual tokens per frame. In the area of image and short-form videos, average pooling is widely used as an effective way to reduce redundant tokens [6, 18, 26, 29, 49]. However, it is argued that the average pooling is not ideal in long-form video understanding, as the high diversity among sparsely sampled frames can lead to information distortion when forced into aggregation. Unlike previous works that rely on knowledge priors or well-aligned video-text pairs to compress video tokens, we propose an autoencoder-based video compressor in this work. This compressor condenses raw video feature representations into a compact latent space,

reducing the number of visual tokens while retaining essential information. The autoencoder architecture ensures the scalability of the compressor, allowing it to be trained using video data alone. In [16], it is pointed out that the average pooling demonstrates superior performance than the convolution-based compressor. In this work, we show that the convolution-based compressor with autoencoder pre-training significantly outperform the average pooling.

To be specific, the compressor model \mathcal{C} removes redundant information among tokens by compressing raw video features f to a compact latent space,

$$h = \mathcal{C}(f), \quad h \in \mathcal{R}^{t \times h \times w \times c} \quad (6)$$

where t, h, w are the temporal length, height, and width of the feature tensor in the latent space, $t \leq T$, $h \leq H$, and $w \leq W$.

The decoder \mathcal{D} model then reconstructs the raw feature representation from h :

$$\hat{f} = \mathcal{D}(h), \quad \hat{f} \in \mathcal{R}^{T \times H \times W \times C} \quad (7)$$

The compressor \mathcal{C} and the video decoder \mathcal{D} are trained together to minimize the difference between \hat{f} and f using the mean absolute loss:

$$\mathcal{L}_{rec} = |f - \hat{f}| = |\mathcal{E}(X) - \mathcal{D}(\mathcal{C}(\mathcal{E}(X)))| \quad (8)$$

Note that when optimizing \mathcal{C} and \mathcal{D} to minimize \mathcal{L}_{rec} , the visual encoder \mathcal{E} remains fixed. The loss forces the encoder to retrain all necessary information in the compressed representation during the compression process.

The reconstruction loss also enables us to pre-train the compressor \mathcal{C} using only video data, thereby learning valuable prior knowledge. In contrast, previous methods train the compressor alongside LLMs for next-text-token prediction tasks, which require costly visual-text pairs for training. The limited scale of such data may constrain the compressor's performance.

After pretraining, we leverage the compressor \mathcal{C} with the visual encoder \mathcal{E} to extract compressed video representation from the video, which is then combined with a text question Q and fed into the LLM for response inference.

$$A = LLM(P(\mathcal{C}(\mathcal{E}(X))), Q) \quad (9)$$

Residual Latent Space Constraint: We empirically observe that aligning the compressor pretrained using Equation 8 with LLMs presents significant challenges. This difficulty stems from the lack of constraints within the compact latent space. Consequently, the learned compressor has limited generalization capabilities and may encode unseen data into gaps or "holes" within the latent space, leading to a loss of meaningful representation. To mitigate this issue, we propose incorporating the average pooled feature of X

with h as constraints, redefining the latent feature h and reconstruction loss \mathcal{L}_{rec} as follows:

$$h = \mathcal{C}(f) + avgpool_{3D}(X) \quad (10)$$

$$\mathcal{L}_{rec} = |\mathcal{E}(X) - \mathcal{D}(\mathcal{C}(\mathcal{E}(X)) + Avgpool_{3D}(X))| \quad (11)$$

With the proposed constraint, the compressor focuses on learning the residual (lost information) from the average pooling process, thus reducing the learning complexity of the autoencoder while implicitly ensuring that the learned feature $\mathcal{C}(f)$ aligns with the same space as $avgpool_{3D}(X)$. Additionally, compared to Variational Autoencoders (VAEs), which enforce the latent feature space to follow a Gaussian distribution, our constraint achieves superior performance by eliminating nondeterministic behavior in the latent space, which can increase the challenges of learning for the autoencoder (see experiments for more results).

Compression Ratio: Compared with the setting where the output features from the visual encoder are used directly as input to the LLM (see equation 2), the compressor reduces the token number by factors of $\frac{T}{t}$, $\frac{H}{h}$, and $\frac{W}{w}$ for the temporal, height, and width dimensions, respectively. This results in a total compression ratio of $\frac{T}{t} \times \frac{H}{h} \times \frac{W}{w}$ times.

Light-weight 3D convolutional kernel: We implement \mathcal{C} using cascaded convolutional residual blocks. *We chose a convolution-based compressor for its simplicity and superior computational efficiency compared to transformer-based models such as Perceiver. Additionally, convolution operations effectively utilize the inductive bias that assumes that local visual tokens have high redundancy.* To reduce the parameter size, we decompose the 3D convolution operation into a 2D spatial convolution and a 1D temporal convolution. Additionally, we apply a channel-wise bottleneck mechanism when the 2D spatial convolution maps the feature to a lower dimension than the original, while the temporal convolution restores the feature to its original dimension. The downsampling is achieved through the convolutional strides (s_h, s_w) for the 2D convolution and s_t for the 1D convolution.

3.3. Implementation Details

We adopt a four-stage training recipe. (1) In the first stage, we train a VAE-based video compressor with supervision of construction loss. (2) Next, a projector between the video features and LLM is trained with our filtered ShutterStock data ($\sim 3M$) to align the visual features from visual space to language space. (3) In the third step, we unfreeze the LLM weights and further jointly train the LLM and projector with $\sim 8M$ data which is collected from subset of ShutterStock, Ego4D [9], Breakfast [16], AVA [17], Vatex [44],

Something-somethingv2 [8], Kinetics [3] datasets. (4) In the last step, we finetune the model with our supervised finetuning (SFT) dataset which is collected from NextQA [48], CLEVRER [52], PerceptionTest [33], and Egoschema [30]. The SFT data adds up to 1M scale.

4. Experimental Results

4.1. Benchmark and Metric

We conduct experiments on general general video understanding benchmarks covering a wide range of video lengths and video types, including: **PerceptionTest** [33] a multimodal benchmark designed to show perceptually interesting situations and defines multiple tasks. Following [33] we report the accuracy of multichoice questions (MCQ). **ActivityNet-QA** [54] contains 58,000 human-annotated QA pairs of 5,800 videos derived from the popular ActivityNet dataset. The dataset provides a benchmark for testing the performance of VideoQA models on spatio-temporal reasoning. Following [45], we report both accuracy and LLM based answer matching scores for above datasets. The **MVBench** [20] is a comprehensive Multi-modal Video understanding benchmark with videos ranging from 5s to 35s. MVBench covers 20 challenging video tasks that require a broad spectrum of temporal skills, ranging from perception to cognition. We return the average score on the downstream tracks as suggested by [20]. The **NExTQA** [48] is a VideoQA benchmark that focuses on the reasoning of causal and temporal actions reasoning and object interactions understanding in daily activities. It contains 48K multi-choice questions with average duration of 44s. For these datasets, we report the multi-choice accuracy following previous practice [45]. For long video understanding, we choose **EgoSchema** [30] which is derived from Ego4D [9], consists of over 5000 human-curated multiple choice question answer pairs, with videos of three-minute duration. And **MLVU** [59] which is designed for very Long Video Understanding (LVU) tasks. It is constructed from a wide variety of long videos, with lengths ranging from 3 minutes to 2 hours. We report the mean accuracy of multi-choice accuracy on all downstream tasks as suggested by benchmark.

4.2. SoTA Comparisons

Comparison with SoTA long-form video understanding methods. First, we benchmark our method against previous long-form video understanding works on the EgoSchema, NextQA, and ActivityNetQA datasets, with results presented in the upper part of the Table 1. Our method demonstrates significant performance improvements on all three benchmarks. Notably, our method outperforms VideoAgent [45], LLoVi [55], and VideoINSTA [25] by 8.6% and 4.9% on the EgoSchema validation set and testing

set. These multi-stage methods, which segment videos into short clips, extract clip captions, and aggregate objects through LLM, suffer from error propagation and information loss during multi-stage aggregation. In contrast, our approach learns to extract information end-to-end from videos. The Adaptive Video Sampler (AVS) efficiently filters out redundant temporal information, while our Spatiotemporal Video Compressor with AE maximally preserves of long-term spatiotemporal information, resulting in superior performance. Our method also outperforms LLaMA-VID and Movie-Chat on the ActivityNet-QA dataset by 4.8%. The two methods rely on conventional, complex pooling-based token compression strategies, such as LLaMA-VID’s Q-former-based compression and Movie-Chat’s memory-buffer approach for long-term information aggregation. In contrast, our method excels by providing a more lightweight yet effective solution, with less information loss.

Comparison with SoTA MLLMs. Furthermore, we evaluate our method against state-of-the-art MLLMs across both general and long-form video understanding benchmarks in the lower part of the Table 1. For fair comparison, we focus on models using 7B LLM backbones. Although we had access to only a subset of the Supervised Fine-Tuning (SFT) data used in previous works [18], our model still achieves impressive results across all benchmarks. Specifically, our approach achieves state-of-the-art performance on three out of six benchmarks and comparable results on the remaining ones, demonstrating the effectiveness of our method. Notably, our model processes EgoSchema and Perception-Test using an average of just 1,440 visual tokens, substantially fewer than the approximately 6,000 tokens required by LLava-OV [18]. This significant reduction in token usage demonstrates how our elegant design, incorporating the proposed sampler and compression module, achieves both computational efficiency and strong performance.

4.3. Ablation Study

To validate our design choices and demonstrate the effectiveness of the proposed modules, we conducted ablation studies across four established long-form video benchmarks [20, 30, 48, 59].

Video frame samplers. We evaluated our adaptive sampler against uniform sampling using an equal number of sampled frames, with results shown in Table 2. For videos without shot changes [20, 48], both sampling methods perform similarly; for example, on NExT-QA, where only 1% of the data contains shot changes, our adaptive sampling showed a modest improvement of 0.1%. This comparable performance is expected, as frames within the same scene tend to be visually similar. However, the adaptive sampler demonstrates notable advantages on longer videos containing shot changes, achieving a 1% improvement on MLVU [59]. Fur-

Model	EgoSchema		NextQA mc	ActivityNetQA acc	MLVU avg	MVbench test	PerceptionTest val
	fullset	subset					
VideoChat [19]	-	-	-	26.5	-	-	-
LLaMA Adapter [19]	-	-	-	34.2	-	-	-
Video-ChatGPT [19] (ACL'24)	-	-	-	35.2	-	-	-
LLaMA-VID [22] (ECCV'24)	-	-	-	47.5	-	-	-
Movie.Chat [38] (CVPR'24)	-	-	-	45.7	-	-	-
VideoAgent [45] (ECCV'24)	54.1	60.2	71.3	-	-	-	-
LLoVi [55] (EMNLP'24)	52.2	58.8	73.8	-	-	-	-
VideoINSTA [25] (EMNLP'24)	-	<u>65.0</u>	72.3	-	-	-	-
ShareGPT4V [5]	-	-	-	-	46.4	51.2	-
LLaVA-NeXT-Video-7B [58]	43.9	-	-	-	-	33.7	-
LongVA-7B [57]	-	-	68.3	50.0	56.3	-	-
LLava-OneVision-7B [18]	60.1	-	79.4	56.6	64.7	56.7	57.1
Ours	62.7	69.6	<u>78.7</u>	<u>52.7</u>	<u>59.6</u>	<u>58.8</u>	60.4

Table 1. Comparison with state-of-the-art MLLM-based methods on both long-form video understanding and general video understanding tasks. The bold and underline texts represent the best and the second best performances, respectively.

Sampler	NextQA MC	MVbench test	MLVU (m-avg)						
			Avg.	Anomaly Det.	Counting	Ego.	Needle	Ordering	PlotQA
Uniform	77.0	50.0	51.9	0.53	0.35	0.50	0.51	0.47	0.52
Ours (AVS)	77.1	50.2	52.6	0.54	0.37	0.51	0.52	0.49	0.53

Table 2. Experimental results of our method on NextQA, MLVU and MVbench with different video frame sampling strategies.

Compressor	Ratio	EgoS.	MLVU	MVbench
Percevier	64 \times	63.0	53.7	49.8
Avgpool-3D	2 \times 4 \times 4	63.5	51.5	48.7
Avgpool-3D	4 \times 4 \times 4	61.4	51.3	46.8
Ours (SVC)	4 \times 4 \times 4	63.6	51.9	50.0

Table 3. Comparison between our proposed AE based VAC and different video compressors.

Ratio (T \times W \times H)	EgoS.	MLVU	MVbench	NextQA
2 \times 2 \times 2	65.1	55.0	51.3	77.1
4 \times 4 \times 4	63.6	51.9	50.0	77.0
1 \times 8 \times 8	60.0	50.1	44.2	73.5
4 \times 8 \times 8	62.0	48.7	46.6	75.5
8 \times 4 \times 4	62.3	48.0	43.3	73.3

Table 4. Experimental results of our method with different video compression ratio.

ther analysis of MLVU task-specific performance reveals significant improvements in tasks that heavily rely on key frames, such as anomaly detection, plot understanding, and needle-in-the-haystack scenarios. These results validate both the effectiveness of our proposed AVS and its robustness, maintaining strong performance even in videos without shot changes. In Figure 3, we demonstrate an example showing the different sampling results between uniform sampling and our AVS. As can be seen from the Fig 3, uniform sampling picks many similar video frames and waste

the token budget while our AVS can pick the key frames that are not overlapped with each other and successfully localize the frame showing a women holding a phone.

Video on Compressor. We first conduct the ablation study on our SVC design. We sample 32 frames as input and utilize the proposed SVC with a default compression ratio of 64 \times (4 \times 4 \times 4 for temporal, width, and height dimensions) across all ablations unless otherwise specified.

Compressor Architecture. We evaluated our proposed SVC method against two baseline approaches: average pooling and Perceiver-based down sampling [13], with results presented in Table 3. Our AE-based method demonstrates superior performance compared to both baselines at the same compression ratio (64 \times), validating the effectiveness of our SVC approach for video information compression. Notably, our method outperforms average pooling even with half number of tokens which suggests that simply increasing the number of tokens doesn't necessarily improve video representation. Rather, the key to effective video representation lies in eliminating redundant information while preserving essential spatial-temporal features. While the Perceiver approach achieves comparable performance to our method, it requires substantially more memory due to its cross-attention operations between queries and the numerous video tokens generated by the visual encoder.

Compression ratio. In Table 4, we analyze how different compression ratios affect benchmark performance. As expected, less aggressive compression yields better perfor-

Question: Who is the female chatting with in the video call from her home living room?



Figure 3. Examples of sampled frames using uniform sampling (top) compared to our AVS (bottom). AVS successfully locate the key frame to answer the question.

Method	EgoS.	MLVU	MVbench	NextQA
Random Init	59.2	49.9	47.2	72.8
Pre-train	63.6	51.9	50.0	77.0

Table 5. Experimental results of our method with pre-trained SVC and random initialized counterpart.

Constrain	EgoS.	MLVU	MVBench	NExT-QA
AE	Does not converge			
VAE	63.7	48.4	44.6	74.7
APool+Res	63.6	51.9	50.0	77.1

Table 6. The effectiveness of different constraints in the autoencoder pretraining of SVC.

mance overall. However, when examining more aggressive compression scenarios, we discovered that spatial compression has a smaller negative impact on model performance compared to temporal compression. This difference might be attributed to the auto-encoder loss function’s effectiveness in preserving spatial semantics during information compression. Furthermore, our experiments reveal that balanced spatial-temporal compression ratios perform better than the spatially or temporally heavy pooling approaches under the same total compression ratio. For instance, a balanced $4 \times 4 \times 4$ compression strategy significantly outperforms the more asymmetric $1 \times 4 \times 4$ configuration.

AE pre-training. We show that the autoencoder-based pre-training is necessary and critical for the compressor. As shown in Table 5, AE pre-training consistently leads to 2%-4% performance boosts on all benchmarks. It is worth mentioning that training the AE with long-form videos that has shot changes is critical.

Latent space constraint. We investigate the impact of different constraints on the latent space of autoencoders (AE) during pre-training, and results are shown in Table 6. Initially, we observed that when the compressor is pre-trained without any constraints, there are significant spikes in training loss, leading to feature alignment failure with LLMs. This result arises because the absence of constraints limits

the generalization of the learnable compressor, potentially causing it to encode unseen data into gaps or "holes" in the encoder’s feature space, resulting in a loss of meaningful representation. Introducing Gaussian distribution constraints on the latent space, as in VAE [14], partially resolves the issue but still yields suboptimal results. One possible reason for this is the increased difficulty in learning feature representations due to the inherent nondeterminism of the features. In contrast, our proposed residual constraints demonstrate superior performance.

5. Conclusion

In this work, we present a novel advancement in video representation learning through our novel dual-component approach. The proposed adaptive video sampler effectively reduces spatial and temporal redundancies by leveraging contextual information, while our VAE-based spatiotemporal video compressor further optimizes video features without significant information loss. The experimental results across multiple video benchmarks demonstrate that our method not only achieves state-of-the-art performance, but also provides a more efficient solution for long-form video understanding. This work opens up new possibilities for processing and analyzing long video content in resource-constrained environments, marking an important step forward in the field of video understanding with large multimodal models. Future work could explore the adaptation of our method to real-time applications and investigate its potential in other multimedia domains.

References

- [1] Daniel Bolya and Judy Hoffman. Token merging for fast stable diffusion. In *CVPR Workshop*, 2023. 3
- [2] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token Merging: Your ViT but faster. In *ICLR*, 2022. 3
- [3] Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018. 6
- [4] Guo Chen, Yifei Huang, Jilan Xu, Baoqi Pei, Zhe Chen, Zhiqi Li, Jiahao Wang, Kunchang Li, Tong Lu, and Limin

Wang. Video mamba suite: State space model as a versatile alternative for video understanding. *arXiv preprint arXiv:2403.09626*, 2024. 3

[5] Lin Chen, Xilin Wei, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Bin Lin, Zhenyu Tang, et al. Sharegpt4video: Improving video understanding and generation with better captions. *arXiv preprint arXiv:2406.04325*, 2024. 7

[6] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-lmms. *arXiv preprint arXiv:2406.07476*, 2024. 3, 4, 5

[7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3

[8] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The “something something” video database for learning and evaluating visual common sense. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5842–5850, 2017. The Something-Something dataset is licensed for non-commercial, research purposes at <https://developer.qualcomm.com/downloads/data-license-agreement-research-use?referrer=node/68935>. 6

[9] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022. 5, 6

[10] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. 3

[11] Md Mohaiminul Islam and Gedas Bertasius. Long movie clip classification with state-space video models. In *ECCV*, 2022. 1, 2

[12] Md Mohaiminul Islam, Ngan Ho, Xitong Yang, Tushar Nagarajan, Lorenzo Torresani, and Gedas Bertasius. Video recap: Recursive captioning of hour-long videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18198–18208, 2024. 2

[13] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021. 7

[14] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 8

[15] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *CVPR*, 2014. 1, 2

[16] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 780–787, 2014. 5

[17] Ang Li, Meghana Thotakuri, David A Ross, João Carreira, Alexander Vostrikov, and Andrew Zisserman. The ava-kinetics localized human actions video dataset. *arXiv preprint arXiv:2005.00214*, 2020. 5

[18] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024. 1, 2, 3, 4, 6, 7

[19] KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. 7

[20] Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22195–22206, 2024. 2, 6

[21] Xirui Li, Chao Ma, Xiaokang Yang, and Ming-Hsuan Yang. VidToMe: Video token merging for zero-shot video editing. In *CVPR*, 2024. 3

[22] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models, 2023. 7

[23] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models. In *European Conference on Computer Vision*, pages 323–340. Springer, 2025. 2, 3

[24] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. *arXiv preprint arXiv:2202.07800*, 2022. 3

[25] Ruotong Liao, Max Erler, Huiyu Wang, Guangyao Zhai, Gengyuan Zhang, Yunpu Ma, and Volker Tresp. Videoinsta: Zero-shot long video understanding via informative spatial-temporal reasoning with llms. *arXiv preprint arXiv:2409.20365*, 2024. 2, 6, 7

[26] Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023. 1, 2, 3, 4

[27] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 3

[28] Xiao Liu, Chenxu Zhang, and Lei Zhang. Vision mamba: A comprehensive survey and taxonomy. *arXiv preprint arXiv:2405.04404*, 2024. 3

[29] 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, 2, 3, 4

[30] 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. 2, 6

[31] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. AdaViT: Adaptive vision transformers for efficient image recognition. In *CVPR*, 2022. 3

[32] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3

[33] Viorica Patraucean, Lucas Smaira, Ankush Gupta, Adria Recasens, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Mateusz Malinowski, Yi Yang, Carl Doersch, et al. Perception test: A diagnostic benchmark for multimodal video models. *Advances in Neural Information Processing Systems*, 36, 2024. 2, 6

[34] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. DynamicViT: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, 2021. 3

[35] Shuhuai Ren, Sishuo Chen, Shicheng Li, Xu Sun, and Lu Hou. TESTA: Temporal-spatial token aggregation for long-form video-language understanding. In *EMNLP*, 2023. 3

[36] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation, parallel distributed processing, explorations in the microstructure of cognition, ed. de rumelhart and j. mcclelland. vol. 1. 1986. *Biometrika*, 71(599–607):6, 1986. 2

[37] Xiaoqian Shen, Yunyang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, Chenchen Zhu, Zechun Liu, Fanyi Xiao, Balakrishnan Varadarajan, Florian Bordes, et al. Longvu: Spatiotemporal adaptive compression for long video-language understanding. *arXiv preprint arXiv:2410.17434*, 2024. 1, 2, 3

[38] Enxin Song, Wenhao Chai, Guanhong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, et al. Moviechat: From dense token to sparse memory for long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18221–18232, 2024. 1, 2, 3, 7

[39] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 2

[40] Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. COIN: A large-scale dataset for comprehensive instructional video analysis. In *CVPR*, 2019. 1, 2

[41] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 3

[42] Junke Wang, Xitong Yang, Hengduo Li, Zuxuan Wu, and Yu-Gang Jiang. Efficient video transformers with spatial-temporal token selection. *arXiv preprint arXiv:2111.11591*, 2021. 3

[43] Jue Wang, Wentao Zhu, Pichao Wang, Xiang Yu, Linda Liu, Mohamed Omar, and Raffay Hamid. Selective structured state-spaces for long-form video understanding. In *CVPR*, 2023. 1, 2, 3

[44] Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4581–4591, 2019. 5

[45] Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. Videoagent: Long-form video understanding with large language model as agent. *arXiv preprint arXiv:2403.10517*, 2024. 2, 6, 7

[46] Olivia Wiles, Joao Carreira, Iain Barr, Andrew Zisserman, and Mateusz Malinowski. Compressed vision for efficient video understanding. In *Proceedings of the Asian Conference on Computer Vision*, pages 4581–4597, 2022. 3

[47] Chao-Yuan Wu and Philipp Krähenbühl. Towards long-form video understanding. In *CVPR*, 2021. 1, 2

[48] 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. 2, 6

[49] 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. *arXiv preprint arXiv:2407.15841*, 2024. 1, 2, 3, 4

[50] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024. 4

[51] Yijun Yang, Zhaohu Xing, and Lei Zhu. Vivim: a video vision mamba for medical video object segmentation. *arXiv preprint arXiv:2401.14168*, 2024. 3

[52] Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B Tenenbaum. Cleverer: Collision events for video representation and reasoning. *arXiv preprint arXiv:1910.01442*, 2019. 6

[53] Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Verhari, Kihyuk Sohn, David Minnen, Yong Cheng, Vighnesh Birodkar, Agrim Gupta, Xiuye Gu, et al. Language model beats diffusion-tokenizer is key to visual generation. *arXiv preprint arXiv:2310.05737*, 2023. 3

[54] Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yuet-ting 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. 2, 6

[55] Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, and Gedas Bertasius. A simple llm framework for long-range video question-answering. *arXiv preprint arXiv:2312.17235*, 2023. 6, 7

[56] Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, and Gedas Bertasius. A simple llm framework for long-range video question-answering. *arXiv preprint arXiv:2312.17235*, 2023. 2

[57] Peiyuan Zhang, Kaichen Zhang, Bo Li, Guangtao Zeng, Jingkang Yang, Yuanhan Zhang, Ziyue Wang, Haoran Tan, Chunyuan Li, and Ziwei Liu. Long context transfer from language to vision. *arXiv preprint arXiv:2406.16852*, 2024.

[7](#)

[58] Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jia Shi Feng, Ziwei Liu, and Chunyuan Li. Llava-next: A strong zero-shot video understanding model, 2024.

[7](#)

[59] Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang, Tiejun Huang, and Zheng Liu. Mlvu: A comprehensive benchmark for multi-task long video understanding. *arXiv preprint arXiv:2406.04264*, 2024. [6](#)