

# Scaling Video-Language Models to 10K Frames via Hierarchical Differential Distillation

Chuanqi Cheng<sup>\*1,2</sup> Jian Guan<sup>\*2</sup> Wei Wu<sup>†2</sup> Rui Yan<sup>†1,3</sup>

## Abstract

Long-form video processing fundamentally challenges vision-language models (VLMs) due to the high computational costs of handling extended temporal sequences. Existing token pruning and feature merging methods often sacrifice critical temporal dependencies or dilute semantic information. We introduce differential distillation, a principled approach that systematically preserves task-relevant information while suppressing redundancy. Based on this principle, we develop ViLAMP, a hierarchical video-language model that processes hour-long videos at “mixed precision” through two key mechanisms: (1) differential keyframe selection that maximizes query relevance while maintaining temporal distinctiveness at the frame level and (2) differential feature merging that preserves query-salient features in non-keyframes at the patch level. Hence, ViLAMP retains full information in keyframes while reducing non-keyframes to their most salient features, resembling mixed-precision training (Micikevicius et al., 2018). Extensive experiments demonstrate ViLAMP’s superior performance across five video understanding benchmarks, particularly on long-form content. Notably, ViLAMP can process ultra-long videos (up to 10K frames) on a single NVIDIA A100 GPU, achieving substantial computational efficiency while maintaining state-of-the-art performance. Code and model are available at <https://github.com/steven-ccq/ViLAMP>.

<sup>\*</sup>Equal contribution <sup>1</sup>Gaoling School of Artificial Intelligence, Renmin University of China <sup>2</sup>Ant Group <sup>3</sup>School of Artificial Intelligence, Wuhan University. Correspondence to: Wei Wu <wuwei19850318@gmail.com>, Rui Yan <ruiyan@ruc.edu.cn>.

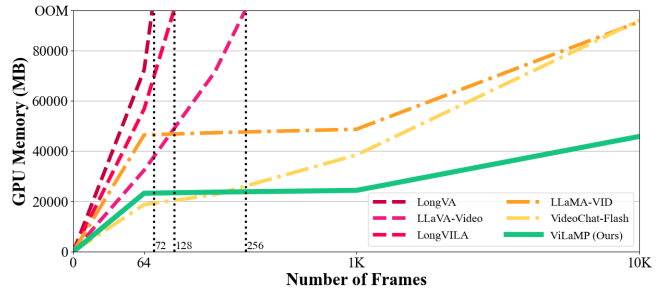


Figure 1. Comparison of the GPU memory consumption with varying input lengths (from 0 to 10K frames) on our VideoNIAH benchmark (c.f. §5.3). OOM indicates out-of-memory errors. All models are evaluated using one NVIDIA A100 GPU.

## 1. Introduction

The rapid advancement of vision-language models (VLMs) continues to drive remarkable progress in processing both visual and textual content (OpenAI, 2024; Liu et al., 2024a; Chen et al., 2024a; Zhang et al., 2024b). However, as models expand their capabilities from processing single images (typically requiring 729 visual tokens (Zhai et al., 2023)) to handling long videos, they inevitably encounter a fundamental challenge: the visual token sequence from a video can easily exceed the context length of LLMs, imposing prohibitive costs in both computational resources and inference latency. For instance, a one-minute video clip at 24 frames per second (FPS) could generate over 1 million visual tokens ( $24 \times 60 \times 729$ ), far surpassing the context capacity of most mainstream LLMs, which typically ranges from 4K to 128K tokens (Dubey et al., 2024; Yang et al., 2024a). The challenge becomes even more critical as video lengths extend to an hour or beyond, despite the prevalence of such lengthy videos in real-world applications, including long-form video analysis (Lin et al., 2023) and continuous robot learning (Wu et al., 2023; Baker et al., 2022).

Prior works have explored two main approaches to handle long videos efficiently: (1) token pruning through uniform or content-aware sampling strategies (Ye et al., 2023; Shen et al., 2024), and (2) feature merging via heuristic or learnable mechanisms (Lin et al., 2024; Bai et al., 2023). How-

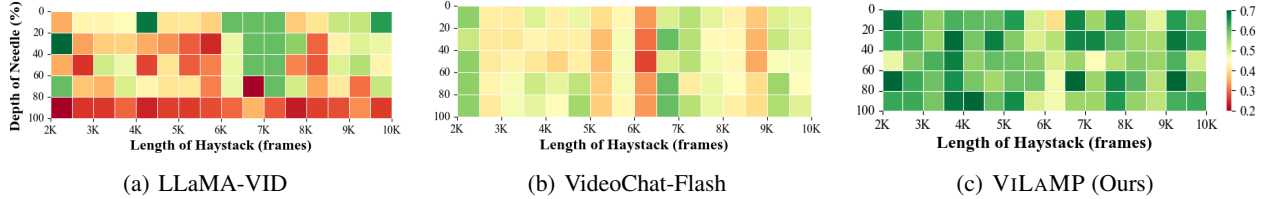


Figure 2. Comparison of ultra-long video understanding capabilities on our VideoNIAH benchmark (c.f. §5.3). Each heatmap shows model performance across different haystack lengths and needle depths. The depth refers to the relative position of the needle video within the haystack with 0% at the front and 100% at the end.

ever, these methods face fundamental limitations: token pruning risks losing critical temporal dependencies, while feature merging often leads to information dilution (Liu et al., 2024e). The challenge of balancing computational efficiency and semantic preservation thus remains unsolved.

In this work, we explore scaling video-language models to handle significantly longer videos by systematically distilling salient signals into minimal visual tokens without sacrificing information integrity. We begin by addressing a fundamental question: *how to identify truly salient signals from long videos?* While the saliency of visual signals is traditionally measured by query relevance (Li et al., 2024d), videos often contain substantial temporal correlations and repeated content (Shen et al., 2024). This suggests that merely considering query relevance might lead to selecting redundant signals. To validate this hypothesis, we conduct a comprehensive analysis at both frame and patch levels (c.f., §3). Our analysis reveals two key findings: (1) At the frame level,  $\sim 90\%$  of query-induced attention weights concentrate on merely 5% of frames, with these attention-heavy frames showing significant internal similarities. (2) At the patch level,  $\sim 50\%$  of patches from less-attended frames command 80% of model attention while exhibiting high visual similarity to those in attention-heavy frames, indicating redundant attention to similar visual patterns. These findings inspire our *differential distillation principle* for efficient video modeling: distilling salient information that exhibits maximal correlation with the query intent while maintaining minimal redundancy with respect to the video’s temporal context.

Building on the principle, we introduce **ViLAMP**, an efficient **V**ideo-**L**anguage model specifically designed to process long videos using “**M**ixed **P**recision” — allocating computational resources according to information saliency. ViLAMP implements this mixed-precision strategy throughout a hierarchical compression framework that operates at two levels: (1) At the frame level, ViLAMP maintains full visual token representations for salient frames (i.e., keyframes) while applying aggressive compression to intermediate frames (i.e., non-keyframes), reducing each to a single token. Frame saliency is determined through our novel *Differential Keyframe Selection* (**DKS**) mechanism,

which identifies keyframes by maximizing query relevance while ensuring temporal distinctiveness. (2) At the patch level, we introduce the *Differential Feature Merging* (**DFM**) strategy for non-keyframe compression, which employs a differentiable learning mechanism to preserve query-salient features while suppressing redundant visual information.

We conduct extensive evaluations of ViLAMP across five video understanding benchmarks, spanning from minutes to nearly 3 hours. ViLAMP outperforms state-of-the-art baselines on five benchmarks. Notably, on the “Long” subset of Video-MME, which features an average duration of 2,386 seconds, ViLAMP achieves a substantial 3.0% and 4.8% absolute accuracy improvement over the previous top-performing model in non-subtitled and subtitled settings, respectively. To rigorously assess VLMs’ capability in ultra-long video understanding, we introduce VideoNIAH, a challenging benchmark inspired by the Needle-in-a-Haystack (NIAH) evaluation paradigm from LLM research (Fu et al., 2024b; Kuratov et al., 2024). This benchmark requires models to both locate and comprehend specific needle videos embedded within extensive video sequences (haystacks) while answering targeted queries about these needles. As demonstrated in Figure 1 and Figure 2, when scaling existing models to process 10K frames (approximately 2.7 hours at 1FPS), they either fail entirely due to memory constraints or suffer severe performance drops compared to the 2K-frame counterpart (e.g.,  $>24.50\%$  accuracy degradation for VideoChat-Flash (Li et al., 2024c)). In contrast, ViLAMP not only successfully processes videos containing up to 10K frames on a single NVIDIA A100 GPU, but also maintains remarkably stable performance with an 12.82% absolute improvement over VideoChat-Flash<sup>1</sup>. These results convincingly demonstrate that ViLAMP establishes a promising direction for both effective and efficient long-form video understanding.

In summary, our primary contributions are threefold:

<sup>1</sup>Unlike NIAH tasks in LLM research where answers can be directly extracted from the needle text, VideoNIAH requires comprehensive video understanding capabilities to interpret visual-temporal events, making it inherently more challenging to achieve the near-perfect accuracy (e.g., 99%) observed in LLM evaluations.

I. We introduce the differential distillation principle that systematically identifies and preserves task-relevant information while actively suppressing redundancy, establishing a principled approach for efficient long video understanding.

II. We propose ViLAMP, a hierarchical architecture that operationalizes differential distillation through keyframe selection and non-keyframe compression, reducing non-keyframes to their most salient features.

III. Experiments show that ViLAMP outperforms similar-sized VLMs on five video understanding benchmarks while efficiently scaling to 10K frames on a single GPU, opening new opportunities for processing ultra-long videos.

## 2. Related Work

### 2.1. Vision-Language Models

Advances in large language models (LLMs) (Dubey et al., 2024; Yang et al., 2024a) has catalyzed the development of VLMs that can process visual and textual inputs, broadly classified into two categories: (1) models jointly trained on visual and textual modalities from scratch, such as Flamingo (Alayrac et al., 2022) and KOSMOS (Peng et al., 2023a;b); and (2) models leveraging pre-trained vision encoders like CLIP (Radford et al., 2021b) or SigLIP (Zhai et al., 2023), connected to LLMs through projection modules. This design, pioneered by LLaVA (Liu et al., 2024a) with a simple two-layer MLP connector, has been widely adopted in Qwen-VL (Bai et al., 2023), MiniCPM (Yao et al., 2024), and InternVL (Chen et al., 2024c).

While recent VLMs can handle multiple images (Wang et al., 2024a; Li et al., 2024b) and videos (Zhang et al., 2024c; Anonymous, 2024; Zhang et al., 2024e), they face significant limitations: multi-image models ignore temporal dependencies, while video models struggle with computational costs for long videos. This highlights the need for more efficient approaches to video understanding.

### 2.2. Long-form Video Understanding

Current approaches to addressing the computational challenges of long video understanding fall into three categories:

(1) Context Window Extension. This line of research focuses on enhancing models’ capacity for extended temporal contexts. Liu et al. (2024b) gradually increases input length during training, while LongVA (Zhang et al., 2024a) transfers long-context capabilities from text to visual understanding. However, these approaches still face efficiency limitations with extensive videos.

(2) Token Pruning. Various pruning strategies have emerged to reduce computational burden. Early models (Lin et al., 2023; Ye et al., 2024; Cheng et al., 2024) used uniform

frame sampling, i.e., selecting fixed-interval frames. Recent works propose more sophisticated content-aware pruning approaches: Goldfish (Ataallah et al., 2024) employs database-inspired retrieval mechanisms, MovieChat (Song et al., 2024) implements dynamic memory management, and LongVU (Shen et al., 2024) introduces content-aware frame grouping and similarity-based patch selection. However, such pruning risks significant information loss.

(3) Feature Merging. Recent research has explored merging video features at various representational levels for video compression, through either heuristic methods (e.g., downsampling (Lin et al., 2024), pooling (Xu et al., 2024; Zhang et al., 2024e; Li et al., 2025), similarity-guided merging (Bolya et al., 2023; Jin et al., 2024; Li et al., 2024c)) or learnable approaches (e.g., Q-Former (Li et al., 2023; Bai et al., 2023; Zhang et al., 2025), Perceiver Resampler (Yao et al., 2024)). However, these strategies often struggle to maintain semantic fidelity (Liu et al., 2024e).

In contrast, ViLAMP preserves salient information through differential distillation, achieving both computational efficiency and information preservation.

## 3. Preliminary Studies

Video understanding tasks process a sequence of frames  $V = \langle f_1, f_2, \dots, f_N \rangle$  to answer a query  $Q$ . Each frame  $f_n$  is divided into  $M$  patches:  $f_n = \langle p_n^1, p_n^2, \dots, p_n^M \rangle$ , where each patch  $p_n^m$  is encoded into a unique visual token  $t_n^m$  through a vision encoder. Consequently, a video sequence generates  $MN$  visual tokens in total. This section analyzes information distribution patterns at both frame level (§3.1) and patch level (§3.2), leading to our differential distillation principle (§3.3).

### 3.1. Frame-Level Redundancy Analysis

We analyze attention patterns of several representative VLMs, including LLaVA-OneVision (Li et al., 2024b), LLaVA-Video (Zhang et al., 2024e), Qwen2-VL (Wang et al., 2024a) and LongVA (Zhang et al., 2024c), on Video-MME (Fu et al., 2024a) by sampling  $N=128$  frames per video. For each frame  $f_n$ , we compute its receiving attention weight  $a_n$  by aggregating cross-attention weights between query  $Q$  and visual tokens  $\{t_n^m\}_{m=1}^M$ :  $a_n = \sum_{m=1}^M a_n^m \in [0, 1]$ . The token-level attention weight  $a_n^m \in [0, 1]$  is obtained by averaging across all query tokens, all layers, and all attention heads.

Figure 3 reveals two key patterns. First, about 90% of attention mass concentrates on less than 5% of frames, as evidenced by the steep rise in the cumulative attention curves

<sup>2</sup>Our study focuses on visual information, excluding subtitles or audio tracks.

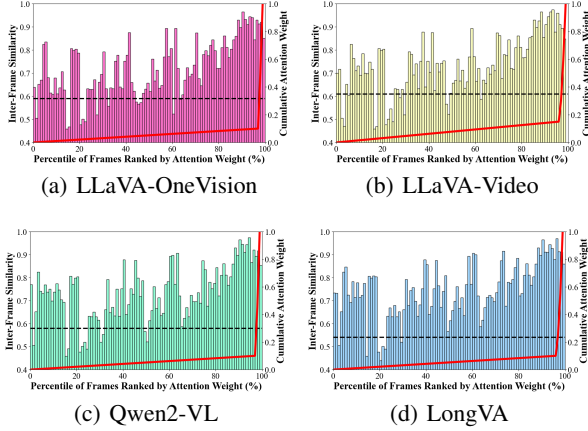


Figure 3. Frame-level attention distribution and frame similarities.  $x$ -axis: percentile of frames ranked by attention weights from the query (%). Red curves: cumulative attention. Bars: average similarity with neighboring three frames of similar attention. Dashed lines: random frame pair similarity baseline (0.54~0.61).

(red lines). This pronounced attention skew suggests **query-relevant information clusters around specific keyframes**. Second, and more intriguingly, these attention-heavy frames show high visual similarity (cosine similarity  $> 0.8$ ) within themselves, substantially exceeding the random baseline similarity (dashed lines). This phenomenon indicates that **frames receiving high attention weights contain significant visual redundancy**. Together, these findings provide compelling empirical evidence for developing more efficient video processing strategies by eliminating redundant computation on visually similar, high-attention frames.

### 3.2. Patch-Level Redundancy Analysis

While frame-level redundancy analysis suggests processing only keyframes (Ye et al., 2023; Shen et al., 2024) for efficiency, this approach faces two critical limitations: (1) potential loss of pivotal moments due to selection uncertainty, and (2) missing temporal context essential for tasks like action recognition and causal reasoning. Although this necessitates effective compression strategies for non-keyframes, naive feature merging approaches (e.g., pooling (Xu et al., 2024)) can be suboptimal due to substantial information overlap between non-keyframes and keyframes, risking attention misattribution to redundant information.

To verify the hypothesis, we analyze video frames by selecting the top 32 attention-weighted frames as keyframes. For each patch in the remaining 96 non-keyframes, we compute its query attention weight and visual similarity to its corresponding patch in the preceding keyframe (or the first keyframe for earlier non-keyframes). The visual similarity is measured using cosine similarity between the model’s visual embeddings. Figure 4 reveals that high-attention patches

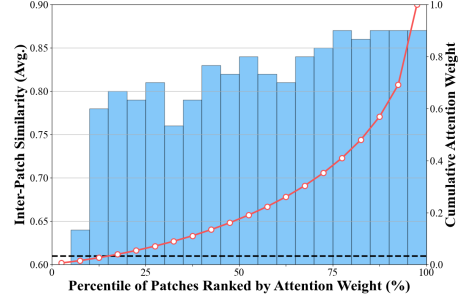


Figure 4. Patch-level attention distribution of LLaVA-OneVision and cross-frame patch similarity.  $x$ -axis: percentile of non-keyframe patches ranked by attention weights from the query (%). Red curve: cumulative attention (normalized across all non-keyframe patches). Bars: similarity to nearest keyframe patches. Dashed line: random inter-patch similarity baseline (0.61).

show strong visual similarity to their keyframe counterparts, with just 50% of patches contributing 80% of total attention. This correlation between attention weights and visual similarity indicates that **substantial computation is spent processing redundant visual patterns across frames**.

### 3.3. Differential Distillation Principle

Our frame- and patch-level analyses motivate our differential distillation principle, fundamentally rethinking how to identify and preserve salient information while reducing redundant computation. The key insight is that truly important information should simultaneously satisfy two criteria: (1) high relevance to the query intent, and (2) low redundancy with respect to its temporal context. Formally, for any video component  $v$  (frame, patch, or feature) and query  $Q$ , we define its differential information saliency score:

$$D(v) = R(v, Q) - T(v, \mathcal{C}(v)), \quad (1)$$

where  $R(v, Q)$  measures query relevance and  $T(v, \mathcal{C}(v))$  captures temporal redundancy with context features  $\mathcal{C}(v)$ . A higher  $D(v)$  indicates more unique, task-relevant information deserving more computational resources.

This formulation provides two key advantages: (1) natural adaptation to varying temporal granularities, enabling unified frame- and patch-level operations, and (2) query-aware processing that dynamically allocates computation based on task requirements.

## 4. Methodology

Following the differential distillation principle, we propose ViLAMP, a hierarchical architecture for mixed-precision video processing, as shown in Figure 5. ViLAMP combines frame-level selection via *Differential Keyframe Selector* (DKS) (§4.1) and patch-level compression via *Dif-*

*differential Feature Merger* (DFM) (§4.2), enabling efficient resource allocation based on information saliency.

#### 4.1. Differential Keyframe Selection

The differential keyframe selection mechanism identifies frames with high query relevance while maintaining temporal distinctiveness. Given frames  $V = \langle f_1, f_2, \dots, f_N \rangle$  and query  $Q$ , we first encode both the query and each frame into normalized  $d$ -dimensional embeddings using a CLIP encoder  $E_f(\cdot)$  (Radford et al., 2021a), and then compute the frame-level query relevance score  $R_f(f_n, Q)$  as the cosine similarity between these embeddings:

$$R_f(f_n, Q) = \cos(\mathbf{f}_n, \mathbf{q}), \quad (2)$$

$$\mathbf{f}_n = E_f(f_n) \in \mathbb{R}^d, \quad \mathbf{q} = E_f(Q) \in \mathbb{R}^d, \quad (3)$$

Then we instantiate the temporal redundancy function in Eq. 1 at the frame level as the maximum cosine similarity between frame  $f_n$  and its context frames  $\mathcal{C}(f_n)$ :

$$T_f(f_n, \mathcal{C}(f_n)) = \max_{f \in \mathcal{C}(f_n)} \cos(\mathbf{f}_n, E_f(f)) \quad (4)$$

While one could naively set  $\mathcal{C}(f_n) = V \setminus \{f_n\}$ , i.e., measuring frame-level redundancy by comparing each frame against all others, this would incur prohibitive quadratic complexity  $O(N^2)$ . Moreover, selecting frames that directly maximize  $R_f - T_f$  without proper prioritization of two terms risks selecting frames that are temporally distinct but irrelevant to the query intent. Therefore, we propose an efficient greedy algorithm (Algorithm 1). It ranks frames by query relevance and selects those maintaining sufficient temporal distance (threshold  $\tau$ ), achieving  $O(\max(NK, N \log N))$  complexity for  $K$  keyframes while ensuring both semantic relevance and temporal diversity.

---

#### Algorithm 1 Differential Keyframe Selection

---

**Data:**  $V = \langle f_1, f_2, \dots, f_N \rangle$ : Video frames;  $Q$ : Query;  $K$ : Maximum keyframes;  $\tau$ : Similarity threshold.

**Result:**  $\mathcal{K}$ : Keyframe Set.

```

// Sort all frames by query-relevance in descending order
1  $\langle \hat{f}_1, \hat{f}_2, \dots, \hat{f}_N \rangle = \text{sorted}(V, \text{key} = \text{lambda } f_n : R_f(f_n, Q))$ 
// Initialize  $\mathcal{K}$  with the most relevant frame
2  $\mathcal{K} = \{\hat{f}_1\}$ 
// Iterate through the sorted frames
3 while  $n \leftarrow 1$  to  $N$  do
4     if  $|\mathcal{K}| < K$  and  $T_f(\hat{f}_n, \mathcal{K}) < \tau$  then
5          $\mathcal{K}.\text{add}(\hat{f}_n)$ 
6 return  $\mathcal{K}$ 
    
```

---

#### 4.2. Differential Feature Merging

While DKS identifies salient keyframes, non-keyframes provide crucial temporal context. Instead of discarding these frames as in mainstream approaches (Yang et al., 2024a;

Chen et al., 2024c), we propose a differential feature merging (DFM) approach that compresses non-keyframes to single tokens while preserving task-relevant information.

Specifically, for a non-keyframe  $f_n = \langle p_n^1, p_n^2, \dots, p_n^M \rangle$  and its nearest preceding keyframe  $f_k = \langle p_k^1, p_k^2, \dots, p_k^M \rangle$ , we compute patch-level embeddings using the VLM’s visual encoder  $E_p(\cdot)$ , that generates a  $d$ -dimensional embedding for each patch. The differential information saliency  $D_p$  of each patch  $p_n^m$  is determined by two factors, i.e., its query relevance  $R_p$  and temporal redundancy  $T_p$ :

$$D_p(p_n^m) = R_p(p_n^m, Q) - \lambda T_p(p_n^m, p_k^m), \quad (5)$$

$$R_p(p_n^m, Q) = \cos(\mathbf{p}_n^m, \mathbf{q}), \quad (6)$$

$$T_p(p_n^m, \mathcal{C}(p_n^m)) = \cos(\mathbf{p}_n^m, \mathbf{p}_k^m), \quad (7)$$

$$\{\mathbf{p}_n^m\}_{m=1}^M = E_p(f_n), \quad \{\mathbf{p}_k^m\}_{m=1}^M = E_p(f_k) \quad (8)$$

where  $m$  is the patch index,  $k/n$  refers to the keyframe/non-keyframe index,  $\lambda$  balances the trade-off between query relevance and temporal uniqueness, and  $\mathbf{p}_n^m/\mathbf{p}_k^m$  and  $\mathbf{q}$  represent the encoded patch and query embeddings, respectively.  $\mathcal{C}(p_n^m)$  is set to  $\{p_k^m\}$  to compute temporal redundancy only with respect to the corresponding patch position in the keyframe, leveraging the strong spatial correspondence prior in video content while maintaining computational efficiency. We set  $T_p = 0$  for non-keyframes without preceding keyframes.

The saliency scores guide our adaptive token merging strategy through a differentially weighted pooling operation:

$$\mathbf{t}_n = \frac{\sum_{m=1}^M w_n^m \mathbf{p}_n^m}{\sum_{m=1}^M w_n^m}, \quad (9)$$

$$w_n^m = \text{softmax}\left(\frac{1}{\alpha} [D_p(p_n^1), \dots, D_p(p_n^M)]\right)|_m, \quad (10)$$

where  $\mathbf{t}_n$  is the compressed representation for the non-keyframe  $f_n$ ,  $w_n^m$  is the pooling weights, and  $\alpha$  controls the sharpness of the weight distribution. This formulation ensures that patches receive high-importance weights when they contain query-relevant information that is not already captured by the corresponding keyframe patches.

#### 4.3. Multimodal Learning

We utilize a two-stream visual connector architecture to integrate heterogeneous representations from keyframes and compressed non-keyframes into the language model:

$$\mathbf{h}_k^m = \text{MLP}_k(\mathbf{p}_k^m), \quad \mathbf{h}_n = \text{MLP}_n(\mathbf{t}_n), \quad (11)$$

where  $\mathbf{h}_k^m$  is the projected embedding for the  $m$ -th patch of keyframe  $f_k$  and  $\mathbf{h}_n$  is the embedding for the compressed representation of non-keyframe  $f_n$ , processed through separate learnable two-layer MLPs. The model is then trained

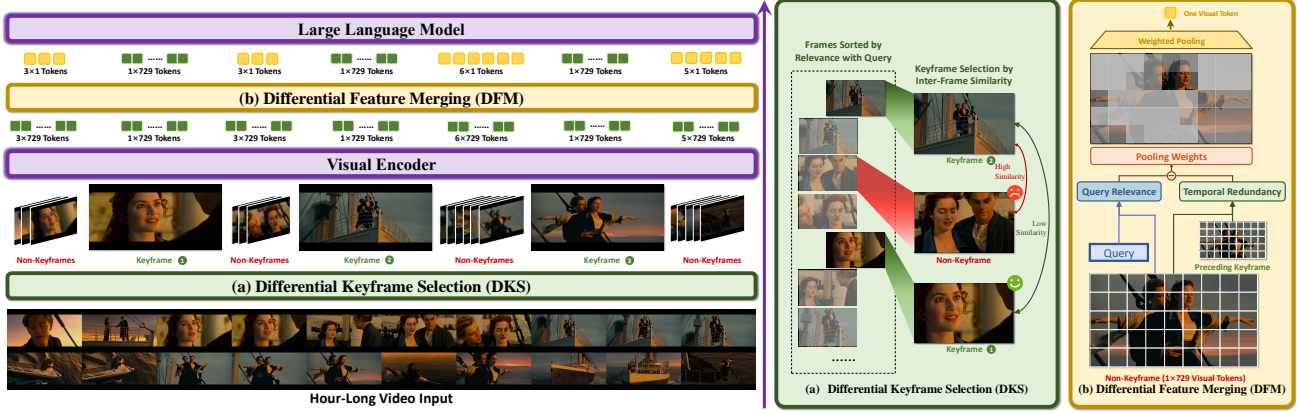


Figure 5. Overview of ViLAMP, which first identifies representative keyframes from videos using the differential keyframe selection algorithm (a) and then adaptively merges visual features from non-keyframes through the differential feature merging mechanism (b).

using a language modeling objective over the concatenated sequence of visual embeddings and the prompt:

$$\mathcal{L} = -\log P(A | \{h_k^m | f_k \in \mathcal{K}\} \cup \{h_n | f_n \notin \mathcal{K}\}, Q), \quad (12)$$

where the keyframe and non-keyframe embeddings are arranged in their temporal order,  $Q$  and  $A$  represent the query and target answer, respectively.

This design of ViLAMP offers three key advantages. First, it dramatically reduces the total number of visual tokens from  $MN$  to  $MK + (N - K)$  where  $K \ll N$ , enabling efficient processing of extended video sequences. Second, the full-token keyframe representations serve as natural anchors for learning temporal relationships, facilitating the model’s interpretation of compressed intermediate frames. Third, the end-to-end optimization of DFM parameters through direct supervision ensures that our compression mechanism adapts to preserve task-relevant information effectively.

## 5. Experiments

### 5.1. Experimental Setup

**Benchmarks.** We thoroughly evaluate ViLAMP on five video understanding benchmarks spanning diverse temporal scales and tasks (See Appendix A.1 for details): LVBench (Wang et al., 2024b) for long-term decision-making, EgoSchema (Mangalam et al., 2023) for natural scenario understanding, LongVideoBench (Wu et al., 2024) for referred reasoning, MLVU (Zhou et al., 2024) for multi-task question answering, and Video-MME (Fu et al., 2024a) for comprehensive video understanding.

**Baselines.** We conduct comprehensive comparisons with state-of-the-art baselines spanning proprietary VLMs, open-source multi-image VLMs, and open-source video-language models. Please refer to Appendix A.2 for more details.

**Implementation Details.** ViLAMP employs the same architecture as LLaVA-OneVision (Li et al., 2024a), utilizing SigLIP-so400m<sup>3</sup> as the visual encoder, a two-layer MLP as the vision-language connector and Qwen2-7B<sup>4</sup> as the language model. We process frames at  $384 \times 384$  resolution and empirically set  $\tau$  in Alg. 1 to 0.85, the keyframe count  $K$  to 32,  $\lambda$  in Eq. 5 to 1, and  $\alpha$  in Eq. 10 to  $10^{-2}$  unless otherwise specified. We employ CLIP-ViT-B-32<sup>5</sup> as  $E_f(\cdot)$  in Eq. 3. More training details are provided in Appendix A.3.

### 5.2. Main Results

Table 1 presents comprehensive evaluation results across five benchmarks. Among open-source models of similar size, ViLAMP establishes new state-of-the-art performance across diverse temporal scales, from short-form (EgoSchema) to medium-form (LongVideoBench and MLVU) and long-form (LVBench and Video-MME) video understanding tasks. Remarkably, ViLAMP performs competitively with or surpasses substantially larger open-source models ( $\sim 70B$  parameters) on these benchmarks. This superior performance is particularly evident in long-form video understanding: On MLVU, ViLAMP achieves 72.6% accuracy, even outperforming the proprietary GPT-4o model (64.6%); On the long-video subset of Video-MME, ViLAMP achieves 67.5% in non-subtitled setting and 73.5% in subtitled setting, outperforming existing similar-sized models by a large margin. These consistent strong results across varying temporal scales, particularly in challenging long-form scenarios, validate ViLAMP’s robust video modeling capabilities.

<sup>3</sup><https://huggingface.co/google/siglip-so400m-patch14-384>

<sup>4</sup><https://huggingface.co/Qwen/Qwen2-7B>

<sup>5</sup><https://huggingface.co/sentence-transformers/clip-ViT-B-32>

Table 1. Accuracy (%) on five video understanding benchmarks. **Size** indicates the number of parameters. **Frames** denotes either the fixed number of frames sampled from each video (e.g., 32 means sampling 32 frames regardless of video length) or the frame sampling rate (e.g., 2 FPS means sampling 2 frames per second). The best and second best results among open-source models of similar size (7~9B) are in **bold** and underlined, respectively. **M-Avg** refers to the average score across all multi-choice tasks. For Video-MME, we report the performance on the whole benchmark (“Overall”) and the long-form video subset (“Long”, videos > 39 minutes). All baseline results are collected from their original papers. “-” indicates results not found.

Models	Size	Frames	LVBench	EgoSchema	LongVideoBench	MLVU	Video-MME (wo / w sub)	
			Overall	Validation	Validation	M-Avg	Overall	Long
Average Duration			4101s	180s	473s	651s	1,010s	2,386s
Proprietary Vision-Language Models								
GPT4-V (Achiam et al., 2023)	Undisclosed		-	-	59.1	49.2	59.9 / 63.3	53.5 / 56.9
GPT4-o (OpenAI, 2024)	Undisclosed		30.8	-	66.7	64.6	71.9 / 77.2	65.3 / 72.1
Gemini-1.5-Pro (Reid et al., 2024)	Undisclosed		33.1	-	64.0	-	75.0 / 81.3	67.4 / 77.4
Open-Source Multi-Image Vision-Language Models								
LLaVA-OneVision (Li et al., 2024b)	72B	32	-	62.0	61.3	66.4	66.3 / 69.6	60.0 / 62.4
InternVL2 (Chen et al., 2024c)	76B	16	-	-	61.0	69.9	61.2 / 67.8	-
LLaVA-OneVision (Li et al., 2024a)	7B	32	-	60.1	56.5	64.7	58.2 / -	-
Oryx-1.5 (Liu et al., 2024d)	7B	128	-	-	56.3	67.5	58.8 / 64.2	-
InternVL2 (Chen et al., 2024c)	8B	16	-	-	54.6	64.0	56.3 / 59.3	-
InternVL2.5 (Chen et al., 2024b)	8B	64	-	-	60.0	68.9	64.2 / 66.9	-
Chat-Univi (Jin et al., 2024)	7B	64	-	-	-	-	40.6 / 45.9	35.8 / 41.8
MiniCPM-v2.6 (Yao et al., 2024)	8B	64	-	-	54.9	37.3	60.9 / 63.7	51.8 / 56.3
mPLUG-Owl3 (Ye et al., 2024)	7B	16	-	-	52.1	63.7	59.3 / -	50.1 / -
Qwen2-VL (Yang et al., 2024a)	7B	2FPS	-	-	55.6	-	63.3 / 69.0	-
NVILA (Liu et al., 2024e)	7B	256	-	-	-	70.1	64.2 / 70.0	54.8 / 63.3
Open-Source Video-Language Models								
VideoLLaMA2 (Cheng et al., 2024)	72B	16	-	63.9	-	61.2	62.4 / 64.7	57.6 / 59.0
LLaVA-Video (Zhang et al., 2024e)	7B	1FPS	-	65.6	58.2	70.8	63.3 / 69.7	-
LLaMA-VID (Li et al., 2024d)	7B	1FPS	23.9	38.5	-	33.2	25.9 / -	-
Video-XL (Shu et al., 2024)	7B	2,048	-	-	49.5	64.9	55.5 / 61.0	49.2 / -
Video-LLaVA (Lin et al., 2023)	7B	8	-	38.4	37.6	47.3	39.9 / 41.6	36.2 / 38.1
VideoLLaMA2 (Zhang et al., 2023)	7B	16	-	51.7	-	48.5	47.9 / 50.3	-
VideoChat2 (KunChang et al., 2023)	7B	16	-	54.4	36.0	47.9	54.6 / -	39.2 / -
Video-CCAM (Fei et al., 2024)	9B	96	-	-	-	58.5	53.2 / 57.4	46.7 / 49.9
PLLaVA (Xu et al., 2024)	7B	16	-	-	40.2	-	-	-
Movie-Chat (Song et al., 2024)	7B	2,048	22.5	-	-	-	-	-
Kangaroo (Liu et al., 2024c)	8B	64	-	62.7	54.8	61.0	56.0 / 57.6	46.7 / 59.3
LongVU (Shen et al., 2024)	7B	1FPS	-	67.6	-	65.4	60.6 / 59.5	-
LongVA (Zhang et al., 2024c)	7B	128	-	-	-	56.3	52.6 / 54.3	46.2 / 47.6
LongVILA (Chen et al., 2024a)	7B	256	-	67.7	57.1	-	60.1 / 65.1	-
ViLAMP(Ours)	7B	1FPS	45.2	70.2	61.2	72.6	67.5 / 73.5	57.8 / 68.1

### 5.3. Scaling Up to 10K Frames

The ability to process hour-long videos is crucial for real-world applications. However, evaluating models’ long video understanding capabilities remains challenging. While existing works often adapt the Needle-in-a-Haystack (NIAH) paradigm from LLM research to assess VLMs (Zhang et al., 2024a; Shen et al., 2024), these adaptations typically exhibit two key limitations: (1) using simplified single-frame “needles” that fail to capture temporal dynamics, and (2) limiting haystack sizes to at most 3K frames, inadequately testing true long-form processing capabilities.

To rigorously address above limitations, we develop VideoNIAH, a more challenging variant of the NIAH task specifically designed for video content. We construct VideoNIAH

by sampling long-form videos from Video-MME to create “haystacks” ranging from 2K to 10K frames (at 1FPS). We then insert “needle” video clips (30~120 seconds) within these haystacks at random positions. Both the needle videos and their corresponding query-answer pairs are sampled from the original Video-MME dataset. Models must not only locate these needles but also comprehend their temporal content to answer targeted queries. We create 3K test cases that distributed across five haystack lengths (2K, 4K, 6K, 8K, and 10K frames), with carefully balanced question types to ensure comprehensive evaluation. The testing was conducted in a non-subtitled setting.

**Efficacy.** Figure 2 compares ViLAMP against LLaMA-VID (Li et al., 2024d) and VideoChat-Flash (Li et al., 2024c)

Table 2. Ablation study results.

Model	LongVideoBench	MLVU	Video-MME
LLaVA-OneVision	56.5	64.7	58.2
+DKS (Ours)	57.0	67.3	61.8
+Query-Guided Sampling	54.8	63.7	55.6
+Uniformly Sampling	57.6	66.8	60.2
+DFM (Ours)	<b>60.5</b>	<b>72.6</b>	<b>65.1</b>
+Q-former	53.9	62.6	59.2
+Mean-Pooling	56.3	67.3	62.1

across varying haystack lengths, two state-of-the-art models specifically designed for ultra-long video understanding. The results reveal several key insights: (1) While all models show some performance degradation as context length increases, ViLAMP maintains significantly higher average accuracy at 10K frames (58.15% vs. 47.25% of VideoChat-Flash). (2) ViLAMP exhibits remarkably stable performance scaling, with 12.82% less drop of average accuracy from 2K to 10K frames compared to VideoChat-Flash, suggesting superior scalability to ultra-long videos.

**Efficiency.** Figure 1 shows that ViLAMP exhibits minimal memory growth with increased input length, achieving  $\sim 50\%$  lower memory consumption than baselines when processing 10K frames. To comprehensively evaluate computational efficiency, we additionally measure the required computational load (FLOPs) and inference latency (Time-to-First-Token, TTFT) under varying input lengths. As reported in Appendix B.1, ViLAMP shows substantial computational efficiency, particularly for longer inputs with FLOPs reduced by  $>80\%$  compared to baseline models at 8,192 frames while maintaining comparable inference speed.

#### 5.4. Discussions

In this section, we conduct a comprehensive investigation to validate our design choices. We first evaluate the effectiveness of DKS and DFM, and then analyze the sensitivity of model performance to the number of keyframes. Appendix B.2 further discusses other configuration settings.

**Ablation Study.** We evaluate the effectiveness of DKS and DFM through comparative experiments. For frame selection, we compare DKS with two baselines: (1) query-guided sampling, which selects 32 frames with the highest query relevance scores (Eq. 2) without considering inter-frame redundancy, and (2) uniform sampling, which extracts 32 frames at fixed intervals. For feature merging, we compare DFM against (1) Q-former (Li et al., 2023), which integrates 729 visual token sequences into 1 fixed-length embeddings through a learnable cross-attention module, and (2) mean-pooling, which averages token embeddings within each non-keyframe into one token. All models are trained on the same dataset as mentioned in Appendix A.3.

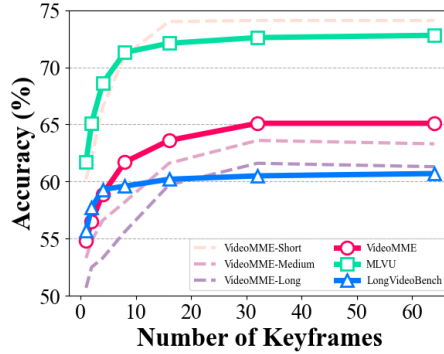


Figure 6. Analysis of model performance with varying numbers of keyframes across different video benchmarks. The test for Video-MME is conducted in non-subtitled setting.

Results in Table 2 demonstrate the effectiveness of both components. For frame selection, DKS shows particularly strong improvements on longer videos (MLVU and Video-MME), while uniform sampling becomes less effective with increased video length due to missing key information, and query-guided sampling underperforms due to locally concentrated frame selection. For feature merging, DFM consistently outperforms alternatives across all datasets by effectively preserving contextual information, whereas mean-pooling suffers from redundant information and Q-former is limited by fixed embedding constraints.

**Influence of Keyframes.** To determine the optimal trade-off between computational efficiency and model performance, we analyze ViLAMP’s behavior with varying numbers of keyframes. Results in Figure 6 show consistent performance improvements as keyframe count increases, until reaching dataset-specific saturation points. Short-form videos (Video-MME-Short and LongVideoBench) plateau around 16 keyframes, while longer videos (Video-MME-Medium/Long and MLVU) require 32 keyframes for optimal performance. This aligns with our findings in § 3.1 that informative frames can be sparsely sampled, though optimal sampling density varies with video length.

**Effectiveness on other video tasks** Although ViLAMP’s effectiveness has been verified on video question answering benchmarks, it is still unclear whether this method can truly preserve necessary information for more temporal-sensitive tasks like action recognition, video grounding or temporal reasoning. Therefore, we add some experiments.

For the video grounding task, we perform experiments on the widely used QVHighlights (Lei et al., 2021) benchmark. This dataset provides annotations for videos, marking the timestamps of query-relevant segments and the positions of keyframes. We evaluated ViLAMP using the Hit@k metric, which measures the probability of the ground-truth keyframe appearing in the top-k keyframes selected by our

method. The results are as follows:

Table 3. Evaluation results on QVHighlights.

Hit@1	Hit@5	Hit@10	Hit@30
26.1	47.3	60.0	83.2

As shown in Table 3, ViLAMP achieves 60% accuracy for Hit@10 and 83% for Hit@30, providing strong empirical support for its Dynamic Keyframe Selection (DKS) process. The current method effectively captures query-relevant video clips.

For the action recognition task, we incorporate two baselines: Adaframe (Wu et al., 2019) and MGSampler (Zhi et al., 2021), and evaluate them alongside ViLAMP on three benchmarks: ActivityNet (Caba Heilbron et al., 2015), Something-Something V2 (Sth-V2) (Goyal et al., 2017), and UCF-101 (Soomro et al., 2012). Since ViLAMP is a generative model, we reformulate the original classification task as a multiple-choice problem during testing. The results are as follows:

Table 4. Evaluation on action recognition tasks. The results for Adaframe and MGSampler are sourced from their papers.

Method	ActivityNet (mAP %)	Sth-V2 (Acc. %)	UCF-101 (Acc. %)
Adaframe	71.5	-	-
MGSampler	-	60.1	95.2
ViLAMP(Ours)	<b>78.6</b>	<b>86.9</b>	<b>97.6</b>

Notably, although ViLAMP is not trained on these datasets, it demonstrates robust performance. This may be attributed to the generalizable knowledge acquired by the large-scale pretrained model, enabling effective zero-shot transfer to action recognition tasks.

## 6. Conclusion

In this work, we address the critical challenge of scaling video-language models (VLMs) to process ultra-long videos efficiently. We present ViLAMP, a video-language model guided by differential distillation. By combining Differential Keyframe Selection (DFS) and Differential Feature Merging (DFM), ViLAMP demonstrate superior performance across five video benchmarks, particularly in long-form scenarios. On the challenging VideoNIAH benchmark, ViLAMP processes videos up to 10K frames ( $\sim 2.7$  hours) on a single NVIDIA A100 GPU, outperforming state-of-the-art models by 12.82% in accuracy while significantly reducing computational overhead. The hierarchical compression strategy enables stable performance across varying video lengths, with minimal memory growth and latency. Our work establishes a scalable paradigm for long-form video understanding, bridging the gap between computational constraints and real-world demands.

## Acknowledgments

This work was supported by the Beijing Outstanding Young Scientist Program NO. BJJWZYJH012019100020098, and Intelligent Social Governance Platform, Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Renmin University of China, the Fundamental Research Funds for the Central Universities, and Public Computing Cloud, Renmin University of China, the fund for building world-class universities (disciplines) of Renmin University of China.

## Impact Statement

This paper aims to propose a video-language model that efficiently handles ultra-long video inputs with relatively low computational cost. We would like to highlight a potential impact related to the generated output of our model, which may contain inaccuracies, biases, or offensive content. These issues may stem from various factors, including inherent biases in the training datasets, or the complex nature of video-language interpretation. Such outputs could potentially lead to misinformation propagation or unintended consequences in real-world applications. Therefore, users are expected to implement rigorous validation protocols when deploying this model.

## References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022.
- Anonymous. LongVILA: Scaling long-context visual language models for long videos. In *Submitted to The Thirteenth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=wCXAlfvCy6>. under review.
- Ataallah, K., Shen, X., Abdelrahman, E., Sleiman, E., Zhuge, M., Ding, J., Zhu, D., Schmidhuber, J., and El-hoseiny, M. Goldfish: Vision-language understanding of arbitrarily long videos. *arXiv preprint arXiv:2407.12679*, 2024.
- Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., and Zhou, J. Qwen-vl: A versatile vision-language model for understanding, localization, text read-

- ing, and beyond. *arXiv preprint arXiv:2308.12966*, 1(2): 3, 2023.
- Bain, M., Nagrani, A., Varol, G., and Zisserman, A. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *IEEE International Conference on Computer Vision*, 2021.
- Baker, B., Akkaya, I., Zhokov, P., Huizinga, J., Tang, J., Ecoffet, A., Houghton, B., Sampedro, R., and Clune, J. Video pretraining (VPT): Learning to act by watching unlabeled online videos. In Oh, A. H., Agarwal, A., Belgrave, D., and Cho, K. (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=AXDNM76Tlnc>.
- Bolya, D., Fu, C.-Y., Dai, X., Zhang, P., Feichtenhofer, C., and Hoffman, J. Token merging: Your vit but faster. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=JroZRrW7Eu>.
- Caba Heilbron, F., Escorcia, V., Ghanem, B., and Carlos Niebles, J. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 961–970, 2015.
- Chen, Y., Xue, F., Li, D., Hu, Q., Zhu, L., Li, X., Fang, Y., Tang, H., Yang, S., Liu, Z., He, E., Yin, H., Molchanov, P., Kautz, J., Fan, L., Zhu, Y., Lu, Y., and Han, S. Longvila: Scaling long-context visual language models for long videos, 2024a. URL <https://arxiv.org/abs/2408.10188>.
- Chen, Z., Wang, W., Cao, Y., Liu, Y., Gao, Z., Cui, E., Zhu, J., Ye, S., Tian, H., Liu, Z., et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024b.
- Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24185–24198, 2024c.
- Cheng, Z., Leng, S., Zhang, H., Xin, Y., Li, X., Chen, G., Zhu, Y., Zhang, W., Luo, Z., Zhao, D., et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Farré, M., Marafioti, A., Tunstall, L., Von Werra, L., and Wolf, T. Finevideo. <https://huggingface.co/datasets/HuggingFaceFV/finevideo>, 2024.
- Fei, J., Li, D., Deng, Z., Wang, Z., Liu, G., and Wang, H. Video-ccam: Enhancing video-language understanding with causal cross-attention masks for short and long videos, 2024. URL <https://arxiv.org/abs/2408.14023>.
- Fu, C., Dai, Y., Luo, Y., Li, L., Ren, S., Zhang, R., Wang, Z., Zhou, C., Shen, Y., Zhang, M., et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. *arXiv preprint arXiv:2405.21075*, 2024a.
- Fu, Y., Panda, R., Niu, X., Yue, X., Hajishirzi, H., Kim, Y., and Peng, H. Data engineering for scaling language models to 128k context. *arXiv preprint arXiv:2402.10171*, 2024b.
- Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., et al. The” something something” video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pp. 5842–5850, 2017.
- Jia, B., Lei, T., Zhu, S.-C., and Huang, S. Egocentric tasks: Understanding human tasks in egocentric videos. *Advances in Neural Information Processing Systems*, 35:3343–3360, 2022.
- Jin, P., Takanobu, R., Zhang, W., Cao, X., and Yuan, L. Chat-univi: Unified visual representation empowers large language models with image and video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13700–13710, 2024.
- KunChang, L., Yinan, H., Yi, W., Yizhuo, L., Wenhai, W., Ping, L., Yali, W., Limin, W., and Yu, Q. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023.
- Kuratov, Y., Bulatov, A., Anokhin, P., Sorokin, D., Sorokin, A., and Burtsev, M. In search of needles in a 11m haystack: Recurrent memory finds what llms miss, 2024. URL <https://arxiv.org/abs/2402.10790>.
- Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.
- Lei, J., Berg, T. L., and Bansal, M. Detecting moments and highlights in videos via natural language queries. *Advances in Neural Information Processing Systems*, 34: 11846–11858, 2021.

- Li, B., Zhang, Y., Guo, D., Zhang, R., Li, F., Zhang, H., Zhang, K., Zhang, P., Li, Y., Liu, Z., and Li, C. Llava-onevision: Easy visual task transfer, 2024a. URL <https://arxiv.org/abs/2408.03326>.
- Li, B., Zhang, Y., Guo, D., Zhang, R., Li, F., Zhang, H., Zhang, K., Zhang, P., Li, Y., Liu, Z., et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024b.
- Li, J., Li, D., Savarese, S., and Hoi, S. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.
- Li, X., Wang, Y., Yu, J., Zeng, X., Zhu, Y., Huang, H., Gao, J., Li, K., He, Y., Wang, C., et al. Videochat-flash: Hierarchical compression for long-context video modeling. *arXiv preprint arXiv:2501.00574*, 2024c.
- Li, Y., Wang, C., and Jia, J. Llama-vid: An image is worth 2 tokens in large language models. 2024d.
- Li, Y., Wang, C., and Jia, J. Llama-vid: An image is worth 2 tokens in large language models. In *European Conference on Computer Vision*, pp. 323–340. Springer, 2025.
- Lin, B., Ye, Y., Zhu, B., Cui, J., Ning, M., Jin, P., and Yuan, L. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023.
- Lin, J., Yin, H., Ping, W., Molchanov, P., Shoeybi, M., and Han, S. Vila: On pre-training for visual language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26689–26699, 2024.
- Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024a.
- Liu, H., Yan, W., Zaharia, M., and Abbeel, P. World model on million-length video and language with blockwise ringattention, 2024b. URL <https://arxiv.org/abs/2402.08268>.
- Liu, J., Wang, Y., Ma, H., Wu, X., Ma, X., Wei, X., Jiao, J., Wu, E., and Hu, J. Kangaroo: A powerful video-language model supporting long-context video input, 2024c. URL <https://arxiv.org/abs/2408.15542>.
- Liu, Z., Dong, Y., Liu, Z., Hu, W., Lu, J., and Rao, Y. Oryx mllm: On-demand spatial-temporal understanding at arbitrary resolution. *arXiv preprint arXiv:2409.12961*, 2024d.
- Liu, Z., Zhu, L., Shi, B., Zhang, Z., Lou, Y., Yang, S., Xi, H., Cao, S., Gu, Y., Li, D., Li, X., Fang, Y., Chen, Y., Hsieh, C.-Y., Huang, D.-A., Cheng, A.-C., Nath, V., Hu, J., Liu, S., Krishna, R., Xu, D., Wang, X., Molchanov, P., Kautz, J., Yin, H., Han, S., and Lu, Y. Nvlla: Efficient frontier visual language models, 2024e. URL <https://arxiv.org/abs/2412.04468>.
- Mangalam, K., Akshulakov, R., and Malik, J. Egoschema: A diagnostic benchmark for very long-form video language understanding. *Advances in Neural Information Processing Systems*, 36:46212–46244, 2023.
- Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., Ginsburg, B., Houston, M., Kuchaiev, O., Venkatesh, G., and Wu, H. Mixed precision training. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=r1gs9JgRZ>.
- Nan, K., Xie, R., Zhou, P., Fan, T., Yang, Z., Chen, Z., Li, X., Yang, J., and Tai, Y. Openvid-1m: A large-scale high-quality dataset for text-to-video generation. *arXiv preprint arXiv:2407.02371*, 2024.
- OpenAI. Hello gpt-4o. In *OpenAI Blog*, 2024. URL <https://openai.com/index/hello-gpt-4o/>.
- Patraucean, V., Smaira, L., Gupta, A., Recasens, A., Markeeva, L., Banarse, D., Koppula, S., Malinowski, M., Yang, Y., Doersch, C., et al. Perception test: A diagnostic benchmark for multimodal video models. *Advances in Neural Information Processing Systems*, 36:42748–42761, 2023.
- Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., and Wei, F. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023a.
- Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., and Wei, F. Kosmos-2: Grounding multimodal large language models to the world, 2023b. URL <https://arxiv.org/abs/2306.14824>.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning transferable visual models from natural language supervision. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 18–24 Jul 2021a. URL <https://proceedings.mlr.press/v139/radford21a.html>.

- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021b.
- Rawal, R., Saifullah, K., Farré, M., Basri, R., Jacobs, D., Somepalli, G., and Goldstein, T. Cinepile: A long video question answering dataset and benchmark. *arXiv preprint arXiv:2405.08813*, 2024.
- Reid, M., Savinov, N., Teplyashin, D., Lepikhin, D., Lillcrap, T., Alayrac, J.-b., Soricut, R., Lazaridou, A., Firat, O., Schrittwieser, J., et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- Shen, X., Xiong, Y., Zhao, C., Wu, L., Chen, J., Zhu, C., Liu, Z., Xiao, F., Varadarajan, B., Bordes, F., et al. Longvu: Spatiotemporal adaptive compression for long video-language understanding. *arXiv preprint arXiv:2410.17434*, 2024.
- Shu, Y., Zhang, P., Liu, Z., Qin, M., Zhou, J., Huang, T., and Zhao, B. Video-xl: Extra-long vision language model for hour-scale video understanding. *arXiv preprint arXiv:2409.14485*, 2024.
- Song, E., Chai, W., Wang, G., Zhang, Y., Zhou, H., Wu, F., Guo, X., Ye, T., Lu, Y., Hwang, J.-N., et al. Moviechat: From dense token to sparse memory for long video understanding. *arXiv preprint arXiv:2307.16449*, 2023.
- Song, E., Chai, W., Wang, G., Zhang, Y., Zhou, H., Wu, F., Chi, H., Guo, X., Ye, T., Zhang, Y., 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*, pp. 18221–18232, 2024.
- Soomro, K., Zamir, A. R., and Shah, M. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- Wang, P., Bai, S., Tan, S., Wang, S., Fan, Z., Bai, J., Chen, K., Liu, X., Wang, J., Ge, W., Fan, Y., Dang, K., Du, M., Ren, X., Men, R., Liu, D., Zhou, C., Zhou, J., and Lin, J. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024a.
- Wang, W., He, Z., Hong, W., Cheng, Y., Zhang, X., Qi, J., Gu, X., Huang, S., Xu, B., Dong, Y., et al. Lvbench: An extreme long video understanding benchmark. *arXiv preprint arXiv:2406.08035*, 2024b.
- Wang, Y., He, Y., Li, Y., Li, K., Yu, J., Ma, X., Li, X., Chen, G., Chen, X., Wang, Y., et al. Internvid: A large-scale video-text dataset for multimodal understanding and generation. *arXiv preprint arXiv:2307.06942*, 2023.
- Wu, B. and Star, S. Y. A benchmark for situated reasoning in real-world videos. *NeurIPS Datasets and Benchmarks*, 2021.
- Wu, H., Li, D., Chen, B., and Li, J. Longvideobench: A benchmark for long-context interleaved video-language understanding, 2024. URL <https://arxiv.org/abs/2407.15754>.
- Wu, P., Escontrela, A., Hafner, D., Abbeel, P., and Goldberg, K. Daydreamer: World models for physical robot learning. In Liu, K., Kulic, D., and Ichnowski, J. (eds.), *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pp. 2226–2240. PMLR, 14–18 Dec 2023. URL <https://proceedings.mlr.press/v205/wu23c.html>.
- Wu, Z., Xiong, C., Ma, C.-Y., Socher, R., and Davis, L. S. Adaframe: Adaptive frame selection for fast video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1278–1287, 2019.
- Xiao, J., Shang, X., Yao, A., and Chua, T.-S. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9777–9786, June 2021.
- Xu, L., Zhao, Y., Zhou, D., Lin, Z., Ng, S. K., and Feng, J. Pillava: Parameter-free llava extension from images to videos for video dense captioning. *arXiv preprint arXiv:2404.16994*, 2024.
- Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Li, C., Liu, D., Huang, F., Wei, H., et al. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*, 2024a.
- Yang, D., Huang, S., Lu, C., Han, X., Zhang, H., Gao, Y., Hu, Y., and Zhao, H. Vript: A video is worth thousands of words. *Advances in Neural Information Processing Systems*, 37:57240–57261, 2024b.
- Yao, Y., Yu, T., Zhang, A., Wang, C., Cui, J., Zhu, H., Cai, T., Li, H., Zhao, W., He, Z., et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024.
- Ye, J., Xu, H., Liu, H., Hu, A., Yan, M., Qian, Q., Zhang, J., Huang, F., and Zhou, J. mplug-owl3: Towards long image-sequence understanding in multi-modal large language models. *arXiv preprint arXiv:2408.04840*, 2024.

- Ye, Q., Xu, H., Xu, G., Ye, J., Yan, M., Zhou, Y., Wang, J., Hu, A., Shi, P., Shi, Y., et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
- Yi, K., Gan, C., Li, Y., Kohli, P., Wu, J., Torralba, A., and Tenenbaum, J. B. Clevrer: Collision events for video representation and reasoning. *arXiv preprint arXiv:1910.01442*, 2019.
- Zala, A., Cho, J., Kottur, S., Chen, X., Oguz, B., Mehdad, Y., and Bansal, M. Hierarchical video-moment retrieval and step-captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23056–23065, 2023.
- Zhai, X., Mustafa, B., Kolesnikov, A., and Beyer, L. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 11975–11986, 2023.
- Zhang, H., Li, X., and Bing, L. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023. URL <https://arxiv.org/abs/2306.02858>.
- Zhang, P., Zhang, K., Li, B., Zeng, G., Yang, J., Zhang, Y., Wang, Z., Tan, H., Li, C., and Liu, Z. Long context transfer from language to vision. *arXiv preprint arXiv:2406.16852*, 2024a.
- Zhang, P., Zhang, K., Li, B., Zeng, G., Yang, J., Zhang, Y., Wang, Z., Tan, H., Li, C., and Liu, Z. Long context transfer from language to vision, 2024b. URL <https://arxiv.org/abs/2406.16852>.
- Zhang, P., Zhang, K., Li, B., Zeng, G., Yang, J., Zhang, Y., Wang, Z., Tan, H., Li, C., and Liu, Z. Long context transfer from language to vision. *arXiv preprint arXiv:2406.16852*, 2024c. URL <https://arxiv.org/abs/2406.16852>.
- Zhang, R., Gui, L., Sun, Z., Feng, Y., Xu, K., Zhang, Y., Fu, D., Li, C., Hauptmann, A., Bisk, Y., et al. Direct preference optimization of video large multimodal models from language model reward. *arXiv preprint arXiv:2404.01258*, 2024d.
- Zhang, S., Fang, Q., Yang, Z., and Feng, Y. Llava-mini: Efficient image and video large multimodal models with one vision token, 2025. URL <https://arxiv.org/abs/2501.03895>.
- Zhang, Y., Wu, J., Li, W., Li, B., Ma, Z., Liu, Z., and Li, C. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*, 2024e.
- Zhi, Y., Tong, Z., Wang, L., and Wu, G. Mgsampler: An explainable sampling strategy for video action recognition. In *Proceedings of the IEEE/CVF International conference on Computer Vision*, pp. 1513–1522, 2021.
- Zhou, J., Shu, Y., Zhao, B., Wu, B., Xiao, S., Yang, X., Xiong, Y., Zhang, B., Huang, T., and Liu, Z. Mlvu: A comprehensive benchmark for multi-task long video understanding. *arXiv preprint arXiv:2406.04264*, 2024.

## A. Experiment Details

### A.1. Benchmarks

**LVBench (Wang et al., 2024b)** consists of publicly sourced long videos with an average duration of approximately 4,101 seconds, significantly longer than existing datasets. The videos are categorized into six major categories and 21 subcategories, covering a wide range of topics such as sports, live streams, TV shows, documentaries, and animations. The dataset is annotated through a combination of manual effort and model assistance.

**EgoSchema (Mangalam et al., 2023)** is a diagnostic benchmark aimed at assessing the long-form video-language understanding capabilities of modern multimodal systems. Derived from the Ego4D dataset, EgoSchema consists of over 5000 human-curated multiple-choice question-answer pairs, covering more than 250 hours of high-quality egocentric video data (180 seconds in average). The dataset is designed to require models to understand and reason over extended temporal contexts, with each question based on a three-minute video clip.

**LongVideoBench (Wu et al., 2024)** comprises 3,763 videos with varying lengths, ranging from 8 seconds to 1 hour, and 6,678 human-annotated multiple-choice questions. The videos cover a wide range of themes, including movies, news, life and knowledge. The benchmark introduces a novel task called "referring reasoning" to address the single-frame bias in existing video understanding metrics. This task requires models to process more frames to improve performance.

**MLVU (Zhou et al., 2024)** consists of 2,593 samples across 9 task categories, designed to comprehensively assess MLLMs' capabilities in long video understanding. The tasks are categorized into holistic LVU (e.g., topic reasoning, anomaly recognition), single-detail LVU (e.g., needle question answering, ego reasoning), and multi-detail LVU (e.g., action order, action count), including both multiple-choice and free-form generation formats. Videos range from 3 minutes to over 2 hours, with an average length of 651 seconds.

**VideoMME (Fu et al., 2024a)** consists of 900 videos spanning 6 primary visual domains (Knowledge, Film & Television, Sports Competition, Artistic Performance, Life Record, and Multilingual) with 30 subfields. The videos vary in length from 11 seconds to 1 hour, covering short, medium, and long durations. Each video is annotated with 3 multiple-choice questions, resulting in a total of 2,700 question-answer pairs. It also includes subtitles and audio tracks for 744 and 900 videos.

### A.2. Baselines

#### A.2.1. PROPRIETARY VISION-LANGUAGE MODELS

**GPT4-V (Achiam et al., 2023)** is an advanced multimodal AI model developed by OpenAI, combining the power of natural language processing with visual understanding. Building on the capabilities of GPT-4, GPT-4-V integrates vision-based functionalities, enabling it to generate responses based on both text and visual inputs.

**GPT4-o (OpenAI, 2024)** is the latest iteration in the Generative Pre-trained Transformer series. Building on the strengths of its predecessors, GPT-4-o offers enhanced capabilities in understanding multimodal inputs.

**Gemini-1.5-Pro (Reid et al., 2024)** is Google's latest breakthrough in artificial intelligence, representing a significant leap in multi-modal and long-context understanding.

#### A.2.2. OPEN-SOURCE MULTI-IMAGE VISION-LANGUAGE MODELS

**LLaVA-OneVision (Li et al., 2024b)** leverages strong transfer learning across different modalities, pushing performance boundaries in tasks such as video understanding and complex reasoning through a unified model architecture.

**Oryx-1.5 (Liu et al., 2024d)** is a novel multi-modal large language model (MLLM) designed for on-demand spatial-temporal understanding of diverse visual inputs, including images, videos, and 3D scenes. It processes visual inputs at arbitrary resolutions and temporal lengths through a pre-trained OryxViT model and a dynamic compressor module.

**InternVL Series (Chen et al., 2024c;b)** is characterized by its large-scale visual encoder, InternViT-6B, and language middleware, QLLaMA, which enables it to handle a wide range of vision-language tasks. InternVL 2.5 is the latest version

in the InternVL series, and it significantly enhances its performance through improvements in training and testing strategies, as well as data quality. The model demonstrates outstanding performance across multiple benchmarks.

**Chat-Univi (Jin et al., 2024)** uses of dynamic visual tokens to represent images and videos in a unified framework, allowing it to capture spatial details of images and temporal relationships of videos efficiently. Additionally, the model employs a multi-scale representation that enables it to perceive both high-level semantic concepts and low-level visual details.

**MiniCPM-v2.6 (Yao et al., 2024)** achieves strong performance comparable to GPT-4V while being optimized for efficiency. It has high-resolution image perception, strong OCR capabilities, multilingual support for over 30 languages, and low hallucination rates, making it suitable for real-world applications with limited computational resources.

**mPLUG-Owl3 (Ye et al., 2024)** is a versatile multi-modal large language model designed to enhance long image-sequence understanding. It introduces Hyper Attention blocks, which integrate vision and language efficiently by allowing parallel cross-attention and self-attention within the transformer architecture.

**Qwen2-VL (Wang et al., 2024a)** is an advanced vision-language model that enhances the perception of visual information at any resolution. Qwen2-VL introduces Naive Dynamic Resolution mechanism, which allows the model to process images of varying resolutions into different numbers of visual tokens, and the Multimodal Rotary Position Embedding (M-RoPE), which effectively fuses positional information across text, images, and videos.

**NVILA (Liu et al., 2024e)** builds on VILA and enhances its architecture by scaling up spatial and temporal resolutions before compressing visual tokens, enabling efficient processing of high-resolution images and long videos. Additionally, NVILA incorporates system-wide optimizations, including dataset pruning, FP8 training, and specialized inference engines, to achieve significant improvements in training speed, memory usage, and inference efficiency, while maintaining competitive performance across various benchmarks.

#### A.2.3. OPEN-SOURCE VIDEO VISION-LANGUAGE MODELS

**LLaVA-Video (Zhang et al., 2024e)** is trained on a high-quality synthetic dataset LLaVA-Video-178K. It processes frames using an optimized video representation technique called LLaVA-Video SlowFast, which allows for more detailed temporal understanding compared to previous models.

**LLaMA-VID (Li et al., 2024d)** represents each frame with two distinct tokens—a context token that captures the overall image context based on user input, and a content token that encapsulates visual details. This dual-token strategy significantly reduces the computational burden of processing long videos while preserving critical information, enabling VLMs to handle hour-long videos.

**Video-XL (Shu et al., 2024)** leverages the inherent key-value sparsification capability of large language models to condense visual inputs through a new special token called the Visual Summarization Token (VST). This approach enables the model to handle long videos efficiently by summarizing visual information into compact representations, while dynamic compression and curriculum learning strategies further enhance its performance and training effectiveness.

**Video-LLaVA (Lin et al., 2023)** uses a LanguageBind encoder to align visual signals from images and videos into a unified visual feature space, followed by joint training of images and videos to enhance multi-modal interactions.

**VideoLLaMA2 (Cheng et al., 2024)** is designed to enhance spatial-temporal modeling and audio understanding in video and audio-oriented tasks. It introduces a Spatial-Temporal Convolution (STC) connector that effectively captures the intricate spatial and temporal dynamics of video data, and an integrated Audio Branch that enriches the model’s multimodal understanding capabilities by seamlessly incorporating audio cues.

**VideoChat2 (KunChang et al., 2023)** utilizes the UMT-L visual encoder, which is specifically designed for spatiotemporal representation learning. It also adopt Q-former to compresses redundant visual tokens into fewer tokens and aligns them

with text tokens. VideoChat2 excels in tasks that require understanding of temporal sequences, such as action prediction, action sequence, and moving direction.

**Video-CCAM (Fei et al., 2024)** is designed to enhance video-language understanding for both short and long videos. It uses Causal Cross-Attention Masks (CCAMs) within the cross-attention layers, which allows the model to effectively process a large number of visual tokens while preserving temporal order in videos.

**PLLaVA (Xu et al., 2024)** is a novel and efficient approach to adapt existing image-language pre-trained models for dense video understanding. PLLaVA employs a simple yet effective pooling strategy to smooth the feature distribution along the temporal dimension, reducing the impact of high-norm visual features that often lead to performance degradation.

**Movie-Chat (Song et al., 2024)** introduces a innovative memory mechanism inspired by the Atkinson-Shiffrin memory model, which employs a combination of short-term and long-term memory to efficiently process and understand long videos.

**Kangaroo (Liu et al., 2024c)** is designed for long-context video input. Kangaroo addresses the challenges of limited training data and excessive compression of visual features in long videos by developing a data curation system to build a large-scale high-quality dataset for vision-language pre-training and instruction tuning. It also employs a curriculum training pipeline that gradually increases the resolution and number of input frames to enhance the model’s ability to process long videos.

**LongVU (Shen et al., 2024)** introduces a spatiotemporal adaptive compression mechanism designed for long video-language understanding. LongVU leverages cross-modal queries and inter-frame dependencies to reduce the number of video tokens while preserving visual details, making it possible to process hour-long videos within the context length constraints of large language models (LLMs).

**LongVA (Zhang et al., 2024c)** leverages a unique approach called "long context transfer," where it extends the context length of a language model by training it on longer text data and then aligns this extended model with visual inputs using short image data. This enables LongVA to process over 2,000 frames or more than 200K visual tokens without explicit long video training.

**LongVILA (Anonymous, 2024)** adopts a five-stage training pipeline that incorporates multi-modal alignment, large-scale pre-training, supervised fine-tuning, context extension, and long video fine-tuning. Additionally, LongVILA utilizes a novel Multi-Modal Sequence Parallelism (MM-SP) system, which efficiently parallelizes training and inference for long videos.

### A.3. Implement Details

Stage	Datasets	Size
Pretrain	WebVid (Bain et al., 2021), InternVid (Wang et al., 2023), ShareGPTVideo (Zhang et al., 2024d), OpenVid (Nan et al., 2024), Vript (Yang et al., 2024b)	7.4M
Short Video Understanding	VideoChat2-IT (KunChang et al., 2023), EgoTaskQA (Jia et al., 2022), CLEVRER (Yi et al., 2019), LLaVA-Video-178K (Zhang et al., 2024e), MovieChat (Song et al., 2023), PerceptionTest (Patraucean et al., 2023), STAR (Wu & Star, 2021), NExTQA (Xiao et al., 2021)	1.3M
Longer Video Understanding	FineVideo (Farré et al., 2024), CinePile (Rawal et al., 2024)	500K

**Training Datasets** We employ a three-stage training strategy to fine-tune ViLAMP and enhance its video comprehension capabilities. During the pretraining phase, the primary objective is to adapt ViLAMP to the interleaved input structure combining keyframes and non-keyframes. Our training data comprises approximately 7.4M video-caption pairs collected from multiple open-source datasets including WebVid (Bain et al., 2021), InternVid (Wang et al., 2023), ShareGPTVideo (Zhang et al., 2024d), OpenVid (Nan et al., 2024), and Vript (Yang et al., 2024b), ensuring broad coverage of diverse visual scenarios. This comprehensive collection trains the model to develop a foundational understanding of video inputs.

Additionally, we further expand the data scale and enhance the model’s capabilities by incorporating diverse types of QA datasets, including VideoChat2-IT (KunChang et al., 2023), EgoTaskQA (Jia et al., 2022), CLEVRER (Yi et al., 2019), LLaVA-Video-178K (Zhang et al., 2024e), MovieChat (Song et al., 2023), PerceptionTest (Patraucean et al., 2023), HiREST (Zala et al., 2023), STAR (Wu & Star, 2021), and NExTQA (Xiao et al., 2021), ultimately constructing a dataset of 1.3M samples. This dataset improves the model’s instruction-following ability, enabling it to better adapt to various downstream tasks. Specifically, for LLaVA-Video-178K, we utilize the *academic* and *youtube* subsets.

Finally, to improve the model’s performance on longer videos, we incorporate FineVideo (Farré et al., 2024) and CinePile (Rawal et al., 2024), two datasets with longer average video durations (280s and 160s, respectively), and get approximately 500K training samples.

**Training Parameters** As mentioned in § B.2, we configure  $\alpha = 10^{-2}$  and  $\tau = 0.85$  to achieve balanced effectiveness. During training, we set learning rate to  $2 \times 10^{-6}$  for the visual encoder and  $10^{-5}$  for the remaining components. We use adamW as the optimizer. The optimization follows a cosine learning rate scheduler with a warmup ratio of 0.03. We set the batch size to 1 and the gradient accumulation steps to 4. We train the model for 1 epoch, which costs approximately two weeks with 32 NVIDIA A100 GPUs.

## B. Discussions

### B.1. Computational Efficiency Result

We compare against four representative video-language models: LongVA, LLaVA-Video, LLaMA-VID, and VideoChat-Flash. As shown in Table 5, ViLAMP demonstrates a significant advantage in memory efficiency. Unlike typical models whose memory usage increases drastically with input length, ViLAMP shows only a marginal increase in memory demand as the input length grows. This characteristic enables ViLAMP to handle extremely long inputs under limited memory constraints, which can be attributed to the effective removal of redundant information through our differential distillation. Furthermore, ViLAMP exhibits a slow growth rate of FLOPs with input expansion, demonstrating remarkable superiority in handling long sequences. When the input length reaches the thousand-frame level, ViLAMP outperforms all baseline models, requiring merely 18.4% of the FLOPs consumed by VideoChat-Flash when processing 8,192 frames.

Moreover, ViLAMP also demonstrates a significant speed advantage, with TTFT results indicating that it exhibits a response latency that is merely 50% or less of typical models. When the number of frames reaches 8,192, ViLAMP exhibits marginally slower performance compared to VideoChat-Flash. This is attributed to the pruning operations implemented in the Large Language Model (LLM) component of VideoChat-Flash, which specifically optimizes inference speed. Nevertheless, ViLAMP achieves comparable performance levels without employing such optimizations.

### B.2. Influence of Hyperparameter Configuration

We conduct a series of experiments under various settings of  $\alpha$  and  $\tau$ . As reported in Table 6, the overall performance of the model first improves and then decreases as  $\tau$  increases. This trend can be attributed to the fact that a too small  $\tau$  leads to insufficient sampling, while an excessively large  $\tau$  results in overly concentrated sampling, both of which negatively impact the model’s effectiveness. We find that setting  $\tau = 0.85$  provides a robust balance between semantic preservation and temporal diversity. Further analysis of  $\alpha$  reveals a similar pattern. To some extent, reducing  $\alpha$  allows the weight distribution to focus more on the target tokens, which helps in eliminating redundant information. However, when  $\alpha$  is too small, the weights concentrate on only a few tokens, leading to information loss that adversely affects the model’s performance. Through extensive experimentation, we find that setting  $\alpha = 10^{-2}$  ensures robust performance across diverse video content.

Table 5. Computational efficiency comparison across different input lengths.

Model	FLOPs (T)	Memory (MB)	TTFT (ms)
<i>Input 64 Frames</i>			
LongVA	316.0	62,262	798
LLaVA-Video	101.8	32,362	203
LLaMA-VID	112.3	46,402	373
VideoChat-Flash	<b>29.6</b>	<b>18,688</b>	521
ViLAMP	61.0	23,201	<b>80</b>
<i>Input 256 Frames</i>			
LongVA	2,946.3	<i>OOM</i>	<i>OOM</i>
LLaVA-Video	1,221.6	71,495	535
LLaMA-VID	245.9	47,294	492
VideoChat-Flash	<b>113.3</b>	<b>22,665</b>	586
ViLAMP	120.3	23,647	<b>157</b>
<i>Input 1,024 Frames</i>			
LongVA	27,605.6	<i>OOM</i>	<i>OOM</i>
LLaVA-Video	18,324.0	<i>OOM</i>	<i>OOM</i>
LLaMA-VID	1,045.3	48,678	785
VideoChat-Flash	489.6	38,427	694
ViLAMP	<b>356.8</b>	<b>24,339</b>	<b>462</b>
<i>Input 8,192 Frames</i>			
LongVA	2,153,901.7	<i>OOM</i>	<i>OOM</i>
LLaVA-Video	1,643,892.3	<i>OOM</i>	<i>OOM</i>
LLaMA-VID	20,906.0	91,438	5,080
VideoChat-Flash	13,896.6	92,034	<b>2,730</b>
ViLAMP	<b>2,565.2</b>	<b>45,819</b>	3,374

Table 6. Performance varying with hyperparameters  $\alpha$  and  $\tau$ .

$\alpha$	$\tau$	LongVideoBench	MLVU	Video-MME
1	0.35	55.6	65.9	59.4
	0.50	58.3	69.5	61.2
	0.85	<b>58.9</b>	<b>70.5</b>	<b>62.7</b>
	1.00	51.3	64.5	58.0
$10^{-1}$	0.35	57.2	63.3	60.5
	0.50	58.0	68.9	62.6
	0.85	<b>59.1</b>	<b>71.5</b>	<b>64.5</b>
	1.00	51.1	63.6	57.3
$10^{-2}$	0.35	56.8	65.3	60.2
	0.50	59.6	69.7	63.2
	0.85	<b>60.5</b>	<b>72.6</b>	<b>65.1</b>
	1.00	50.9	62.3	58.8
$10^{-3}$	0.35	49.3	61.6	45.6
	0.50	52.2	63.4	47.0
	0.85	<b>54.4</b>	<b>65.6</b>	<b>48.2</b>
	1.00	46.5	60.9	43.6