
∞ -VIDEO: A Training-Free Approach to Long Video Understanding via Continuous-Time Memory Consolidation

Saul Santos^{1,2} António Farinhas^{1,2} Daniel C. McNamee³ André F. T. Martins^{1,2,4,5}

Abstract

Current video-language models struggle with long-video understanding due to limited context lengths and reliance on sparse frame subsampling, often leading to information loss. This paper introduces ∞ -VIDEO, which can process arbitrarily long videos through a continuous-time long-term memory (LTM) consolidation mechanism. Our framework augments video Q-formers by allowing them to process unbounded video contexts efficiently and without requiring additional training. Through continuous attention, our approach dynamically allocates higher granularity to the most relevant video segments, forming “sticky” memories that evolve over time. Experiments with Video-LLaMA and VideoChat2 demonstrate improved performance in video question-answering tasks, showcasing the potential of continuous-time LTM mechanisms to enable scalable and training-free comprehension of long videos.

1. Introduction

Multimodal large language models have driven progress in video-language tasks through the integration of pretrained visual encoders with powerful text-based models (Li et al., 2023b; Zhang et al., 2023b; Cheng et al., 2024; Li et al., 2024). However, current video models are constrained by short context lengths (Li et al., 2023b; Maaz et al., 2024; Liu et al., 2023) and often rely on sparse frame subsampling for longer sequences, which limits their ability to fully process and understand long videos. This contrasts with the high-capacity persistence of human memory (Brady et al., 2008) and the cognitive principles by which humans store information over long timescales, which involve consoli-

ation processes adaptively integrating important episodic events into long-term memory (McGaugh, 2013; Cowan et al., 2021). *How can we ensure models are able to fully understand and grasp information from arbitrarily long videos while they process them without losing critical details?*

Transformers offer significant potential for extracting spatio-temporal features from videos (Zhang et al., 2023b; Li et al., 2024). While recent video-language models have prioritized simpler methods such as projection layers (Li et al., 2023d; Liu et al., 2023; Li et al., 2023b; Liu et al., 2024c; Ye et al., 2024) and temporal pooling (Luo et al., 2023; Maaz et al., 2024) for the sake of efficiency and scalability, these approaches often sacrifice transformer’s representational depth. Additionally, training video-language models presents significant difficulties, whereas traditional approaches mainly focus on scaling model parameters (Liu et al., 2024c; Cheng et al., 2024; Li et al., 2024), which requires substantial computational resources. A recent alternative explored the usage of additional computation during inference (Zhang et al., 2023a; Wang et al., 2024a;c). However, these methods assume that the spatio-temporal module is intrinsic to the LLM, which limits their ability to effectively specialize in capturing spatio-temporal content. Furthermore, these methods often rely on subsampling and are designed to process the entire video whenever they need to answer a question.

In this paper, we take an alternative approach inspired by human cognition, where memory consolidation processes enable the retention and efficient handling of long-term dependencies (Frankland & Bontempi, 2005; Preston & Eichenbaum, 2013; Song et al., 2023; 2024; Balazevic et al., 2024). Namely, we develop a new framework with **continuous-time visual memory representations**. Our framework adapts the ∞ -former architecture previously developed for textual data (Martins et al., 2020; 2022b), which we leverage to extend the capabilities of pre-trained short-context multimodal LLMs, making them able to process unbounded video contexts in a training-free manner. Shifting from a discrete to a continuous attention framework parallels the recent evolution in theories of human working memory mediated by prefrontal cortex—from the discrete “slot-based” model to the continuous “shared resource” approach (Ma et al., 2014). Consequently, our method aims to cultivate

¹Instituto de Telecomunicações ²Instituto Superior Técnico, Universidade de Lisboa ³Champalimaud Research ⁴ELLIS Unit Lisbon ⁵Unbabel. Correspondence to: Saul Santos <saul.r.santos@tecnico.ulisboa.pt>.

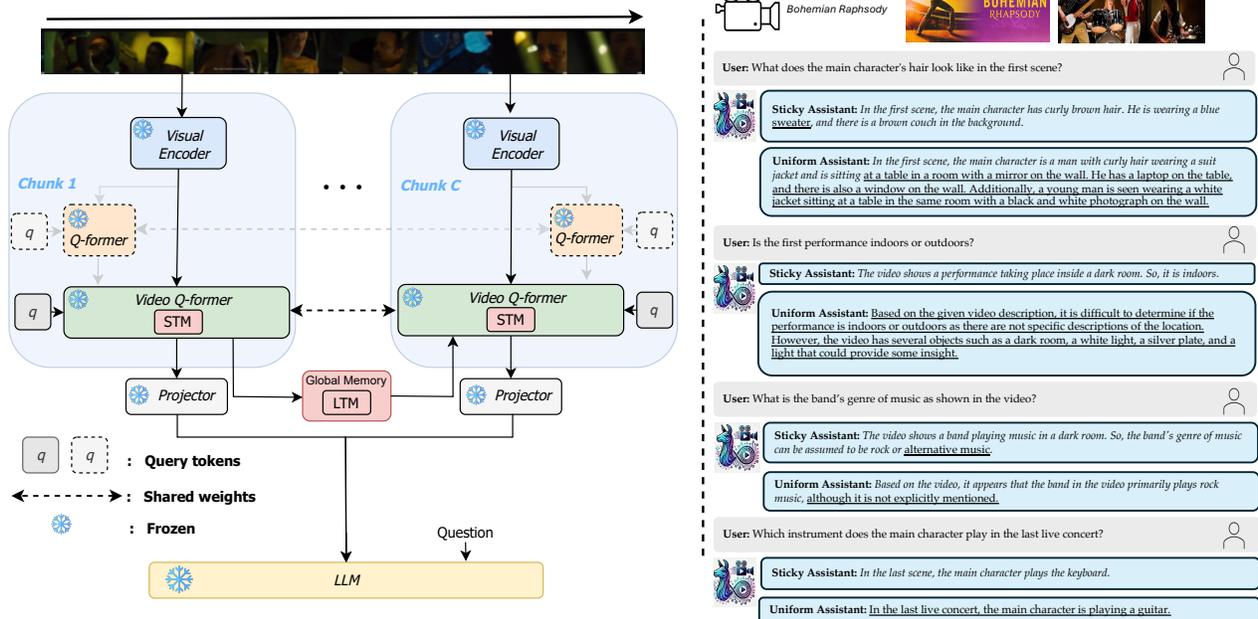


Figure 1. (Left) Overview of ∞-VIDEO (our approach) using Video LLaMA (Zhang et al., 2023b, gray arrows), which uses an additional spatial Q-former module, and VideoChat2 (Li et al., 2024, black arrows). We split the video into frame chunks and apply these models to each chunk. The Video Q-former module combines a weighted average of the STM, which is the attention for an individual chunk, with a continuous LTM that takes into account previous chunks. The outputs of the Video Q-Former are projected and then averaged. The LLM takes as input visual tokens, generated by our modified video Q-former, alongside the corresponding question to obtain the answer. (Right) Examples of ∞-Video LLaMA answers, equipped with our LTM, with uniform sampling and sticky memories for short and ultra-long videos. *Italicized* corresponds to the correct answer, while underlined corresponds to a wrong answer or hallucination.

similar insights within the realm of episodic memory processing, where dynamic handling of memory is essential. This leads to ∞-VIDEO models (Fig. 1), which are able to process and organize information as it comes, with only one pass over the video. Our main contributions are:¹

- We equip the current attention mechanism of video Q-formers (short-term memory, STM) with a continuous-time LTM that consolidates video information by dynamically allocating higher granularity to the most relevant parts of a video.
- We develop a new continuous-time attention mechanism which is more powerful than the Gaussian model of Martins et al. (2022b) by considering the Gibbs density based on the continuous query-key similarity function.
- We show that architectures with spatio-temporal feature extractors designed for short videos can generalize to long-video understanding in a simple, training-free manner without requiring task-specific fine-tuning or training on

¹The code is made available in <https://github.com/deep-spin/Infinite-Video>.

long-video datasets.

- We validate our approach using the Video-LLaMA (Zhang et al., 2023b) and VideoChat2 (Li et al., 2024) models by processing a stream of video frames with a single pass, where the STM considers one chunk at the time, and the LTM maintains global information from past chunks. Our proposed model shows improved performance on video question-answering tasks when using the LTM and competitive results with other training-free models.

2. Background

2.1. Discrete Attention

Attention mechanisms (Bahdanau et al., 2015) act as a memory component in modern neural networks, enabling them to dynamically focus on key parts of the input and capture long-range dependencies, improving performance across various tasks (Vaswani et al., 2017; Dosovitskiy et al., 2021).

Consider two sequences $\mathbf{X} \in \mathbb{R}^{L \times e}$ and $\mathbf{Y} \in \mathbb{R}^{R \times e}$, where L and R are the sequence lengths and e is the embedding size. A vanilla attention mechanism in a transformer works

as follows. First, we obtain queries (\mathbf{Q}), keys (\mathbf{K}), and values (\mathbf{V}) by linearly projecting \mathbf{X} and \mathbf{Y} for each attention head h :

$$\mathbf{Q}^h = \mathbf{Y}\mathbf{W}_Q^h, \quad \mathbf{K}^h = \mathbf{X}\mathbf{W}_K^h, \quad \mathbf{V}^h = \mathbf{X}\mathbf{W}_V^h, \quad (1)$$

where $\mathbf{W}_Q^h \in \mathbb{R}^{e \times d}$, $\mathbf{W}_K^h \in \mathbb{R}^{e \times d}$, and $\mathbf{W}_V^h \in \mathbb{R}^{e \times d}$ are head-specific learnable projection matrices, $d = e/|h|$, and $|h|$ is the number of attention heads. For each head, the context representation $\mathbf{Z}^h \in \mathbb{R}^{L \times d}$ is computed as:

$$\mathbf{Z}^h = \text{softmax} \left(\frac{\mathbf{Q}^h (\mathbf{K}^h)^\top}{\sqrt{d}} \right) \mathbf{V}^h. \quad (2)$$

The outputs from all heads are then concatenated to obtain the final context representation $\mathbf{Z} \in \mathbb{R}^{L \times e}$:

$$\mathbf{Z} = [\mathbf{Z}^1 \quad \mathbf{Z}^2 \quad \dots \quad \mathbf{Z}^{|h|}] \mathbf{W}_Z, \quad (3)$$

where $\mathbf{W}_Z \in \mathbb{R}^{e \times e}$ is another learnable projection matrix.

2.2. Continuous Attention

Instead of splitting the input object into a finite set of pieces (e.g., tokens in text or pixels in images), continuous attention mechanisms (Martins et al., 2020) assume an underlying continuous domain, suitable for arbitrarily long temporal signals, such as audio or video data. This is done by replacing the attention probability mass function by a probability density function (PDF) over a continuous signal.

In continuous attention, the input is assumed to be a continuous signal $\mathbf{x}(t)$. Although video data is ‘‘continuous-time’’ in nature, it comes as a stream of L discrete frames $\mathbf{X} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_L^\top] \in \mathbb{R}^{L \times e}$, and therefore it is necessary to convert this sequence into a smooth continuous signal. This can be done by expressing the continuous signal $\mathbf{x}(t) \in \mathbb{R}^e$ as a linear combination of N basis functions $\psi(t) \in \mathbb{R}^N$:

$$\mathbf{x}(t) = \mathbf{B}^\top \psi(t), \quad (4)$$

where $\mathbf{B} \in \mathbb{R}^{N \times e}$ is a coefficient matrix. For compression, it is appealing to use a smaller number of basis functions than frames, $N \ll L$. \mathbf{B} can be computed with multivariate ridge regression (Brown & Zidek, 1980), where we set the domain of the continuous-time signal to the unit interval $[0, 1]$. Frames are associated with time instants in this unit interval, $t_1 \leq t_2 \leq \dots \leq t_L$, with each $t_\ell \in [0, 1]$, and we set our design matrix as $\mathbf{F} = [\psi(t_1), \dots, \psi(t_L)] \in \mathbb{R}^{N \times L}$. The coefficients \mathbf{B} are computed such that $\mathbf{x}(t_\ell) \approx \mathbf{x}_\ell$ for each frame $\ell \in \{1, \dots, L\}$, with $\lambda > 0$, leading to:

$$\mathbf{B}^\top = \mathbf{X}^\top \mathbf{F}^\top (\mathbf{F} \mathbf{F}^\top + \lambda \mathbf{I})^{-1}. \quad (5)$$

The final step involves attending to $\mathbf{x}(t)$. In this approach, a PDF $p(t)$ replaces the probability mass function of the

discrete attention.² The context is then computed as the expected value of the values $\mathbf{v}(t) = (\mathbf{W}_V)^\top \mathbf{x}(t)$:

$$\mathbf{Z} = \mathbb{E}_p[\mathbf{v}(t)]. \quad (6)$$

∞ -former. In a transformer with discrete attention, handling a long context (large L) becomes impractical due to excessive memory demands. The ∞ -former (Martins et al., 2022b) overcomes this limitation by means of an unbounded LTM leveraging continuous attention (§2.2). It allows for unbounded context without increasing memory usage by trading off the number of basis functions that fit into memory with the granularity of their representations. It achieves this by sampling points within the $[0, 1]$ interval, either uniformly or based on prior attention, at which $\mathbf{x}(t)$ is evaluated. This process, as we will see in §3, can be seen as the memory consolidation step, allowing new information from the short-term memory $\mathbf{x}(t)$ to be incorporated by scaling it down with a forgetting factor τ . The past context is associated with positions in $[0, \tau]$, while the new context is associated in $(\tau, 1]$, followed by ridge regression over the new $\mathbf{x}(t)$ and computation of the output context as in Eq. 6.

2.3. Video Q-former

A video Q-former (Zhang et al., 2023b) is a specialized variant of a transformer architecture designed to enable LLMs to process and understand video content. Developed initially to capture spatial features in images (Li et al., 2023a), its primary function is to map $L \times P$ spatial embeddings, where P can be the number of patch vectors from a visual encoder, to R spatio-temporal video representations. It operates by stacking layers that first apply self-attention between R learned queries, followed by cross-attention between this self-attention output and the $L \times P$ spatial embeddings, leading to R spatio-temporal representations.

3. Unbounded Memory Video Q-former

In this section, we describe our approach to endow video models with a continuous-time memory mechanism. We adapt existing video Q-former models (Zhang et al., 2023b; Li et al., 2024) by first splitting the full sequence of frames into chunks and then processing each chunk individually. Each chunk includes a discrete STM, which is the already present cross-attention in that chunk. However, building the LTM assumes continuity within the embeddings, which is not the case since each frame contains P distinct embeddings. To address this, we perform average pooling over the P embeddings, resulting in a discrete sequence $\mathbf{X} \in \mathbb{R}^{M \times e}$, where M denotes the number of frames in the chunk. We

²For instance, in Martins et al. (2022b), the PDF is modeled as a Gaussian distribution $\mathcal{N}(t; \mu, \sigma^2)$, where the parameters μ and σ^2 are learned by a neural network. In §3, we propose a different, more powerful strategy.

introduce a global and dynamic LTM, which works with a modified, more powerful version of the continuous attention mechanism in §2.2. The LTM update enables increased granularity in memory regions with higher cumulative attention density.

3.1. Long-Term Memory

The first step in building our continuous LTM is to project the long-term continuous input $\mathbf{x}(t) = \mathbf{B}^\top \boldsymbol{\psi}(t)$, leading to the continuous keys $\mathbf{k}^h(t) \in \mathbb{R}^d$ and values $\mathbf{v}^h(t) \in \mathbb{R}^d$, with the same projection matrices (1) of the STM. This allows us to perform attention over the same embedding space as the vanilla video Q-former as:

$$\mathbf{k}^h(t) = (\mathbf{W}_K^h)^\top \mathbf{x}(t) = (\mathbf{W}_K^h)^\top \mathbf{B}^\top \boldsymbol{\psi}(t), \quad (7)$$

$$\mathbf{v}^h(t) = (\mathbf{W}_V^h)^\top \mathbf{x}(t) = (\mathbf{W}_V^h)^\top \mathbf{B}^\top \boldsymbol{\psi}(t). \quad (8)$$

We compute the query $\mathbf{Q}^h = [\mathbf{q}_1^\top, \dots, \mathbf{q}_R^\top] = \mathbf{Y} \mathbf{W}_Q^h$ as in (1). For each query $\mathbf{q}_i \in \mathbb{R}^d$, we compute the continuous query-key similarity $s_i^h(t)$ for each head and query as

$$s_i^h(t) = \mathbf{q}_i^\top \mathbf{k}(t) = \mathbf{q}_i^\top (\mathbf{W}_K^h)^\top \mathbf{B}^\top \boldsymbol{\psi}(t), \quad (9)$$

and compute a Gibbs PDF as

$$p_i^h(t) = \frac{\exp(s_i^h(t))}{\int \exp(s_i^h(t')) dt'}, \quad (10)$$

where the integral is approximated with the trapezoidal rule. Given the value function $\mathbf{v}^h(t)$, we compute the attention-specific representation vectors as described in (6):

$$\mathbf{Z}_i^h = \mathbb{E}_{p_i^h}[\mathbf{v}^h(t)] = (\mathbf{W}_V^h)^\top \mathbf{B}^\top \int p_i^h(t) \boldsymbol{\psi}(t) dt. \quad (11)$$

Finally, we obtain the LTM representation \mathbf{Z}_{LTM} by concatenating the context heads and projecting them as in (3).

3.2. Continuous-Time Memory Consolidation

As we continue processing chunks, the LTM is progressively updated, as illustrated in Fig. 2. This is done by first sampling T locations within the interval $[0, 1]$ and then evaluating the continuous signal $\mathbf{x}(t)$, or current LTM, at these sampled points. The sampling can be performed either uniformly in $[0, 1]$ or based on prior attendance, as we will explain in detail in §3.3. The next step involves concatenating this LTM context with the new one coming from the current chunk. To do this, we first ‘‘contract’’ the LTM as

$$\mathbf{x}'(t) = \mathbf{x}(t/\tau) = \mathbf{B}^\top \boldsymbol{\psi}(t/\tau). \quad (12)$$

We then compute $\mathbf{x}(t)$ at the T locations $0 \leq t_1, \dots, t_T \leq \tau$:

$$\mathbf{x}_i = \mathbf{B}^\top \boldsymbol{\psi}(t_i/\tau) \quad \forall i \in [T], \quad (13)$$

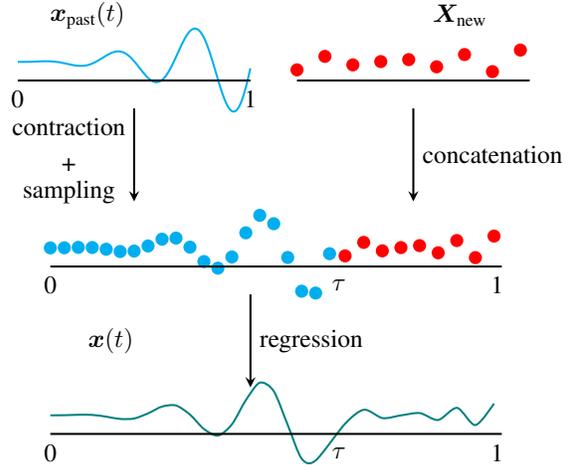


Figure 2. Proposed Memory Consolidation Mechanism.

and build the matrix $\mathbf{X}_{\text{past}} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]^\top \in \mathbb{R}^{T \times e}$. Following this, we concatenate the current step context $\mathbf{X}_{\text{new}} \in \mathbb{R}^{M \times d}$ with the previous context $\mathbf{X}_{\text{past}} \in \mathbb{R}^{T \times e}$, resulting in the combined sequence:

$$\mathbf{X} = [\mathbf{X}_{\text{past}}, \mathbf{X}_{\text{new}}]^\top \in \mathbb{R}^{(T+M) \times e}. \quad (14)$$

With the new and previous chunk contexts, we compute \mathbf{B} , as described in Eq. 5, where the continuous signal is approximated using a linear combination of rectangular functions. To achieve this, we first adjust the contribution of the previous context using the factor τ . Specifically, we associate \mathbf{X}_{past} with positions in the interval $[0, \tau]$ and \mathbf{X}_{new} with positions in $(\tau, 1]$. This process yields a matrix $\mathbf{G} \in \mathbb{R}^{(M+L) \times N}$. During each step, the previous context is contracted by a factor of τ , which regulates the extent of long-term memory used for attention and induces a gradual ‘‘forgetting’’ process. After the computation of \mathbf{B} , continuous attention is performed as described in §3.1.

3.3. Sticky Memories

When extending the LTM, the signal is evaluated at T locations within $[0, 1]$, as described in §3.2. While these locations can be uniformly distributed, allocating more ‘‘memory space’’ to regions of higher relevance ensures that critical information is prioritized in line with continuous resource allocation conceptualizations of human memory (Ma et al., 2014). This selective allocation compresses the signal into an efficient representation whereby essential video details are retained, and less significant ones are compressed or discarded. The recurrent and adaptive nature of this process is reminiscent of memory consolidation and reconsolidation brain mechanisms for relevance-based long-term memory transformation (Hardt et al., 2010; Dudai et al., 2015).

To address this, we propose selecting the T locations based

on the relevance of the signal in each region as done by Martins et al. (2022b). This process begins by constructing a histogram of the previous attention across intervals. Specifically, the signal is divided into D linearly spaced bins, $\{d_1, \dots, d_D\}$. The probability assigned to each bin, $p(d_j)$ for $j \in [D]$, is computed by integrating Eq. 10 over each bin interval using the trapezoidal rule:

$$p(d_j) \propto \sum_{h=1}^H \sum_{i=1}^R \int_{d_j} p_i^h(t). \quad (15)$$

Finally, T locations are sampled based on the resulting distribution, followed by the LTM update of §3.2. In our model, this sampling process is analogous to the phenomenon of non-local discontinuous “replay” in the brain whereby past experiences are reactivated for the purposes of consolidation (Carr et al., 2011; McNamee, 2024).

3.4. Model Architecture

We define the output context for the cross-attention layers of the video Q-former as a weighted sum of two components: the “local” vanilla video Q-former output context $\mathbf{Z}_{\text{STM}} \in \mathbb{R}^{R \times d}$, which corresponds to the already present attention over the current chunk, and the “global” LTM context \mathbf{Z}_{LTM} , which takes into account information from previous chunks as described in §3.2. The overall context is computed as:

$$\mathbf{Z} = \alpha \mathbf{Z}_{\text{STM}} + (1 - \alpha) \mathbf{Z}_{\text{LTM}}, \quad (16)$$

where α is a weighting factor that balances the contribution of short-term and long-term memories. The video Q-formers considered in this work were trained with the R tokens (or a subset) concatenated with the prompt tokens. In our approach, where the video context is divided into frame chunks processed by the long-term mechanism, it is necessary to aggregate information from all the C chunks, in $\mathbb{R}^{C \times R}$, into a fixed set of R tokens. To achieve this, we run the modified Video-LLaMA (Zhang et al., 2023b) and VideoChat2 (Li et al., 2024) models for each chunk and compute a running average of the embeddings as we process each chunk. The current video token embedding $\mathbf{E}_c \in \mathbb{R}^{R \times d}$, defined as $\mathbf{E}_c = \mathbf{W}_{\text{proj}} \mathbf{Z}_c$, is updated incrementally, enabling the model to handle arbitrarily long contexts without storing all embeddings of chunks in memory. The updated embedding is calculated as:

$$\bar{\mathbf{E}}_c = \frac{C-1}{C} \bar{\mathbf{E}}_{c-1} + \frac{1}{C} \mathbf{E}_c, \quad (17)$$

Upon reaching the final chunk, the video token embeddings \mathbf{E} are fed to the LLM that generates an answer. The full architecture diagram is shown in Fig. 1.

4. Experiments

In this section, we evaluate our proposed method on video question answering tasks, including multiple choice ques-

tion answering (§4.2) and open-ended generation (§4.3).

4.1. Implementation Details

The Video-LLaMA (Zhang et al., 2023b) architecture that we adapt here with our LTM module was initially designed for short videos and employs a dual Q-Former architecture (Li et al., 2023a)—one for spatial and another for temporal feature extraction. We use the Video-LLaMA2-7B finetuned model,³ leveraging EVA-CLIP’s ViT-G/14 (Fang et al., 2022) as visual encoder and Vicuna 7B (Chiang et al., 2023) as our LLM.

We also adapt VideoChat2 (Li et al., 2024), a stronger short-video model equipped with a single video Q-Former and trained on extended instruction data. We use UMT-L (Li et al., 2023c), which captures both spatial and temporal dependencies but requires more memory. Thus, we use chunks with fewer frames. Finally, we follow Jiang et al. (2023) and use Mistral-7B (Jiang et al., 2023). Additional implementation details can be found in App. A.

In all our experiments, we approximate the integrals of Eqs. 10 and 11 using the trapezoidal rule with 1000 sampling points. For ∞ -VIDEO with Video-LLaMA, we use 8 chunks of 256 frames with 1024 basis functions, except in the case of NeXT-QA (Xiao et al., 2021), where the total number of available frames is reduced. For ∞ -VIDEO with VideoChat2, we use 8 chunks of 16 frames with $N = 256$ basis functions. We experiment with several values of α in Eq. 16. Full hyperparameters can be seen in App. A.2.

4.2. Multiple-Choice Question Answering

4.2.1. COMPARISON WITH OTHER TRAINING-FREE METHODS

We consider systems of three kinds: (1) training-free approaches leveraging a GPT-4 backbone (Zhang et al., 2023a; Wang et al., 2024a;c); (2) models most similar to ours, which share the same underlying architecture, such as Video-LLaMA (Zhang et al., 2023b) and its training-free variants like MovieChat (Song et al., 2023) and MovieChat+ (Song et al., 2024); and (3) VideoChat2 (Li et al., 2024). For (2) and (3), we test ∞ -VIDEO variants without LTM (corresponding to $\alpha = 1.0$) and those using uniform sampling and sticky memories with $\alpha = 0.9$.

NeXT-QA. We evaluate our models in NeXT-QA (Xiao et al., 2021), a dataset with questions and 5 multiple choice options about **short** videos with average duration of 44 seconds. As ∞ -Video LLaMA can process a higher number of frames, we anticipate that using all the video information may lead to improvement over sub-sampling. For ∞ -Video

³<https://huggingface.co/DAMO-NLP-SG/Video-LLaMA-2-7B-Finetuned>

Table 1. Evaluation on Multiple Choice Datasets. Evaluation accuracies on NeXT-QA (NeXT) (Xiao et al., 2021) and Egoschema subset (Ego) (Mangalam et al., 2023). * denotes results obtained by running the models on both datasets. † indicates models run on Egoschema but not on NeXT-QA. ◊ highlights models trained on NeXT-QA. The remaining results are from Song et al. (2024) or Wang et al. (2024c). We bold the best-performing models for each category.

Method	LLM	#Frames	NeXT	Ego
<i>Based on Proprietary LLMs</i>				
LLoVi (Zhang et al., 2023a)	GPT-4	-	67.7	61.2
VideoAgent (Wang et al., 2024a)	GPT-4	-	71.3	60.2
VideoTree (Wang et al., 2024c)	GPT-4	-	75.6	66.2
<i>Video LLaMA-Based Models</i>				
Video LLaMA* (Zhang et al., 2023b)	Vicuna-7B	32	30.7	20.2
MovieChat† (Song et al., 2023)	Vicuna-7B	2048	34.4	41.6
MovieChat+† (Song et al., 2024)	Vicuna-7B	2048	35.2	37.4
∞ -Video LLaMA (No LTM)*	Vicuna-7B	all/2048	37.6	40.8
∞ -Video LLaMA (Uniform)*	Vicuna-7B	all/2048	37.5	42.6
∞ -Video LLaMA (Sticky)*	Vicuna-7B	all/2048	41.1	46.8
<i>VideoChat2-Based Models</i>				
VideoChat2*◊ (Li et al., 2024)	Mistral-7B	16	78.7	64.2
∞ -VideoChat2 (No LTM)*◊	Mistral-7B	128	78.1	64.6
∞ -VideoChat2 (Uniform)*◊	Mistral-7B	128	78.1	64.4
∞ -VideoChat2 (Sticky)*◊	Mistral-7B	128	78.1	64.8

LLaMA, we use all available frames in the videos, which we then split into chunks of up to 256 frames. In contrast, baseline models like Video-LLaMA (Zhang et al., 2023b) are limited to processing only 32 frames, while VideoChat2 achieves optimal performance for 16 frames, as shown in Li et al. (2024).

As shown in Tab. 1, ∞ -Video LLaMA with sticky memories outperforms other approaches. This includes MovieChat+, which employs a heuristic to merge adjacent frames, fed to the video’s Q-former cross-attention with question knowledge. In contrast, our model encounters the question only within the LLM prompt. Surprisingly, our method using uniform sampling of the continuous signal performs slightly worse than ∞ -Video LLaMA without the LTM ($\alpha = 1$). For the VideoChat2 variants, we observe no significant performance increase. We attribute this to the in-domain nature of the evaluation, where the original model is already highly optimized for the given tasks.

Egoschema. We evaluate the training-free question-answering performance of our models on EgoSchema (Mangalam et al., 2023), a medium-length benchmark for ego-centric planning designed to test long-context video understanding with 3-minute videos.

Tab. 1 presents the results of our methods compared to strong parameter-free baselines. ∞ -Video LLaMA, equipped with the LTM module ($\alpha = 0.9$) and using sticky memories, significantly outperforms both the uniform LTM variant, also for $\alpha = 0.9$ and the model without LTM ($\alpha = 1$), achieving a notable accuracy improvement of +6 points over the latter. A similar trend is observed with ∞ -

Table 2. VideoMME results. Baseline results are taken from (Fu et al., 2024). We bold the best performing models.

Method	LLM	#Frames	Medium	Long	Avg
ST-LLM	Vicuna-7B	64	36.8	31.1	37.9
Video-LLaVA	Vicuna-7B	8	38.0	36.2	39.9
ShareGPT4Video 8B	-	16	36.3	35.0	39.9
Chat-UniVi-v1.5	Vicuna-7B	64	40.3	35.8	40.6
Qwen-VL-Chat	Qwen-7B	4	38.7	37.8	41.1
VideoChat2	Mistral-7B	32	37.9	38.0	42.1
∞ -VideoChat2 (no LTM)	Mistral-7B	128	39.6	38.8	42.3
∞ -VideoChat2 (uniform)	Mistral-7B	128	40.0	38.8	42.4
∞ -VideoChat2 (sticky)	Mistral-7B	128	40.2	38.9	42.4

VideoChat2, though the improvements are less pronounced. We attribute this to the fact that VideoChat2 is inherently a stronger model, having been trained on a larger and more recent datasets, offering less room for improvement compared to the Video-LLaMA ∞ -variants. However, gains are observed for both uniform sampling and sticky memories, with the latter showing superior performance. Moreover, despite having significantly fewer parameters, ∞ -VideoChat2 demonstrates competitive results against proprietary LLMs based on ChatGPT-4.

4.2.2. EVALUATION ON VERY LONG VIDEOS

To emphasize the effectiveness of our approach on extended video content, we also present results on Video-MME (Fu et al., 2024), which features a diverse collection of lengthy videos, ranging up to 1 hour in duration. We compare our method against baseline models of similar size such as ST-LLM (Liu et al., 2024b), Video-LLaVA (Liu et al., 2023), ShareGPT4Video 8B (Chen et al., 2024a), Chat-UniVi-v1.5 (Jin et al., 2023) and Qwen-VL-Chat (Bai et al., 2023) as well as our ∞ -VIDEO variants with $\alpha = 1$ and the base architecture VideoChat2 (Li et al., 2024).

In Tab. 2, we show the results for VideoMME. Although Video-LLaMA performs well on earlier datasets, its description-focused training dataset makes the performance as good as random guessing. Including it in the evaluation would detract from more relevant models. However, for the VideoChat2 category, sticky memories outperform others, followed by uniform sampling.

4.3. Long-Term Open-Ended Question Answering

We now investigate the performance of our models on open-ended question answering using the MovieChat-1K dataset (Song et al., 2023), a benchmark comprising long videos with an average duration of around 8 minutes. We compare our models with other baselines based on Vicuna-7B (Chiang et al., 2023), including VideoChat (Li et al., 2023b), Video-ChatGPT (Maaz et al., 2024), MovieChat (Song et al., 2023; 2024), MovieChat+ (Song et al., 2024). Furthermore,

Table 3. **MovieChat Results.** **Score** measures the overall answer score, **CI** stands for correctness of information, **DO** stands for detail orientation, and **CU** stands for contextual understanding. We bold the best results and underline the best within a category. We omit the temporal and consistency metrics due to the absence of the subset specific to these metrics. Except for Moviechat+, results were taken from (Song et al., 2023).

Method	LLM	Number of Frames	Accuracy	Score	CI	DO	CU
Video Chat (Li et al., 2023b)	Vicuna-7B	32	61.0	3.34	3.26	3.20	3.38
Video-ChatGPT (Maaz et al., 2024)	Vicuna-7B	100	44.2	2.71	2.48	2.78	3.03
<i>Video LLaMA-Based Models</i>							
Video LLaMA (Zhang et al., 2023b)	Vicuna-7B	32	51.4	3.10	3.30	2.53	3.28
MovieChat (Song et al., 2023)	Vicuna-7B	2048	67.8	3.81	3.32	3.28	3.44
MovieChat+ (Song et al., 2024)	Vicuna-7B	2048	66.4	3.67	3.70	3.30	3.62
∞ -Video LLaMA (no LTM)	Vicuna-7B	2048	68.0	3.76	3.72	3.33	3.71
∞ -Video LLaMA (uniform)	Vicuna-7B	2048	66.5	3.69	3.60	3.31	3.58
∞ -Video LLaMA (sticky)	Vicuna-7B	2048	72.2	3.88	3.89	3.47	3.79
∞ -Video LLaMA (no STM uniform)	Vicuna-7B	2048	62.4	3.75	3.36	3.38	3.52
∞ -Video LLaMA (no STM sticky)	Vicuna-7B	2048	59.2	3.68	3.30	3.30	3.44
<i>VideoChat2-Based Models</i>							
VideoChat2	Mistral-7B	16	62.2	3.72	3.46	3.60	3.69
∞ -VideoChat2 (no LTM)	Mistral-7B	128	63.9	3.74	3.54	3.60	3.73
∞ -VideoChat2 (uniform)	Mistral-7B	128	64.1	3.73	3.54	3.60	3.75
∞ -VideoChat2 (sticky)	Mistral-7B	128	63.9	3.74	3.55	3.63	3.74
∞ -VideoChat2 (no STM uniform)	Mistral-7B	128	65.7	3.78	3.65	3.60	3.84
∞ -VideoChat2 (no STM sticky)	Mistral-7B	128	66.5	3.85	3.71	3.68	3.96

we evaluate our ∞ -VIDEO variants with $\alpha = 0.9$ for both uniform sampling and sticky memories, as well as with $\alpha = 0$ (i.e., without STM). Following the standard evaluation method for open-ended questions, we prompt GPT-3.5 (OpenAI et al., 2024) for a yes/no answer prediction, a confidence score (0 to 5), and other qualitative metrics. The prompts are shown in App. A.3.

As shown in Tab. 3, our ∞ -VIDEO LLaMA with sticky memories outperforms all models in its category, as well as VideoChat and Video-ChatGPT, across all metrics. It also surpasses MovieChat, which is designed for training-free long-context video understanding. In contrast, ∞ -Video LLaMA with uniform sampling performs worse than both sticky memories and the model without LTM. Additionally, using only the LTM does not improve performance, highlighting that a weighted combination of STM and LTM yields the best results for the Video-LLaMA category. The same does not apply to the VideoChat2 category, where replacing the STM with the LTM yields top results. For α values other than 0, the performance of ∞ -VIDEO with VideoChat2 remains nearly unchanged. Surprisingly, VideoChat2 underperforms compared to the Video LLaMA category despite being trained on a larger dataset. We hypothesize this is due to differences in the training datasets: Video LLaMA was trained on video descriptions with multiple sentences, encouraging more context, which increases the probability of correct predictions, while VideoChat2 was fine-tuned on concise datasets, favouring brief, open-ended predictions. Ablation studies on ∞ -VIDEO with Video LLaMA are presented in App. B.1.

4.4. Qualitative Analysis

In Fig. 3, we show the continuous attention density map over the LTM in the final layer of the video Q-Former for the last chunk of ∞ -VIDEO LLaMA. The example uses the *Interstellar* trailer, divided into 8 chunks of 256 frames each, with $\tau = 0.75$, $N = 1024$, and $\alpha = 0.9$. The bottom heatmap reveals that the continuous attention favours frames after $t = \tau$, where the sticky LTM exhibits peaks before this point. In contrast, the uniform LTM shows vanishing density as t decreases, likely due to context contraction across chunks, which might explain the superior results for sticky memories.

In Fig. 4, we present the attention density as a function of the number of frames, with $\alpha = 0.9$, $N = 256$, and $\tau = 0.5$, for 3 chunks of 256 frames each, spanning 3 contraction steps. We also display 6 representative frames from high-density regions by identifying the top 26 frames within each interval and selecting the 6 non-redundant frames. The selected frames appear to align with visually striking or narratively significant scenes within the trailer, which suggests that ∞ -VIDEO effectively might capture key moments in the video and discard irrelevant parts. For example, in the final interval, the region with the lowest attention density corresponds to the credits section of the video. We show a similar figure but for the uniform sampling in App. B.2.

5. Related Work

There are several recent advances in long-context video understanding (Li et al., 2023d; Liu et al., 2024a; Balazevic et al., 2024; Wang et al., 2024b; Shu et al., 2024; Ye

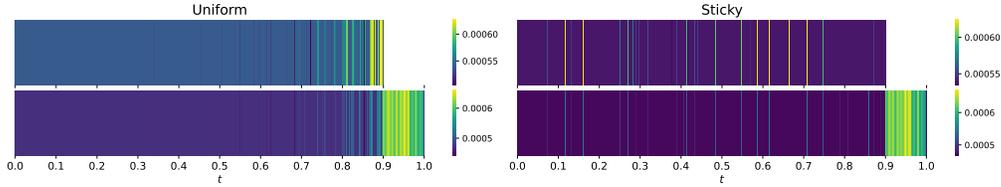


Figure 3. **(Top)** LTM attention density on the $[0, \tau]$ interval for the *Interstellar* trailer, using sticky memories in the final chunk of the ∞ -Video LLaMA video Q-former’s last layer. **(Bottom)** The same attention density map, extended over the full t interval.

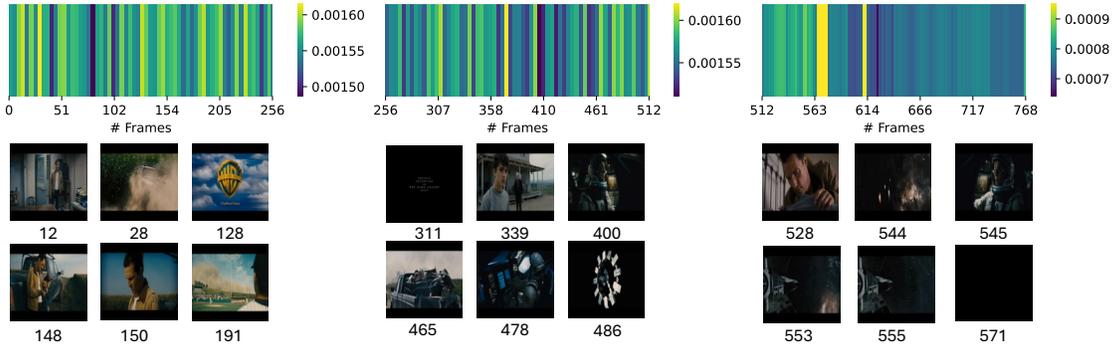


Figure 4. Highest continuous attention density frames selected using sticky memories in the *Interstellar* trailer for ∞ -Video LLaMA across 3 chunks. **(Left)** Interval: $[0, \tau^2]$. **(Middle)** Interval: $(\tau^2, \tau]$. **(Right)** Interval: $(\tau, 1]$.

et al., 2024; Chen et al., 2024b), but few adapting transformers to leverage temporal information in a training-free setting. Closest to our work is Song et al. (2023; 2024), which extends Video LLaMA with memory consolidation by using a heuristic to merge similar frames, which alters frame embedding-level information. In contrast, our method retains embedding integrity, equipping the video Q-former’s cross-attention with an LTM to efficiently process an arbitrary number of frames in one single pass over the video.

Moreover, works such as Zhang et al. (2024), Shu et al. (2024), and Chen et al. (2024b) address the challenge of long video contexts. However, these approaches involve fine-tuning or training models from scratch, which can be computationally expensive and time-intensive. Our approach, in contrast, enables the seamless adaptation of short video models to arbitrary long contexts without the need for sparse subsampling or discarding important information.

Our approach builds on continuous attention mechanisms, originally proposed by Martins et al. (2020), and later applied to image, speech, and natural language processing tasks (Farinhas et al., 2021; Martins et al., 2022a;b). We extend these ideas to video data by replacing the continuous softmax with the Gibbs PDF, which better replicates the discrete softmax used in the video Q-Former attention.

6. Conclusions

We introduced a lightweight extension to short-video vision LLMs, enabling arbitrary-long video understanding by augmenting the video Q-former’s cross-attention mechanism with a long-term memory module that consolidates global information dynamically. Our approach adapts continuous attention to perform visual memory consolidation, allocating higher granularity to the most relevant frames. This ensures an efficient and focused representation of critical moments while maintaining scalability. Additionally, our method enables the sequential processing of videos with a single pass. Despite being training-free, our approach also paves the way for scalable long-context video understanding with transformers as spatio-temporal feature extractors.

Our work takes inspiration from cognitive and mechanistic theories of memory (re)consolidation in brains (Hardt et al., 2010; Preston & Eichenbaum, 2013; Ma et al., 2014), and a deeper integration with such theories may be pursued. For example, memory reactivation or “replay” is thought to be a key component of systems consolidation in “offline” states such as sleep. Our model could be extended to incorporate such with further training and schema-driven fine-tuning for the purposes of continual learning (Cai et al., 2024). Furthermore, as a neural architecture, our model goes beyond current brain models of episodic memory processing, which focus on discrete low-dimensional sequences of static images and relatively simple functionalities such as mem-

ory recall and event segmentation (Franklin et al., 2020; Chandra et al., 2025). Given our model’s integration of rich and streaming visual input with flexible and sophisticated querying via text prompting, interpretability analyses of our architecture may provide insights regarding how episodic memory may be interrogated for complex inferences in the human brain (Tulving, 2002; Radvansky & Zacks, 2014).

Impact Statement

We discuss the broader implications of our work, including ethical considerations and potential societal consequences. Our framework extends the capabilities of existing short-context multimodal language models, enabling them to process unbounded video contexts without requiring retraining. This is particularly relevant given concerns about the energy consumption of training large models (Strubell et al., 2019). However, we must also acknowledge the societal risks associated with video models, especially their potential use in privacy-violating surveillance. As our approach enables scaling to longer videos, there is a concern that it could be applied in undesirable domains. While current state-of-the-art models are often trained on datasets with documented biases, we intentionally focus on applications using standard benchmark applications, aiming to distance ourselves from harmful uses.

Acknowledgments

We would like to thank Marcos Treviso, Giuseppe Attanasio, Sweta Agrawal, Chryssa Zerva and the SARDINE lab team for helpful discussions. This work was supported by EU’s Horizon Europe Research and Innovation Actions (UTTER, contract 101070631), by the project DECOLLAGE (ERC-2022-CoG 101088763), by the Portuguese Recovery and Resilience Plan through project C645008882-00000055 (Center for Responsible AI), and by FCT/MECI through national funds and when applicable co-funded EU funds under UID/50008: Instituto de Telecomunicações.

References

Atreja, S., Ashkinaze, J., Li, L., Mendelsohn, J., and Hemphill, L. Prompt design matters for computational social science tasks but in unpredictable ways, 2024. URL <https://arxiv.org/abs/2406.11980>.

Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. In *Proc. of International Conference on Learning Representations*, 2015.

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 reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023.

- Balazevic, I., Shi, Y., Papalampidi, P., Chaabouni, R., Koppala, S., and Henaff, O. J. Memory consolidation enables long-context video understanding. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 2527–2542. PMLR, 21–27 Jul 2024.
- Brady, T. F., Konkle, T., Alvarez, G. A., and Oliva, A. Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, 105(38):14325–14329, 2008. doi: 10.1073/pnas.0803390105.
- Brown, P. J. and Zidek, J. V. Adaptive multivariate ridge regression. *The Annals of Statistics*, 1980.
- Cai, C., Wang, Z., Gao, J., Liu, W., Lu, Y., Zhang, R., and Yap, K.-H. Empowering large language model for continual video question answering with collaborative prompting, 2024. arXiv:2410.00771v2.
- Carr, M. F., Jadhav, S. P., and Frank, L. M. Hippocampal replay in the awake state: A potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14: 147–153, 2011.
- Chandra, S., Sharma, S., Chaudhuri, R., et al. Episodic and associative memory from spatial scaffolds in the hippocampus. *Nature*, XX(XX):XX–XX, 2025. doi: 10.1038/s41586-024-08392-y.
- Chen, L., Wei, X., Li, J., Dong, X., Zhang, P., Zang, Y., Chen, Z., Duan, H., Lin, B., Tang, Z., et al. Sharegpt4video: Improving video understanding and generation with better captions. *arXiv preprint arXiv:2406.04325*, 2024a.
- 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, 2024b.
- Cheng, Z., Leng, S., Zhang, H., Xin, Y., Li, X., Chen, G., Zhu, Y., Zhang, W., Luo, Z., Zhao, D., and Bing, L. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024. URL <https://arxiv.org/abs/2406.07476>.
- Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J. E., and et al. Vicuna: An open-source chatbot impressing gpt-4

- with 90% chatgpt quality. <https://vicuna.lmsys.org>, 2023.
- Cowan, E. T., Schapiro, A. C., Dunsmoor, J. E., and Murty, V. P. Memory consolidation as an adaptive process. *Psychonomic Bulletin & Review*, 28:1796–1810, 2021. doi: 10.3758/s13423-021-01978-x. URL <https://doi.org/10.3758/s13423-021-01978-x>.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- Dudai, Y., Karni, A., and Born, J. The consolidation and transformation of memory. *Neuron*, 88:20–32, 2015.
- Fang, Y., Wang, W., Xie, B., Sun, Q., Wu, L., Wang, X., Huang, T., Wang, X., and Cao, Y. Eva: Exploring the limits of masked visual representation learning at scale. 2022.
- Farinhas, A., Martins, A. F. T., and Aguiar, P. M. Q. Multi-modal continuous visual attention mechanisms. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pp. 1047–1056, October 2021.
- Frankland, P. W. and Bontempi, B. The organization of recent and remote memories. *Nature Reviews Neuroscience*, 6(2):119–130, 2005.
- Franklin, N. T., Norman, K. A., Ranganath, C., Zacks, J. M., and Gershman, S. J. Structured event memory: A neuro-symbolic model of event cognition. *Psychological Review*, 127(3):327–361, 2020. doi: 10.1037/rev0000177.
- Fu, C., Dai, Y., Luo, Y., Li, L., Ren, S., Zhang, R., Wang, Z., Zhou, C., Shen, Y., Zhang, M., Chen, P., Li, Y., Lin, S., Zhao, S., Li, K., Xu, T., Zheng, X., Chen, E., Ji, R., and Sun, X. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis, 2024.
- Hardt, O., Einarsson, E. O., and Nader, K. A bridge over troubled water: reconsolidation as a link between cognitive and neuroscientific memory research traditions. *Annual Review of Psychology*, 61:141–167, 2010. doi: 10.1146/annurev.psych.093008.100455.
- hwchase17. Langchain, 2023. URL <https://github.com/hwchase17/langchain>. Accessed: 2023-12-20.
- Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., de las Casas, D., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., Lavaud, L. R., Lachaux, M.-A., Stock, P., Scao, T. L., Lavril, T., Wang, T., Lacroix, T., and Sayed, W. E. Mistral 7b, 2023.
- Jin, P., Takanobu, R., Zhang, C., Cao, X., and Yuan, L. Chat-univi: Unified visual representation empowers large language models with image and video understanding. *arXiv preprint arXiv:2311.08046*, 2023.
- 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 *Proceedings of the 40th International Conference on Machine Learning*, 2023a.
- Li, K., He, Y., Wang, Y., Li, Y., Wang, W., Luo, P., Wang, Y., Wang, L., and Qiao, Y. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023b.
- Li, K., Wang, Y., Li, Y., Wang, Y., He, Y., Wang, L., and Qiao, Y. Unmasked teacher: Towards training-efficient video foundation models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 19891–19903, 2023c.
- Li, K., Wang, Y., He, Y., Li, Y., Wang, Y., Liu, Y., Wang, Z., Xu, J., Chen, G., Luo, P., Wang, L., and Qiao, Y. Mvbench: A comprehensive multi-modal video understanding benchmark, 2024. URL <https://arxiv.org/abs/2311.17005>.
- Li, Y., Wang, C., and Jia, J. Llama-vid: An image is worth 2 tokens in large language models, 2023d.
- Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning, 2023.
- 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, 2024a.
- Liu, R., Li, C., Tang, H., Ge, Y., Shan, Y., and Li, G. St-llm: Large language models are effective temporal learners, 2024b.
- 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. Nvila: Efficient frontier visual language models, 2024c.
- Luo, R., Zhao, Z., Yang, M., Dong, J., Qiu, M., Lu, P., Wang, T., and Wei, Z. Valley: Video assistant with large language model enhanced ability, 2023.
- Ma, W., Husain, M., and Bays, P. Changing concepts of working memory. *Nature Neuroscience*, 17:347–356,

2014. doi: 10.1038/nn.3655. URL <https://doi.org/10.1038/nn.3655>.
- Maaz, M., Rasheed, H., Khan, S., and Khan, F. S. Videochatpt: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*, 2024.
- Mangalam, K., Akshulakov, R., and Malik, J. Egoschema: A diagnostic benchmark for very long-form video language understanding. *arXiv preprint arXiv:2308.09126*, 2023.
- Martins, A., Farinhas, A., Treviso, M., Niculae, V., Aguiar, P., and Figueiredo, M. Sparse and continuous attention mechanisms. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 20989–21001. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/f0b76267fbel2b936bd65e203dc675c1-Paper.pdf.
- Martins, A. F. T., Treviso, M., Farinhas, A., Aguiar, P. M. Q., Figueiredo, M. A. T., Blondel, M., and Niculae, V. Sparse continuous distributions and fenchel-young losses. *Journal of Machine Learning Research*, 23(257): 1–74, 2022a. URL <http://jmlr.org/papers/v23/21-0879.html>.
- Martins, P. H., Marinho, Z., and Martins, A. F. ∞ -former: Infinite memory transformer. In *Proc. ACL*, 2022b.
- McGaugh, J. Making lasting memories: Remembering the significant. *Proceedings of the National Academy of Sciences of the United States of America*, 110:10402–10407, 2013. doi: 10.1073/pnas.1301209110. URL <https://doi.org/10.1073/pnas.1301209110>.
- McNamee, D. C. The generative neural microdynamics of cognitive processing. *Current Opinion in Neurobiology*, 85:102855, 2024. doi: 10.1016/j.conb.2024.102855. Epub 2024 Feb 29.
- OpenAI, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., Avila, R., Babuschkin, I., Balaji, S., Balcom, V., Baltescu, P., Bao, H., Bavarian, M., Belgum, J., Bello, I., Berdine, J., Bernadett-Shapiro, G., Berner, C., Bogdonoff, L., Boiko, O., Boyd, M., Brakman, A.-L., Brockman, G., Brooks, T., Brundage, M., Button, K., Cai, T., Campbell, R., Cann, A., Carey, B., Carlson, C., Carmichael, R., Chan, B., Chang, C., Chantzis, F., Chen, D., Chen, S., Chen, R., Chen, J., Chen, M., Chess, B., Cho, C., Chu, C., Chung, H. W., Cummings, D., Currier, J., Dai, Y., Decareaux, C., Degry, T., Deutsch, N., Deville, D., Dhar, A., Dohan, D., Dowling, S., Dunning, S., Ecoffet, A., Eleti, A., Eloundou, T., Farhi, D., Fedus, L., Felix, N., Fishman, S. P., Forte, J., Fulford, I., Gao, L., Georges, E., Gibson, C., Goel, V., Gogineni, T., Goh, G., Gontijo-Lopes, R., Gordon, J., Grafstein, M., Gray, S., Greene, R., Gross, J., Gu, S. S., Guo, Y., Hallacy, C., Han, J., Harris, J., He, Y., Heaton, M., Heidecke, J., Hesse, C., Hickey, A., Hickey, W., Hoeschele, P., Houghton, B., Hsu, K., Hu, S., Hu, X., Huizinga, J., Jain, S., Jain, S., Jang, J., Jiang, A., Jiang, R., Jin, H., Jin, D., Jomoto, S., Jonn, B., Jun, H., Kaftan, T., Łukasz Kaiser, Kamali, A., Kanitscheider, I., Keskar, N. S., Khan, T., Kilpatrick, L., Kim, J. W., Kim, C., Kim, Y., Kirchner, J. H., Kiros, J., Knight, M., Kokotajlo, D., Łukasz Kondraciuk, Kondrich, A., Konstantinidis, A., Kosic, K., Krueger, G., Kuo, V., Lampe, M., Lan, I., Lee, T., Leike, J., Leung, J., Levy, D., Li, C. M., Lim, R., Lin, M., Lin, S., Litwin, M., Lopez, T., Lowe, R., Lue, P., Makanju, A., Malfacini, K., Manning, S., Markov, T., Markovski, Y., Martin, B., Mayer, K., Mayne, A., McGrew, B., McKinney, S. M., McLeavey, C., McMillan, P., McNeil, J., Medina, D., Mehta, A., Menick, J., Metz, L., Mishchenko, A., Mishkin, P., Monaco, V., Morikawa, E., Mossing, D., Mu, T., Murati, M., Murk, O., Mély, D., Nair, A., Nakano, R., Nayak, R., Neelakantan, A., Ngo, R., Noh, H., Ouyang, L., O’Keefe, C., Pachocki, J., Paino, A., Palermo, J., Pantuliano, A., Parascandolo, G., Parish, J., Parparita, E., Passos, A., Pavlov, M., Peng, A., Perelman, A., de Avila Belbute Peres, F., Petrov, M., de Oliveira Pinto, H. P., Michael, Pokorny, Pokrass, M., Pong, V. H., Powell, T., Power, A., Power, B., Proehl, E., Puri, R., Radford, A., Rae, J., Ramesh, A., Raymond, C., Real, F., Rimbach, K., Ross, C., Rotsted, B., Roussez, H., Ryder, N., Saltarelli, M., Sanders, T., Santurkar, S., Sastry, G., Schmidt, H., Schnurr, D., Schulman, J., Sel-sam, D., Sheppard, K., Sherbakov, T., Shieh, J., Shoker, S., Shyam, P., Sidor, S., Sigler, E., Simens, M., Sitkin, J., Slama, K., Sohl, I., Sokolowsky, B., Song, Y., Staudacher, N., Such, F. P., Summers, N., Sutskever, I., Tang, J., Tezak, N., Thompson, M. B., Tillet, P., Tootoonchian, A., Tseng, E., Tuggle, P., Turley, N., Tworek, J., Uribe, J. F. C., Vallone, A., Vijayvergiya, A., Voss, C., Wainwright, C., Wang, J. J., Wang, A., Wang, B., Ward, J., Wei, J., Weinmann, C., Welihinda, A., Welinder, P., Weng, J., Weng, L., Wiethoff, M., Willner, D., Winter, C., Wolrich, S., Wong, H., Workman, L., Wu, S., Wu, J., Wu, M., Xiao, K., Xu, T., Yoo, S., Yu, K., Yuan, Q., Zaremba, W., Zellers, R., Zhang, C., Zhang, M., Zhao, S., Zheng, T., Zhuang, J., Zhuk, W., and Zoph, B. Gpt-4 technical report, 2024.
- Preston, A. R. and Eichenbaum, H. The interplay of hippocampus and prefrontal cortex in memory-based decision making. *Current Biology*, 23(17):R764–R773, 2013.
- Radvansky, G. A. and Zacks, J. M. *Event Cognition*. Oxford

- University Press, New York, NY, 2014.
- Shu, Y., Liu, Z., Zhang, P., Qin, M., Zhou, J., Liang, Z., Huang, T., and Zhao, B. Video-xl: Extra-long vision language model for hour-scale video understanding, 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., Ye, T., Hwang, J.-N., Li, X., and Wang, G. Moviechat+: Question-aware sparse memory for long video question answering. *arXiv preprint arXiv:2404.17176*, 2024.
- Strubell, E., Ganesh, A., and McCallum, A. Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3645–3650, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1355. URL <https://aclanthology.org/P19-1355>.
- Tulving, E. Episodic memory: From mind to brain. *Annual Review of Psychology*, 53:1–25, 2002. doi: 10.1146/annurev.psych.53.100901.135114. URL <https://doi.org/10.1146/annurev.psych.53.100901.135114>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- Wang, X., Zhang, Y., Zohar, O., and Yeung-Levy, S. Videoagent: Long-form video understanding with large language model as agent. *European Conference on Computer Vision (ECCV)*, 2024a.
- Wang, Y., Xie, C., Liu, Y., and Zheng, Z. Videollamb: Long video understanding with recurrent memory bridges. *arxiv*, 2024b.
- Wang, Z., Yu, S., Stengel-Eskin, E., Yoon, J., Cheng, F., Bertasius, G., and Bansal, M. Videotree: Adaptive tree-based video representation for llm reasoning on long videos. *arXiv preprint arXiv:2405.19209*, 2024c.
- 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*, pp. 9777–9786, 2021.
- 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, 2024.
- Zhang, C., Lu, T., Islam, M. M., Wang, Z., Yu, S., Bansal, M., and Bertasius, G. A simple llm framework for long-range video question-answering, 2023a.
- Zhang, H., Li, X., and Bing, L. Videollama: An instruction-tuned audio-visual language model for video understanding. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023)*, 2023b.
- 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*, 2024. URL <https://arxiv.org/abs/2406.16852>.
- Zhu, D., Chen, J., Shen, X., Li, X., and Elhoseiny, M. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.

A. Implementation Details and Hyperparameters

A.1. Additional Implementation Details

Video LLaMA-Based Models. As discussed in §3, Video-LLaMA (Zhang et al., 2023b) serves as one of the adapted models in this work. It employs a cascade of two Q-formers—one for spatial feature extraction and the other for temporal feature extraction—with the latter enhanced by our LTM integration. The video Q-former and projection layer parameters are consistent with the Video-LLaMA-2-7B-Finetuned model (Zhang et al., 2023b), which was fine-tuned using instruction-tuning data from MiniGPT-4 (Zhu et al., 2023), LLaVA (Liu et al., 2023), and VideoChat (Li et al., 2023b). For visual feature extraction, we utilize the ViT-G/14 encoder from EVA-CLIP (Fang et al., 2022), while spatial dependency features rely on the Q-former from BLIP-2 (Li et al., 2023a). This lightweight ViT facilitates efficient processing of a large number of frames per chunk. The LLM used in this model is Vicuna-7B (Chiang et al., 2023).

Atreja et al. (2024) and our empirical validation have shown limitations in respecting the format of the multiple-choice answer asked by the prompt, influencing the accuracy of this method in multiple-choice datasets. To address this, and for these datasets we follow the approach of (Song et al., 2023; 2024) and provide our modified model, ∞ -Video LLaMA, exclusively with the questions in the prompt. Using LangChain (hwchase17, 2023), we calculate the similarity between ∞ -Video LLaMA’s open-ended responses and the given options, selecting the option that best aligns with the expected answer.

VideoChat2-Based Models. Building on Video-LLaMA, we also evaluated our methods using VideoChat2 (Li et al., 2024), a more advanced short-video model. Like Video-LLaMA, it incorporates a video Q-former, but it benefits from additional instruction-tuning data (Li et al., 2024). For visual encoding, we use UMT-L (Li et al., 2023c), which captures both spatial and temporal features specifically designed for video data. This model’s higher memory requirements necessitated the use of smaller frame chunks in our experiments. For the LLM component, we used the stage-3 Mistral-7B version of VideoChat2,⁴ which demonstrated the best performance in its original paper.

A.2. Hyperparameters

We report in Tab. 4 the hyperparameters used in our experiments. For ∞ -Video LLaMA and for NeXT-QA we use all the frames available with variable number of chunks of 256 frames.

Table 4. Hyperparameters used in our ∞ -variants for the different datasets.

Parameter	NeXT-QA		Egoschema		VideoMME	MovieChat	
	∞ -Video LLaMA	∞ -VideoChat2	∞ -Video LLaMA	∞ -VideoChat2	∞ -VideoChat2	∞ -Video LLaMA	∞ -VideoChat2
# chunks	-	8	8	8	8	8	8
# frames	all	16	256	16	16	256	256
N	256	256	1024	256	256	1024	256
τ	0.75	0.75	0.75	0.75	0.5	0.75	0.75

A.3. Evaluation

We further show in List. 1, 2, 3, 4 the prompts used for evaluation on open-ended question answering tasks for the Moviechat dataset.

Listing 1. ChatGPT-3.5 prompt for the overall accuracy and score metric.

```
"role": "system",
"content":
  "You are an intelligent chatbot designed for evaluating the correctness of generative outputs for question-answer pairs. "
  "Your task is to compare the predicted answer with the correct answer and determine if they match meaningfully. Here's how you can
  accomplish the task:"
  "-----"
  "##INSTRUCTIONS: "
  "- Focus on the meaningful match between the predicted answer and the correct answer.\n"
  "- Consider synonyms or paraphrases as valid matches.\n"
  "- Evaluate the correctness of the prediction compared to the answer."
},
{
```

⁴https://huggingface.co/OpenGVLab/VideoChat2_stage3_Mistral_7B

```
"role": "user",
"content":
  "Please evaluate the following video-based question-answer pair:\n\n"
  f"Question: {question}\n"
  f"Correct Answer: {answer}\n"
  f"Predicted Answer: {pred}\n\n"
  "Provide your evaluation only as a yes/no and score where the score is an integer value between 0 and 5, with 5 indicating the
  highest meaningful match. "
  "Please generate the response in the form of a Python dictionary string with keys 'pred' and 'score', where value of 'pred' is a
  string of 'yes' or 'no' and value of 'score' is in INTEGER, not STRING."
  "DO NOT PROVIDE ANY OTHER OUTPUT TEXT OR EXPLANATION. Only provide the Python dictionary string. "
  "For example, your response should look like this: {'pred': 'yes', 'score': 4.8}."
}
```

Listing 2. ChatGPT-3.5 prompt for the contextual understanding (CI) metric.

```
"role": "system",
"content":
  "You are an intelligent chatbot designed for evaluating the factual accuracy of generative outputs for video-based question-answer
  pairs. "
  "Your task is to compare the predicted answer with the correct answer and determine if they are factually consistent. Here's how
  you can accomplish the task:"
  "-----"
  "##INSTRUCTIONS: "
  "- Focus on the factual consistency between the predicted answer and the correct answer. The predicted answer should not contain
  any misinterpretations or misinformation.\n"
  "- The predicted answer must be factually accurate and align with the video content.\n"
  "- Consider synonyms or paraphrases as valid matches.\n"
  "- Evaluate the factual accuracy of the prediction compared to the answer."
},
{
"role": "user",
"content":
  "Please evaluate the following video-based question-answer pair:\n\n"
  f"Question: {question}\n"
  f"Correct Answer: {answer}\n"
  f"Predicted Answer: {pred}\n\n"
  "Provide your evaluation only as a factual accuracy score where the factual accuracy score is an integer value between 0 and 5,
  with 5 indicating the highest level of factual consistency. "
  "Please generate the response in the form of a Python dictionary string with keys 'score', where its value is the factual accuracy
  score in INTEGER, not STRING."
  "DO NOT PROVIDE ANY OTHER OUTPUT TEXT OR EXPLANATION. Only provide the Python dictionary string. "
  "For example, your response should look like this: {'score': 4.8}."
}
```

Listing 3. ChatGPT-3.5 prompt for the detailed orientation (DO) metric.

```
"role": "system",
"content":
  "You are an intelligent chatbot designed for evaluating the detail orientation of generative outputs for video-based question-
  answer pairs. "
  "Your task is to compare the predicted answer with the correct answer and determine its level of detail, considering both
  completeness and specificity. Here's how you can accomplish the task:"
  "-----"
  "##INSTRUCTIONS: "
  "- Check if the predicted answer covers all major points from the video. The response should not leave out any key aspects.\n"
  "- Evaluate whether the predicted answer includes specific details rather than just generic points. It should provide comprehensive
  information that is tied to specific elements of the video.\n"
  "- Consider synonyms or paraphrases as valid matches.\n"
  "- Provide a single evaluation score that reflects the level of detail orientation of the prediction, considering both completeness
  and specificity."
},
{
"role": "user",
"content":
  "Please evaluate the following video-based question-answer pair:\n\n"
  f"Question: {question}\n"
  f"Correct Answer: {answer}\n"
  f"Predicted Answer: {pred}\n\n"
  "Provide your evaluation only as a detail orientation score where the detail orientation score is an integer value between 0 and 5,
  with 5 indicating the highest level of detail orientation. "
  "Please generate the response in the form of a Python dictionary string with keys 'score', where its value is the detail
  orientation score in INTEGER, not STRING."
  "DO NOT PROVIDE ANY OTHER OUTPUT TEXT OR EXPLANATION. Only provide the Python dictionary string. "
  "For example, your response should look like this: {'score': 4.8}."
}
```

Listing 4. ChatGPT-3.5 prompt for the contextual understanding (CU) metric.

```
"role": "system",
"content":
  "You are an intelligent chatbot designed for evaluating the contextual understanding of generative outputs for video-based question
  -answer pairs. "
  "Your task is to compare the predicted answer with the correct answer and determine if the generated response aligns with the
```

```

    overall context of the video content. Here's how you can accomplish the task:"
    "-----"
    "##INSTRUCTIONS: "
    "- Evaluate whether the predicted answer aligns with the overall context of the video content. It should not provide information
      that is out of context or misaligned.\n"
    "- The predicted answer must capture the main themes and sentiments of the video.\n"
    "- Consider synonyms or paraphrases as valid matches.\n"
    "- Provide your evaluation of the contextual understanding of the prediction compared to the answer."
  },
  {
    "role": "user",
    "content":
      "Please evaluate the following video-based question-answer pair:\n\n"
      f"Question: {question}\n"
      f"Correct Answer: {answer}\n"
      f"Predicted Answer: {pred}\n\n"
      "Provide your evaluation only as a contextual understanding score where the contextual understanding score is an integer value
        between 0 and 5, with 5 indicating the highest level of contextual understanding. "
      "Please generate the response in the form of a Python dictionary string with keys 'score', where its value is contextual
        understanding score in INTEGER, not STRING."
      "DO NOT PROVIDE ANY OTHER OUTPUT TEXT OR EXPLANATION. Only provide the Python dictionary string. "
      "For example, your response should look like this: {'score': 4.8}."
  }

```

B. Additional Experiments

B.1. Ablation Studies on ∞-Video LLaMA

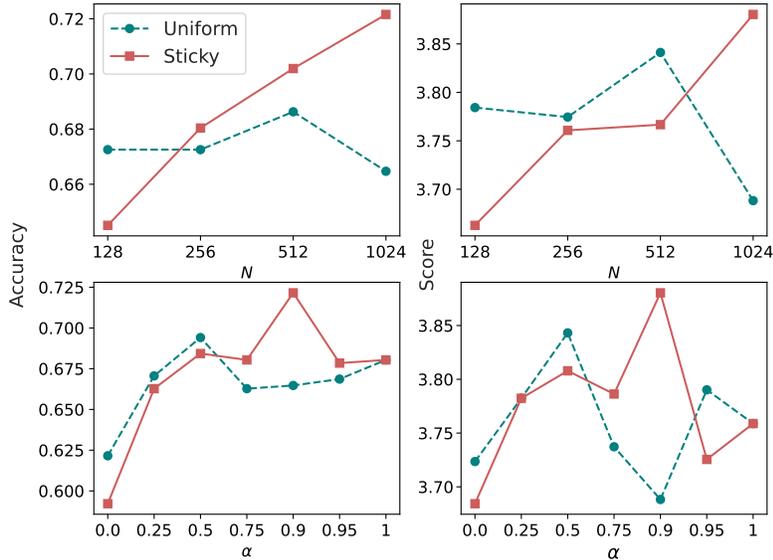


Figure 5. Ablation studies on the MovieChat dataset: Evaluation of accuracy and score metrics for various values of the number of basis functions N and the contribution of long-term memory α .

In this section, we conduct ablation studies on ∞-Video LLaMA on the MovieChat dataset. We ablate three hyperparameters: the percentage of the long-term memory used α , the number of basis functions N and the sampling method. We explore $\alpha \in \{0, 0.25, 0.5, 0.75, 0.95, 1\}$, covering the full spectrum from exclusively using the LTM ($\alpha = 0$) to exclusively using the STM ($\alpha = 1$). Additionally, we vary the number of basis functions with $N \in \{128, 256, 512, 1024\}$ to evaluate the impact of this parameter on performance and vary the sampling as either uniform or sticky.

Fig. 5 presents the accuracy and score provided by ChatGPT 3.5 as functions of the explored hyperparameters. The results reveal a general trend of increasing accuracy and score with α , up to a certain point, after which a slight decline is observed as α approaches 1. This trend is somewhat explained by the inherent variability in ChatGPT’s outputs. Additionally, both accuracy and score tend to be higher for sticky memories compared to uniform sampling from $\alpha = 0.75$ onward, whereas uniform sampling demonstrates superior performance for lower values of α . Another observed trend is that as the number of basis functions increases, both metrics improve. However, for uniform sampling, performance slightly decreases beyond $N = 512$.

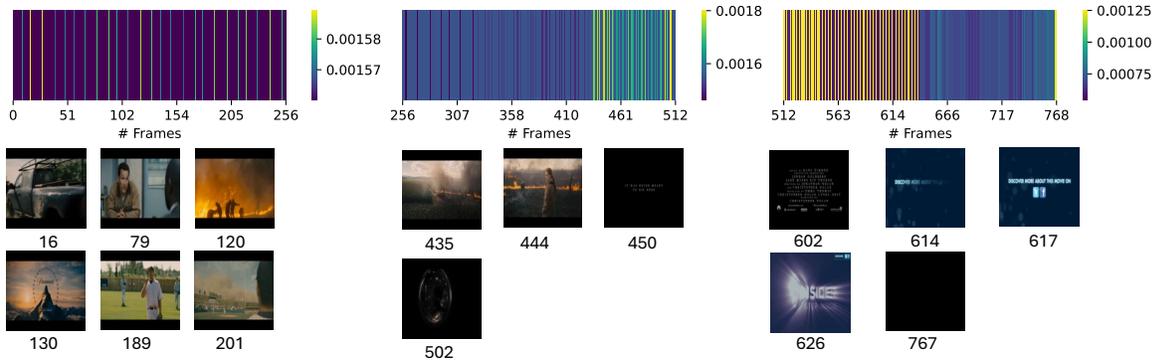


Figure 6. Highest continuous attention density frames selected using uniform memories in the *Interstellar* trailer for ∞ -Video LLaMA across 3 chunks. (Left) Interval: $[0, \tau^2]$. (Middle) Interval: $(\tau^2, \tau]$. (Right) Interval: $(\tau, 1]$.

B.2. Qualitative Analysis

In Fig. 6, we illustrate the attention density as a function of the number of frames for the uniform sampling LTM configuration, using $\alpha = 0.9$, $N = 256$, and $\tau = 0.5$. This analysis spans 3 chunks of 256 frames each, corresponding to 3 contraction steps. Additionally, we showcase representative frames from high-density regions by identifying the top 10 frames in each interval and selecting non-redundant examples.

The results reveal that, while the model effectively identifies key moments during the first two contraction steps, it disproportionately focuses on the credits scene in the final step. This behaviour contrasts with the results in Fig. 4, where such focus is avoided.

We leave in Fig. 7 additional examples of predictions of our modified ∞ -Video LLaMA. We divide the video into 8 chunks of 256 frames with $N=1024$, $\tau = 0.75$ and $\alpha = 0.9$ both for uniform sampling and sticky memories.

3 min Video
Interstellar trailer



User: What is the color of the jacket worn by the main character in the video?

Sticky Assistant: In the video, *the man wearing the brown jacket is standing near a black background with a black sky. The color of the jacket is brown.*

Uniform Assistant: *The man in the video is wearing a brown jacket.*

User: Which vehicles are visible in the video?

Sticky Assistant: *There is a truck, a car, a motorcycle, and a bicycle visible in the video.*

Uniform Assistant: *In the first scene of the video, we can see a dark sky and a black background with people in the background... Finally, in the ninth scene, we can see a dark sky and a black background with a car driving down the street on a dark night.*

User: Where is the spaceship with the astronauts entering?

Sticky Assistant: *In the video, the astronauts enter the spaceship and the spaceship is on a dark black background. The astronauts are wearing spacesuits, and the spaceship is carrying them into the atmosphere.*

Uniform Assistant: *In the first scene of the video, we can see a dark sky with a bright object in the center. The object appears to be a star, and we can see a small airplane flying in the sky (...)*

User: What emotions can be inferred from the girl's expressions when her dad leaves home?

Sticky Assistant: *In this video, the girl is sitting in a room with a door open, looking sad as her father leaves home. Her facial expression is a combination of disappointment, sadness, and fear. She is sad that her father is leaving, and she feels that she is alone in the room. She may be worried about how she will cope with being alone, or she may miss her father and be sad that he is leaving.*

Uniform Assistant: *The girl in the video looks sad when her father leaves, and her expression is difficult to read. However, the concept of family and the bond between parent and child can be inferred from the emotional aspect of the situation. Additionally, the girl's appearance and posture are not specified, but the video does depict a young girl in a tank top on a black background.*

2 hour Movie
Bohemian Rhapsody



User: What does the main character's hair look like in the first scene?

Sticky Assistant: *In the first scene, the main character has curly brown hair. He is wearing a blue sweater, and there is a brown couch in the background.*

Uniform Assistant: *In the first scene, the main character is a man with curly hair wearing a suit jacket and is sitting at a table in a room with a mirror on the wall. He has a laptop on the table, and there is also a window on the wall. Additionally, a young man is seen wearing a white jacket sitting at a table in the same room with a black and white photograph on the wall.*

User: Is the first performance indoors or outdoors?

Sticky Assistant: *The video shows a performance taking place inside a dark room. So, it is indoors.*

Uniform Assistant: *Based on the given video description, it is difficult to determine if the performance is indoors or outdoors as there are not specific descriptions of the location. However, the video has several objects such as a dark room, a white light, a silver plate, and a light that could provide some insight.*

User: What is the band's genre of music as shown in the video?

Sticky Assistant: *The video shows a band playing music in a dark room. So, the band's genre of music can be assumed to be rock or alternative music.*

Uniform Assistant: *Based on the video, it appears that the band in the video primarily plays rock music, although it is not explicitly mentioned.*

User: Which instrument does the main character play in the last live concert?

Sticky Assistant: *In the last scene, the main character plays the keyboard.*

Uniform Assistant: *In the last live concert, the main character is playing a guitar.*

Figure 7. Examples of ∞-Video LLaMA answers with uniform sampling and sticky memories for short and ultra-long videos. Italicized corresponds to the correct answer while underlined corresponds to the wrong answer or hallucination.