

Self-ReS: Self-Reflection in Large Vision-Language Models for Long Video Understanding

João Pereira^{1,2}, Vasco Lopes^{2,3}, David Semedo¹, João Neves³

¹NOVA LINCS, NOVA@FCT, Caparica, Portugal.

²DeepNeuronic, Covilhã, Portugal.

³NOVA LINCS, University of Beira Interior, Covilhã, Portugal.
jaca.pereira@campus.fct.unl.pt

Abstract—Large Vision-Language Models (LVLMs) demonstrate remarkable performance in short-video tasks such as video question answering, but struggle in long-video understanding. The linear frame sampling strategy, conventionally used by LVLMs, fails to account for the non-linear distribution of key events in video data, often introducing redundant or irrelevant information in longer contexts while risking the omission of critical events in shorter ones. To address this, we propose Self-ReS, a non-linear spatiotemporal self-reflective sampling method that dynamically selects key video fragments based on user prompts. Unlike prior approaches, Self-ReS leverages the inherently sparse attention maps of LVLMs to define reflection tokens, enabling relevance-aware token selection without requiring additional training or external modules. Experiments demonstrate that Self-ReS can be seamlessly integrated into strong base LVLMs, improving long-video task accuracy and achieving up to 46% faster inference speed within the same GPU memory budget. The code can be found here: <https://github.com/jaca-pereira/Self-ReS>.

Index Terms—Large Vision-Language Models, Video Understanding

I. INTRODUCTION

Recently, Large Vision-Language Models (LVLMs) emerged as a way to leverage the strong zero- and few-shot performance of Large Language Models (LLMs) in downstream vision tasks. These models combine vision and language by projecting visual features from a vision encoder into the LLM’s embedding space, followed by fine-tuning on multimodal data. In the video domain, the adoption of strong backbone models [1] [2] and the development of high-quality pre-training and instruction-tuning datasets [3] has led to improved performance on tasks such as video question answering [4]. While LVLMs demonstrate strong capabilities in understanding short videos, their performance on longer videos remains a significant challenge, due to difficulties in modeling larger numbers of frames. In the reference benchmark Video-MME [4], several state-of-the-art models such as LLaVA-Video [3] report strong performance for short videos, but struggle to obtain competitive performance for longer inputs, especially without the textual hints present in subtitles.

Efforts to improve performance for longer videos can be categorized into two main approaches: a) increasing context




Fig. 1. **Linear Frame Sampling vs Self-Reflective Sampling (Self-ReS).** The linear frame sampling method, used by most LVLMs, samples a dense or uniform subset of frames, increasing the likelihood of sampling redundant and background information or missing key events. In contrast, Self-ReS increases sampling granularity, optimizing context-window usage with prompt-relevant fragments of the video identified through spatiotemporal self-reflection.

size, and b) applying token compression techniques. Often, models propose a combination of the two. Despite significantly increasing the number of sampled frames - e.g., up to 768 for Qwen2-VL [5] and 1024 for NVILA [6]- and applying aggressive token compression to minimize the cost of reasoning over larger contexts, these models still struggle with long videos, being unable to outperform models with significantly smaller context sizes - e.g., LLaVA-Video with 64 frames only.

We argue that current frame sampling strategies overlook video non-linearity, where critical events, such as a sudden action or a key object appearance, occur in a concentrated spatiotemporal region, rather than being evenly distributed throughout the video - e.g., Figure 1. Moreover, linearly scaling the number of frames to accommodate longer videos will likely introduce background and redundant information, inefficiently allocating more computational resources.

An important line of research focuses on proposing improved sampling strategies for video LVLMs. Widely adopted techniques such as token pooling [7] [6] or merging [8] allow for extended context sizes, but also fail to account for the non-linear nature of video data. More recent approaches attempt to address this non-linearity through dynamic memory banks [9], adaptive token compression [10] [11] or by learning a summarization token [12]. Notably, the proposed solutions rely on external, trainable modules, introducing added computational requirements that may affect their adoption in real-world scenarios.

In this paper, we introduce **Self-ReS**, a novel *Self-Reflective Sampling* approach designed to enhance LVLM performance

This work is supported by NOVA LINCS ref. UIDB/04516/2020 (<https://doi.org/10.54499/UIDB/04516/2020>) and ref. UIDP/04516/2020 (<https://doi.org/10.54499/UIDP/04516/2020>) with the financial support of FCT/IP; and Fundação para a Ciência e Tecnologia ref. 2023.03647.BDANA.

on long video understanding. Self-ReS addresses video non-linearity by selecting key spatiotemporal fragments, guided by textual cues present in the user prompt. Our approach builds on the idea that LLMs can self-reflect and critique their own generations [13] to improve their final answer. In particular, Self-ReS leverages the inherent self-reflective ability of base LLMs to identify relevant video fragments according to a user-specified task.

Self-ReS introduces independent reflection paths, where each path processes, in parallel, a different segment of the video. These paths converge at specific depths within the LLM, enabling unified inference across the entire video. The self-reflective nature of the LLM is expressed in *reflection tokens*, which serve as judges to evaluate the relevancy of visual tokens based on the user query. By leveraging the inherent ability of LVLMs to generate sparse attention maps [14], which naturally focus on a subset of visual tokens, Self-ReS defines reflection tokens to prioritize critical information. Once the reflection paths converge, the LLM performs inference using its default context size, concentrating on the most salient visual tokens, which are non-linearly selected according to the textual cues in the user prompt. With this strategy, Self-ReS achieves up to **46%** faster inference speed and higher accuracy than the base model while maintaining GPU memory budget.

Our contributions can be summarized as follows:

- We propose **Self-ReS**, a *Self-Reflective Sampling* method that dynamically identifies and prioritizes key spatiotemporal fragments of video data. By leveraging the inherent self-reflective capabilities of LLMs, Self-ReS addresses spatiotemporal non-linearity without requiring additional training or external modules.
- Self-ReS enables LVLMs to process significantly more frames while reducing computational overhead. It achieves higher accuracy with up to **46% faster inference speed** and identical GPU memory usage, making it practical for resource-constrained scenarios.
- Using Self-ReS, LLaVA-Video outperforms existing open and closed-source models on two reference video benchmarks.

II. RELATED WORK

A. Large Vision-Language Models

LVLMs extend the zero-shot and instruction-following capabilities of LLMs to visual inputs, enabling multimodal tasks across images and videos. Foundational LVLMs such as LLaVA [15] and BLIP [16] connect vision encoders to LLMs via linear projections or learned queries, forming the basis for pioneer video-specific models such as VideoChat-GPT [7]. These models are pre-trained on video-captioning and instruction-tuning datasets [3], enabling strong performance in short video tasks.

Recent advancements scale LVLMs to longer videos using state-of-the-art backbones [1] [2] and increased frame context sizes [6] [5]. However, these improvements often come at the cost of significant computational overhead. Moreover, these

methods fail to address spatiotemporal non-linearity in video data, where key events are concentrated in small regions rather than distributed uniformly. Self-ReS enables foundational models to process longer videos within the default context size, dynamically prioritizing salient tokens while maintaining computational efficiency.

B. Video Sampling Strategies for LVLMs

Sampling strategies are critical for extending LVLM performance to longer videos. Early approaches such as VideoChat-GPT [7] and ChatUniVi [8] employed token pooling or merging to handle the larger number of video frames. Memory-based strategies such as MovieChat [9] store long context video information in memory tokens. LLaMA-ViD [17] learns a single token per frame. Recent proposals include learnable video summarization tokens in VideoXL [12] and spatiotemporal adaptive token compression in LongVU [10], which uses high-resolution vision encoders to compute frame similarity and cross-modal relevance. While effective to an extent, these methods often rely on additional backbones and require extensive training [10] [12] [17], increasing computational costs. Moreover, they may fail to adequately address video non-linearity [9] [7] [8], where key events occur in sparse spatiotemporal fragments of the video. In contrast, Self-ReS leverages the LLM’s self-reflective abilities, expressed in the form of native attention weights that dynamically identify and prioritize critical visual tokens. Our plug-and-play approach eliminates the need for external modules or additional training, offering a computationally efficient solution that is faster than the baseline and surpasses prior methods in accuracy.

III. METHODOLOGY

In this section, we detail the methodology behind Self-ReS, illustrated in Figure 2. Self-ReS is designed to address three key challenges faced by LVLMs in long video understanding: 1) **Non-Linearity of video data**, where critical events, such as short actions or key object appearances, are often concentrated in small spatial and temporal regions of a video, rather than being evenly distributed; 2) **Limited context sizes** which restrict the number of frames that can be processed, risking the omission of key events in shorter temporal windows; 3) **Linearly increasing context size** introduces **background noise and redundant information**, leading to inefficient resource utilization.

Self-ReS addresses these challenges through a process called **spatiotemporal self-reflection**. This process begins by sampling a large number of frames from a long video and distributing them across parallel reflection paths. Each path processes a segment of the video independently up to a specified layer depth in the LLM. The paths are then merged into a single inference path, where the most salient visual tokens are non-linearly selected based on the user task. By dynamically prioritizing key information and minimizing redundancy, Self-ReS effectively expands the context size without inefficiently scaling computational resources.

A. Spatiotemporal Self-Reflection

Inspired by the concept that **LLMs can critique their own predictions** through parallel and independent inferences [13], we design reflection tokens to judge the relevancy of visual tokens for a given task. Our approach leverages the observation that the **last input token of LVLMs inherently generates sparse attention maps for visual tokens** [14]. By adopting the last input token as the reflection token, Self-ReS avoids the need for additional training.

Without loss of generality, we assume a standard LVLM architecture composed of a vision encoder f_V , a linear projection matrix $W \in \mathbb{R}^{d_v \times d}$, and an LLM f , where d_v and d correspond to the visual encoder and LLM token dimensionality, respectively. We also assume the default input sequence formatting, with a system prompt X_{sys} being concatenated with the video features, followed by the user question X_q .

To address the context size limitations, Self-ReS emulates a larger context size by sampling T frames from a long video and distributing them across R independent reflection paths. Each path processes $S = \frac{T}{R}$ frames, where S matches the base model's default context size. The frames are first processed by the Vision Encoder f_V , generating visual features $X_V \in \mathbb{R}^{R \times N_V \times d_v}$, where N_V represents the number of visual tokens per reflection path. These features are then projected into the LLM embedding space using the matrix W , producing $\hat{X}_V \in \mathbb{R}^{R \times N_V \times d}$. The inputs for each reflection path r_i are defined as:

$$X_{r_i} = [X_{\text{sys}} \oplus \hat{X}_V^i \oplus X_q], \quad (1)$$

with $X_{r_i} \in \mathbb{R}^{N_r \times d}$, where $\hat{X}_V^i$ corresponds to the projected visual tokens of path r_i , N_r is the total number of tokens per reflection path, and $\oplus$ denotes concatenation.

Each reflection path sequence X_{r_i} is forwarded independently, thus having its own attention mask and position embeddings. Depending on the available memory resources, Self-ReS can operate in parallel (batch mode) or sequentially (memory-efficient mode), ensuring scalability across diverse computational environments. For simplicity, we assume parallel inference mode in this explanation.

Our self-reflective approach begins with the LLM's prefill stage, performing a forward pass up to a designated self-reflection layer. At this layer, Self-ReS applies **Self-Reflective Attention (SRA)**, a modified version of Scaled Dot-Product Attention designed to perform a non-linear selection of visual tokens. This mechanism determines query-key similarities without the memory-intensive value multiplication and softmax operations. Also, introducing a new attention mechanism only at the self-reflection layer allows compatibility with all existing LVLM setups, including those using Flash Attention [18], which typically does not expose attention weights.

Specifically, we designate the hidden embedding of the last token of the user query X_q , at self-reflection layer l , as our reflection token. The reflection token is then projected as the query vector $q_l \in \mathbb{R}^{N_V \times d_q}$. Likewise, the hidden embeddings of the visual tokens are projected into the key matrix $K_{V_l} \in \mathbb{R}^{N_V \times d_k}$, where $d_k = d_q$. SRA is computed as:




Fig. 2. **Illustration of Self-ReS methodology.** Self-ReS splits the long-context video into reflection paths, performing batch inference to a specific depth. Reflection tokens then assess the relevancy of visual tokens for the user task. After self-reflection, the paths merge into a single inference path using either Regular or Smooth Inference modes. Regular mode samples salient tokens and restarts the forward pass from the first layer, while Smooth mode continues from the self-reflection layer with the most salient hidden embeddings.

$$\text{SRA}(q_l, K_{V_l}) = \frac{q_l \cdot K_{V_l}^T}{\sqrt{d_k}}. \quad (2)$$

We purposely omit the attention head dimension for simplicity. The final saliency scores are computed by averaging the scores for each token over all attention heads, isolating individual contributions.

The last stage of spatiotemporal self-reflection utilizes the judgments produced by SRA to evaluate the relevancy of visual tokens across reflection paths. These judgments are derived from the scores calculated from the sparse attention maps of the reflection tokens. The expressiveness of these scores arises from the fact that reflection tokens q_l act as highly informational embeddings for next-token prediction [14].

Once the attention weights are calculated, Self-ReS selects the top- N_V visual tokens:

$$S_c = \text{Top}_{N_v} \left(\left\{ \text{SRA}_{r_i}(q_l, K_{V_l}) \right\}_{r_i=1}^{r_i=R} \right), \quad (3)$$

where S_c contains the indices of the most salient visual tokens over all reflection paths. This selection process introduces non-linearity in both the spatial and temporal dimensions and eliminates background information, retaining only the most critical tokens. The stored indices in S_c are then used to sample tokens in the next phase, where the reflection paths converge into a unified inference path.

B. Converging Reflection Paths

The final stage of Self-ReS merges the reflection paths into a single inference path, enabling the model to operate within its default context size while reasoning over the non-linear salient video fragments selected during spatiotemporal self-reflection. This step allows the model to effectively capture relationships across video segments. To achieve this, Self-ReS implements two complementary strategies for converging reflection paths, tailored to balance computational efficiency and full model reasoning.

a) *Regular Inference*: This mode initiates a new forward pass from the first layer using only the top- N_v visual tokens selected during the spatiotemporal self-reflection phase. The input sequence is constructed as:

$$X = [X_{\text{sys}} \oplus \hat{X}_r \oplus X_q], \quad (4)$$

where $\hat{X}_r \in \mathbb{R}^{N_v \times d}$ represents the set of projected visual tokens with the highest scores S_c , from the flattened sequence of visual tokens across all reflection paths. The original temporal order of the visual tokens is maintained, enabling the LLM to perform a full inference pass on the selected salient tokens.

b) *Smooth Inference*: This mode leverages the hidden embeddings after the Self-Reflection layer l , and continues the forward pass with the top- N_v visual tokens, instead of restarting from the first layer. We define the input sequence for the final inference path as:

$$H = [H_{\text{sys}_0} \oplus H_r \oplus H_{q_0}], \quad (5)$$

where H_{sys_0} and H_{q_0} are the hidden embeddings of the system and user tokens from the first reflection path, and $H_r \in \mathbb{R}^{N_v \times d}$ represents the top- N_v self-reflected hidden embeddings of salient visual tokens, conditioned on the user query across all reflection paths. As shown in Figure 3, the hidden embeddings of text tokens generated with the same prompt across different reflection paths converge to equivalent representations. We wield this insight to define our smooth converging strategy and select the textual tokens of a single reflection path. Smooth Inference minimizes redundant computations, maintaining high inference efficiency.

c) *Final Inference Path*: Regardless of the chosen strategy, Self-ReS converges into a unified new inference path sequence. This sequence is forwarded using vanilla attention and reassigned positional embeddings (see section A of the Appendix). This path focuses exclusively on the most salient visual tokens, identified through the LLM’s self-reflective capabilities guided by the user query. By leveraging these tokens, the final model prediction effectively captures global inter-segment dependencies, enabling reasoning over a larger context of the input video. Notably, this is achieved without actually increasing the model’s default context size or memory requirements, ensuring efficient handling of non-linear video events.

C. Complexity of Self-ReS

Our proposed SRA effectively computes saliency scores for visual tokens, similar to regular attention mechanisms but with reduced computational overhead. In our baseline LVLM, the computational cost of attention is quadratic over the total number of input tokens - $O(R * N_r^2)$, where R is the number of reflection paths and N_r is the number of input tokens per reflection path. In Self-ReS, while attention computation remains quadratic, it is computed over a smaller number of input tokens split between reflection paths: $O(R * N_r^2) = O(R * (N/R)^2) = O(N^2/R)$. Thus, while Self-ReS maintains quadratic complexity, it is reduced by a factor of R . SRA




Fig. 3. **Consistency of system and textual hidden representations over reflection paths.** The hidden representations of the system and textual input have been reduced to a low-dimensional space (2D), evidencing the residual difference between the hidden representations over different reflection paths.

(Eq. 2) has linear complexity $O(N_v)$, which is then applied to each reflection path $O(R * N_v)$. Therefore, it does not affect the overall complexity of the LVLM inference flow.

IV. EXPERIMENTS

In this section, we thoroughly evaluate the impact of integrating Self-ReS in a state-of-the-art foundational LVLM using two reference video understanding benchmarks. We benchmark Self-ReS against existing foundational LVLMs and state-of-the-art sampling methods to demonstrate its effectiveness in improving accuracy and efficiency.

A. Datasets.

a) *Video-MME [4]*: Encompasses multiple choice questions over challenging short, medium and long videos. The videos cover a wide range of topics including Knowledge, Film & Television, Sports Competition, Artistic Performance, Life Record, and Multilingual. We focus on the benchmark without subtitles because it closely simulates real-world scenarios where subtitles are unavailable, and ensures that the evaluation better reflects the model’s ability to comprehend visual spatiotemporal information.

b) *MLVU [19]*: The Multi-task Long Video Understanding benchmark is used to evaluate performance of the proposed strategy on diverse and complex long-video understanding tasks. MLVU comprises a collection of videos ranging from 3 minutes to 2 hours, spanning diverse genres such as movies, documentaries, and surveillance footage. The benchmark features nine distinct task categories, including topic reasoning, anomaly recognition, summarization, and action understanding, designed to test holistic, single-detail, and multi-detail video comprehension.

B. Implementation Details

For all experiments, we use a greedy decoding strategy. These values ensure deterministic and reproducible results. For Self-ReS specific parameters, we explore different combinations of total frames sampled (T), segment size (S), and the self-reflection layer depth, as detailed in the ablation study. We rely on an open-source library [20] to implement SRA at a designated self-reflection layer and batch processing for reflection paths.

TABLE I
PERFORMANCE COMPARISON ON VIDEO-MME AND MLVU[†].

	Model	Sampling Strategy	T	Video-MME		MLVU
				Long	All	
Large >= 72B	VideoLLaMA2 [23]	Q-Former	32	57.6	62.4	45.6
	LLaVA-OV [24]	Pooling	32	60.0	66.3	68.0
	GPT-4o mini [22]	-	250	58.6	64.8	-
	GPT-4o [21]	-	384	65.3	71.9	64.6
Mid-Sized 7B	MovieChat [9]	Memory	2048	-	-	18.0
	LLaMA-ViD [17]	CrossAt	1fps	-	-	17.2
	ChatUniVI [8]	Merging	64	35.8	40.6	-
	VideoXL [12]	VST	2048	55.5	61.0	64.9
	LongVU [10]	SVA	1fps	-	-	65.4
	Oryx [11]	CrossAt	128	-	58.8	67.5
	Qwen2-VL [5]	MLP	738	-	63.3	-
	NVILA [6]	Pooling	1024	54.8	64.2	70.1
	LLaVA-Video [3]	Pooling	32	52.2	62.5	66.5
			64	53.4	64.3	70.3
			128	54.1	64.6	70.3
	LLaVA-Video	Self-ReS	128 (32) 128 (64)	54.6 56.3	64.3 65.5	67.9 70.7

[†] T denotes the total number of frames sampled for input. VST is Video Summarization Token, SVA is Spatiotemporal Visual Aggregation and CrossAt is Cross-Attention. '-' means value not reported.

C. Long Video Understanding

Table I compares the performance of Self-ReS against other LVLMs with different sampling strategies. Results clearly demonstrate the superiority of Self-ReS in comparison to methods using similar sized LLMs, achieving state-of-the-art results. LLaVA-Video equipped with Self-ReS surpasses both the base model and models with much larger context sizes. Additionally, our spatiotemporal self-reflective sampling methodology helps LLaVA-Video outperform much larger closed and open source methods, including GPT-4o [21] and GPT-4o mini [22].

These results evidence the benefit of non-linear token selection at both the spatial and temporal dimensions of the video. Self-ReS uses non-linear spatiotemporal self-reflection to compress task-relevant information into a significantly smaller context window, contrasting with current approaches that linearly extend the context across numerous frames, leading to redundancy and background noise. Furthermore, our training-free approach, leveraging the naturally sparse attention maps of the last user query token, allows for our method to surpass other trainable, non-linear sampling methods, by building on and improving a strong base model.

D. Ablation Study

This section analyzes the impact of Self-ReS's parameters on accuracy. Figure 4 reveals the effect of the depth of the self-reflection layer in the performance of Self-ReS over long videos, using the Video-MME (Long) [4] benchmark. For consistency, we fix the segment size S to 32 and the total number of sampled frames T to 64.

We focus our ablation study in the shallow layers of the LLM. Considering the role of different layers in the forward pass of an LLM, early layers typically capture syntactic or input-level features, which are better for predictive tasks [25] such as our designated spatiotemporal self-reflection, aiming




Fig. 4. **Ablation study on self-reflection layer depth (left) and context size width (right).** Experiments reveal that optimal self-reflection performance is achieved in earlier layers. Emulating larger context sizes with increased T values also benefits accuracy.

to predict scores for visual tokens. We leave most deeper layers of the LLM out of our study, as they provide semantic and abstract representations geared towards text generation.

Our results demonstrate that the *Regular Inference* strategy, where the model resets the forward pass after self-reflection, observes a steady behavior throughout different depths. However, our analysis reports peak accuracy for *Smooth Inference* at depth 5, and a rapid decline in deeper layers, corroborating the claims in [25]. Since the forward pass continues after self-reflection, and the visual hidden embeddings from different reflection paths are concatenated, the hidden embeddings of earlier layers represent more input-specific representations, while later layers represent abstract features that are relative for the answer being generated for each reflection path, and therefore are both worse for predicting visual token scores and for path convergence. *Regular Inference* is more robust to this factor, since the forward pass restarts from the first layer using input tokens.

We also conduct additional experiments in Video-MME (Long) using varied configurations of S and T , with a fixed self-reflection layer depth of 5 for optimal performance-efficiency trade-off. The results are provided in Table 4 and indicate a steady gain in accuracy by emulating a larger context size of T frames and consequently more reflection paths R . This experiment confirms that we can effectively emulate larger context sizes within the base model's default context window by performing non-linear sampling.

E. Computational Efficiency Analysis

This section evaluates the impact of Self-ReS on the inference speed of the baseline LLaVA-Video [3] model. We measure these metrics on 70 videos from the Video-MME [4] benchmark, after a warm-up round of 30 videos, using a setup of 16GB of RAM and one NVIDIA A100 GPU with 40GB of VRAM.

Figure 5 plots the accuracy in Video-MME (Long) and frames per second (fps) for Self-ReS and the base model across different context sizes. For Self-ReS, we fix the segment size S to 32 and 64, and vary the number of reflection paths R , emulating a larger context size of T frames. We use Self-ReS in batch mode, ensuring a fair comparison to the baseline. We purposely omit GPU memory consumption since we observed similar values when using batch mode. For information on




Fig. 5. Comparison of accuracy and inference speed on Video-MME for base model with and without Self-ReS, across varying context sizes. Results show evidence that Self-ReS has a positive effect on both the efficiency and accuracy of the base model.

inference speed and GPU memory consumption when using the memory efficient sequential mode, we refer the reader to the supplementary material.

We observe up to 46% faster inference speed and 2.2% higher accuracy compared to the base model when using an emulated context size of $T = 128$. As a result of our non-linear spatiotemporal sampling method, our solution can compress task-relevant information into a much smaller context size. Across all setups, Self-ReS achieves similar or superior accuracy in faster inference speed.

F. Qualitative Evaluation

Our proposed spatiotemporal self-reflection process, guided by the user prompt, can effectively focus on the small spatiotemporal window where the main event occurs. In Figure 1, we see that Self-ReS allows the model to correctly infer that the man blowtorches the door handles, despite this event happening in a small spatial and temporal window of the video. Other sampling strategies that perform linear sampling can introduce noise - e.g., frames that are not related to the man or the door handles - or miss key details - e.g., frames where the man interacts with the door lock and not the handles.

V. CONCLUSION

In this paper, we introduce **Self-ReS**, a novel self-reflective sampling method tailored to address the challenges of long-video understanding in LVLMs. By leveraging the inherent self-reflective abilities of LLMs, Self-ReS selects key spatiotemporal fragments guided by an user query, effectively addressing video non-linearity without requiring additional training or external modules. Experiments on Video-MME and MLVU benchmarks demonstrate that Self-ReS achieves state-of-the-art results among mid-sized LVLMs, surpassing models with significantly larger context sizes while reducing base model inference speed by up to 46%. Self-ReS can be seamlessly integrated into existing LVLm architectures, enabling more efficient and accurate video understanding. In future work, we plan to assess whether Self-ReS can cope

with augmented textual contexts and elicit chain-of-thought reasoning.

REFERENCES

- [1] Xiaohua Zhai, Basil Mustafa, and Alexander et al. Kolesnikov, “Sigmoid loss for language image pre-training,” in *ICCV*, 2023.
- [2] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, and et.al, “Qwen2 technical report,” *arXiv:2407.10671*, 2024.
- [3] Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li, “Video instruction tuning with synthetic data,” *arXiv:2410.02713*, 2024.
- [4] Chaoyou Fu, Yuhao Dai, and Yondong et al. Luo, “Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis,” *arXiv:2405.21075*, 2024.
- [5] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, and Jinze Bai et al., “Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution,” *arXiv:2409.12191*, 2024.
- [6] Zhijian Liu, Ligeng Zhu, Baifeng Shi, Zhuoyang Zhang, Yuming Lou, and Shang Yang, “Nvlla: Efficient frontier visual language models,” *arXiv:2412.04468*, 2024.
- [7] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Khan et al., “Video-chatgpt: Towards detailed video understanding via large vision and language models,” in *ACL*, 2024.
- [8] Peng Jin, Ryuichi Takanobu, Wancai Zhang, and et al. Cao, “Chat-univ: Unified visual representation empowers large language models with image and video understanding,” in *CVPR*, 2024.
- [9] Enxin Song, Wenhao Chai, Guanhong Wang, and et al. Zhang, “Moviechat: From dense token to sparse memory for long video understanding,” in *CVPR*, 2024.
- [10] Xiaoqian Shen, Yuniang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, and Chenchen Zhu et al., “Longvu: Spatiotemporal adaptive compression for long video-language understanding,” *arXiv:2410.17434*, 2024.
- [11] Zuyan Liu, Yuhao Dong, Ziwei Liu, and Hu, “Oryx mllm: On-demand spatial-temporal understanding at arbitrary resolution,” *arXiv:2409.12961*, 2024.
- [12] Yan Shu, Peitian Zhang, and Zheng Liu, “Video-xl: Extra-long vision language model for hour-scale video understanding,” *arXiv:2409.14485*, 2024.
- [13] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi, “Self-RAG: Learning to retrieve, generate, and critique through self-reflection,” in *ICLR*, 2024.
- [14] Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang, “An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models,” in *ECCV*, 2024.
- [15] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee, “Visual instruction tuning,” in *NeurIPS*, 2023.
- [16] Junnan Li, Dongxu Li, Silvio Savarese, and Hoi et al., “BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models,” in *ICML*, 2023.
- [17] Yanwei Li, Chengyao Wang, and Jiaya Jia, “Llama-vid: An image is worth 2 tokens in large language models,” 2024.
- [18] Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré, “Flashattention: Fast and memory-efficient exact attention with io-awareness,” *arXiv:2205.14135*, 2022.
- [19] Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, and Xi Yang et al., “MLVU: A comprehensive benchmark for multi-task long video understanding,” *arXiv:2406.04264*, 2024.
- [20] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, and Anthony Moi et. al., “Huggingface’s transformers: State-of-the-art natural language processing,” *arXiv:1910.03771*, 2020.
- [21] OpenAI, “Gpt-4o,” 2024, Accessed: 2024-09-07.
- [22] OpenAI, “Gpt-4o mini,” 2024, Accessed: 2024-12-16.
- [23] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, and Lidong Bing, “Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms,” *arXiv:2406.07476*, 2024.
- [24] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li, “Llava-onevision: Easy visual task transfer,” *arXiv:2408.03326*, 2024.
- [25] Zhu Liu, Cunliang Kong, and Ying et al. Liu, “Fantastic semantics and where to find them: Investigating which layers of generative llms reflect lexical semantics,” *arXiv:2403.01509*, 2024.

APPENDIX

A. Reconstructing Context After Convergence

We reconstruct the context using only the most relevant visual tokens, reassigning position embeddings as a sequential 1D index. Our strategy shifts the positions of tokens beyond the first reflection path without disrupting their spatial-temporal structure. LVLs process inputs as 1D token sequences, and spatial structure is only enforced by the vision encoder, which remains unchanged. Although token distances may be altered, their relative order is maintained, ensuring coherence in the final representation. We validate our approach using our best setup against three alternatives: 1) **Retaining original positional embeddings** – Results in a non-incremental sequence, duplicating positions for tokens from different reflection paths, and introduces varying distances between tokens; 2) **Retaining original positional embeddings incremented according to their reflection path** – Avoids duplicate tokens but introduces varying distances between tokens; 3) **Using a single position embedding** – Removes temporal position information from vision tokens. As shown in Table II, our original strategy of reassigning positional embeddings yields the best performance.

TABLE II
POSITION EMBEDDING ASSIGNMENT STRATEGY (PERFORMANCE ON VIDEO-MME).

Strategy	Long	Overall
1)	51.8	62.1
2)	52.3	62.3
3)	41.7	50.9
Ours	54.6	64.3

B. Sequential Mode

This section evaluates the impact of Self-ReS on the GPU memory budget and inference speed of the base LLaVA-Video [3] model. We measure these metrics on 70 videos from the Video-MME [4] benchmark, after a warm-up round of 30 videos, using the same setup used in batch mode, with 16GB of RAM and one NVIDIA A100 GPU with 40GB of VRAM.

Figure 6 plots the accuracy in Video-MME (Long) and maximum GPU memory consumption during inference for Self-ReS and the base model across different context sizes. For Self-ReS, we fix the segment size S to 32 and 64, and vary the number of reflection paths R , emulating a larger context size of T frames. We use Self-ReS in sequential mode, and split the reflection paths over two mini-batches. During spatiotemporal self-reflection, Self-ReS processes each mini batch sequentially in the GPU, storing the remaining mini-batch in RAM. After convergence into a single path, Self-ReS conducts inference in GPU only.

Results show Self-ReS reduces the maximum GPU memory consumption by up to **23.6%**. Notably, despite sequential processing, our approach is still **16%** faster than the base




Fig. 6. GPU Memory vs Accuracy.




Fig. 7. Inference Speed vs Accuracy.

model for a context of $T = 128$ - Figure 7. Therefore, Self-ReS allows faster inference over emulated larger context sizes in restrained GPU memory environments.

C. Expanding Baselines

We extend Self-ReS to LLaVA-OneVision [24] and show that our approach has a positive impact on performance (Table III).

TABLE III
PERFORMANCE OF SELF-RE S EXTENDED LLaVA-ONEVISION (VIDEOMME).

Model	Size	Long	Overall
LLaVA-OV [24]	7B	48.6	58.2
w/ Self-ReS		50.7	59.1

D. Additional Dataset Details

1) **Video-MME**: The Video-MME [4] dataset contains 900 videos with a total of 254 hours. The videos, ranging from 11 seconds to 1 hour, were manually selected and annotated by human experts, generating a total of 2,700 question-answer pairs. Questions are split across the following categories: Object Recognition, Action Recognition, OCR Problems, Counting Problems, Temporal Reasoning, Spatial Reasoning, Action

Reasoning, Object Reasoning, Information Synopsis, Temporal Perception, Spatial Perception and Attribute Perception.

2) *MLVU*: The MLVU [19] dataset contains a total of 1334 instances of long videos ranging from 3 minutes to 2 hours. The dataset contains 2593 questions, split into 2175 multiple-choice and 418 free-form generation questions. These questions were either automatically generated based on previously existing annotations –e.g., Anomaly Recognition category– or manually annotated by human experts.

E. Extended Performance on all Video Lengths

Self-ReS delivers superior performance across different video lengths. Table IV shows that Self-ReS improves performance for medium and long videos. In short videos, performance is maintained since most short videos require a single reflection path, resulting in the same setup as the base model.

TABLE IV
PERFORMANCE ACROSS VIDEO LENGTHS (VIDEOMME).

Model	T	Short < 2min	Medium 4-15 min	Long 30-60 min	Overall
Baseline	32	76.3	59.0	52.2	62.5
	64	77.2	62.1	53.4	64.3
Self-ReS	128 (32)	76.2	62.0	54.6	64.3
	128 (64)	77.2	63.7	56.3	65.5

F. Extra Qualitative Examples

Figures 8 and 9 provide additional visual examples of the effectiveness of the proposed approach. In Figure 8, the base sampling strategy misses the small spatial and temporal window where the road accident occurs. In Figure 9, the base sampling strategy samples a background picture of a camera, leading the model towards an incorrect answer. With Self-ReS integration, the model can leverage textual cues present in the user prompt and reason over key spatiotemporal fragments of the video that will guide the model towards the correct answer.




Fig. 8. **Linear Sampling vs Self-ReS**. With Self-ReS, the model can detect the crash.




Fig. 9. **Linear Sampling vs Self-ReS**. With Self-ReS, the model can focus on the visual cue indicating led screens with larger pixel pitch.