

VideoStir: Understanding Long Videos via Spatio-Temporally Structured and Intent-Aware RAG

Honghao Fu¹, Miao Xu¹, Yiwei Wang², Dailing Zhang³, Jun Liu⁴, Yujun Cai^{1*}

¹University of Queensland, ²University of California, Merced

³Institute of Automation, CAS, ⁴Lancaster University

{honghao.fu, miao.xu, yujun.cai}@uq.edu.au, yiweiwang2@ucmerced.edu

zhangdailing2023@ia.ac.cn, j.liu81@lancaster.ac.uk

Abstract

Scaling multimodal large language models (MLLMs) to long videos is constrained by limited context windows. While retrieval-augmented generation (RAG) is a promising remedy by organizing query-relevant visual evidence into a compact context, most existing methods (i) flatten videos into independent segments, breaking their inherent spatio-temporal structure, and (ii) depend on explicit semantic matching, which can miss cues that are implicitly relevant to the query’s intent. To overcome these limitations, we propose VideoStir, a structured and intent-aware long-video RAG framework. It firstly structures a video as a spatio-temporal graph at clip level, and then performs multi-hop retrieval to aggregate evidence across distant yet contextually related events. Furthermore, it introduces an MLLM-backed intent-relevance scorer that retrieves frames based on their alignment with the query’s reasoning intent. To support this capability, we curate IR-600K, a large-scale dataset tailored for learning frame–query intent alignment. Experiments show that VideoStir is competitive with state-of-the-art baselines without relying on auxiliary information, highlighting the promise of shifting long-video RAG from flattened semantic matching to structured, intent-aware reasoning. Codes and checkpoints are available at <https://github.com/RomGai/VideoStir>.

1 Introduction

Video understanding has emerged as a core frontier in multimodal intelligence (Li et al., 2025a). While recent Multimodal Large Language Models (MLLMs) have achieved remarkable success in interpreting short clips, extending this capability to long videos remains a formidable challenge (Luo et al., 2025). The primary difficulty lies not merely in the video’s length, but in its information density and structural complexity (Shen et al., 2025a).

*Corresponding Author.

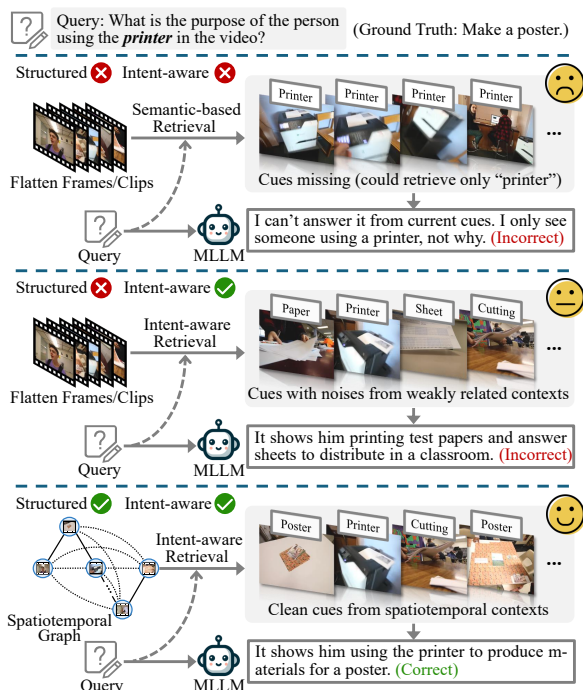


Figure 1: Paradigm shift from flat semantic matching to structured, intent-aware long-video RAG. **(Top)** Semantic-based retrieval relies on explicit semantic overlap, often missing cues that are only implicitly relevant to the query intent. **(Middle)** Intent-aware retrieval utilizes MLLM’s reasoning capability to identify cues that may be relevant to the query intent (detailed in Sec. 3.4); however, it can be distracted by noisy cues when operating over flattened contexts scattered across long temporal ranges. **(Bottom)** VideoStir structures the video as a spatio-temporal graph to retrieve cues from coherent contexts, enabling more reliable evidence aggregation. We further discuss the benefits of structured retrieval in Fig. 3.

Unlike short clips that typically encapsulate a single event, long videos consist of interleaved narratives where critical cues could be sparsely scattered across the timeline. To handle the extensive contexts, state-of-the-art (SOTA) models often resort to expanding context windows to support uniform sampling at second level (Comanici et al., 2025). However, such methods face a trade-off: they could either miss fine-grained details due to temporally sparse sampling or struggle with accurate reasoning when overwhelmed by redundant visual noise.

To address the contextual intractability of pro-

cessing long videos, Retrieval-Augmented Generation (RAG) has emerged as a promising solution by retrieving a small subset of segments to fit within limited context windows (Ren et al., 2025b). However, current long-video RAG methods typically flatten a video into a set of independent segments for query-based retrieval (Tang et al., 2025; Zhu et al., 2025), and rank candidates primarily by embedding similarity computed with contrastive vision–language models (Luo et al., 2025; Cao et al., 2025). Yet this flattened semantic matching paradigm leaves two critical gaps that hinder downstream reasoning.

The first gap is **spatio-temporal structure decoupling**. Flattening a video into independent segments disrupts its intrinsic spatio-temporal structure, decoupling contexts that should remain linked. Because text queries may lack the fine-grained cues needed to reconnect dispersed evidence, direct matching often struggles to retrieve contexts that are spatio-temporally relevant yet share little explicit semantic overlap with the query, limiting downstream models’ ability to form coherent reasoning chains (Shen et al., 2025a; Liao et al., 2024). The second gap is **insufficient intent modeling**. Similarity computed from contrastive embeddings is optimized for vision–language alignment and tends to capture “what looks similar” rather than “why it matters” for addressing the query’s reasoning intent (Pan et al., 2024). As shown in Figure 1, the query “For what *purpose* does the recorder use the *printer*?” leads a semantic retriever to select frames showing the “*printer*” due to high semantic overlap, whereas the actual “*purpose*” is revealed in a different event (e.g., making a poster), which semantic overlap alone fails to surface.

In light of these gaps, we advocate two shifts in long-video RAG paradigms: (i) *From Flattened to Structured*: moving beyond isolated segment retrieval to reconstruct the intrinsic spatio-temporal topology of the video; and (ii) *From Semantic to Intent*: going beyond surface semantic to model the alignment between query’s intent and visual cues.

To realize these shifts, we propose **VideoStir**, a framework designed to advance long-video RAG from flattened semantic matching to **Structured, Intent-aware Reasoning**. Inspired by human episodic memory where recall proceeds coarse-to-fine (locating relevant episodes before examining details), VideoStir operates in two phases, from clips to frames. First, a coarse clip-retrieval phase addresses structure decoupling by represent-

ing the video as a spatio-temporal graph, whose nodes are semantically coherent short clips segmented by an event boundary detector. Temporal edges preserve narrative continuity, while spatial edges bridge distant yet semantically related clips via inter-clip similarity. Traversing this topology enables multi-hop retrieval from query-matched anchors to their spatio-temporal neighbors, leveraging intrinsic correlations among semantically dense clips to aggregate dependencies and mitigate the semantic sparsity of direct query matching. Second, a fine-grained frame-retrieval phase addresses insufficient intent modeling with an **Intent-Relevance** scorer trained on our curated **IR-600K** dataset. Rather than considering semantic similarity alone, this lightweight MLLM-based scorer estimates frame–intent alignment, ensuring retrieved frames are not only semantically relevant but also pivotal for downstream reasoning.

Consequently, VideoStir provides the downstream MLLM with spatio-temporally context-coherent and intent-aligned visual cues, enabling more accurate long-video understanding. The scorer and curated IR-600K dataset further provide a reusable foundation for future research on intent-oriented long-video RAG systems. Our contributions are summarized as follows:

- We propose VideoStir, a novel long-video RAG framework that leverages structured retrieval to reconstruct video’s spatio-temporal topology and aggregate contextually coherent evidence, which overcomes the spatio-temporal decoupling of flattened retrieval.
- We introduce an intent-relevance scorer and IR-600K, a large-scale dataset modeling the relevance between frames and query intent. This shifts retrieval from semantic matching to intent awareness, establishing a foundation for intent-oriented long-video RAG.
- Extensive experiments show that VideoStir is competitive with SOTA baselines without relying on auxiliary information, highlighting the promise of the paradigm shift from flattened semantic matching to structured and intent-aligned reasoning in long-video RAG.

2 Related Works

2.1 Video-Language Models

Long-video understanding has attracted increasing attention in multimodal intelligence. Recent

general-purpose MLLMs such as Gemini (Anil et al., 2023) and Long-VITA (Shen et al., 2025b), along with specialized video models including Video-XL (Shu et al., 2025) and LongVLM (Weng et al., 2024), extend visual context to enable hour-scale video reasoning via uniform sampling. However, this paradigm faces a major bottleneck: limited context windows necessitate sparse sampling across the timeline, which can yield semantically redundant frames while simultaneously risking the loss of fleeting, query-critical cues. In parallel, contrastive video–language models (Luo et al., 2021) (e.g., Video-CLIP (Xu et al., 2021), X-CLIP (Ma et al., 2022), and PE (Bolya et al., 2025)) extend CLIP-style contrastive embeddings to video, enabling query-conditioned retrieval of salient clips and frames. Nevertheless, contrastive objectives primarily optimize for semantic similarity and may miss cues that are implicitly relevant to the query intent but lack an explicit semantic match.

2.2 Agentic Long-Video Understanding

To alleviate the contextual intractability inherent in processing long-duration videos, a growing body of recent work has shifted towards leveraging external systems. These systems are designed to reconstruct and distill extensive long-range context into compact, semantically rich cues tailored for downstream MLLMs. A prominent trajectory within this domain involves the exploration of agentic frameworks (Lin et al., 2023; Zhang et al., 2024b), which employ (M)LLMs as decision-making policies equipped with advanced modules for captioning, multi-step reasoning, and external memory (Jin et al., 2025). By simulating human-like curation processes (Wang et al., 2024b; Min et al., 2024), these agents dynamically filter irrelevant data. For example, DrVideo (Ma et al., 2025) synthesizes disjoint video frames into coherent long-form textual narratives to facilitate precise query answering, while Vgent (Shen et al., 2025a) abstracts complex entity interactions into structured symbolic representations to bolster reasoning. However, despite their demonstrated efficacy in enhancing understanding, deploying such sophisticated MLLM-based policies incurs substantial inference overhead.

2.3 RAG for Long-Video Understanding

To reduce the procedural inference overhead of agentic workflows, another line of work explores long-video retrieval-augmented generation (RAG)

frameworks (Ren et al., 2025b; Xue et al., 2025; Zeng et al., 2025), which externally select cues to fit limited context windows (Jeong et al., 2025; Cao et al., 2025; Zhu et al., 2025). For example, Video-RAG (Luo et al., 2025) retrieves keyframes via semantic similarity, and enriches entity information with external tools like optical character recognition and object detectors. E-VRAG (Xu et al., 2025) introduces a retrieval-reranking pipeline, leveraging Qwen3-Reranker-style binary judgments for fine-grained selection, while AKS (Tang et al., 2025) selects keyframes by jointly optimizing query-frame semantic similarity and temporal uniformity. However, existing methods often treat videos as independent segments, decoupling intrinsic spatio-temporal structure and limiting coherent reasoning. Moreover, current paradigms favor explicit semantic matching over intent-level alignment, risking the omission of cues implicitly relevant to the query intent.

3 Method

3.1 Overview

We propose VideoStir, a long-video RAG framework that explicitly models video spatio-temporal topology and query intent to retrieve pivotal evidence, thereby advancing downstream MLLMs’ long-video understanding.

As illustrated in Figure 2, VideoStir comprises three phases: (1) Spatio-temporal topology modeling. It converts the input video into a structured graph. An event boundary detector partitions the video into clip-level nodes connected by temporal edges (chronological adjacency) and spatial edges (inter-clip semantic similarity). This design reconstructs the topology of the video’s spatio-temporal context. (2) Graph-based clip retrieval. It first identifies anchor clips semantically aligned with the query, then expands the context via multi-hop traversal on the spatio-temporal graph. Consequently, the retrieved evidence includes not only query-matched moments but also their spatio-temporally relevant neighbors that are critical for reasoning. (3) Intent-aware frame retrieval. It refines evidence at the frame level. Since semantic similarity often fails to capture the query’s underlying intent, we introduce an intent-relevance scorer. Specifically, we leverage a lightweight MLLM fine-tuned on our curated IR-600K dataset to assign each frame a relevance score, computed as the weighted expectation over predicted relevance lev-

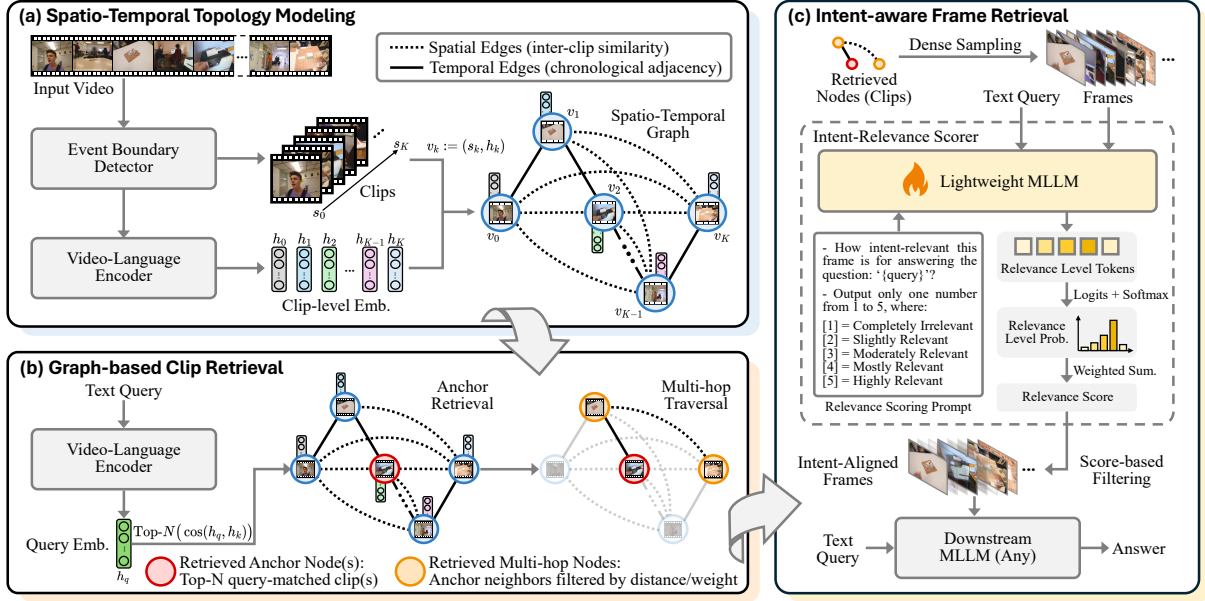


Figure 2: **Overview of VideoStir.** VideoStir achieves long-video understanding via spatio-temporal structuring and two-stage retrieval. **(a) Spatio-Temporal Topology Modeling.** An event boundary detector segments input video into semantically coherent clips. Their embeddings are used to construct a spatio-temporal graph with temporal edges linking adjacent clips and spatial edges weighted by inter-clip similarity. **(b) Graph-based Clip Retrieval.** Input query is embedded in the video-language space to retrieve top- N matched anchor nodes; multi-hop traversal then expands along spatio-temporal links under hop-distance and edge-weight constraints to aggregate contextual clips. **(c) Intent-aware Frame Retrieval.** Frames densely sampled from the retrieved clips are assessed by an intent-relevance scorer (LoRA-tuned on IR-600K), which outputs a probability distribution over discrete relevance levels. These distributions are then aggregated via probability-weighted expectation to produce a continuous relevance score; high-scoring, intent-aligned frames are retained and fed to a downstream MLLM to produce the final answer.

els. It mitigates the loss of implicit cues inherent to explicit semantic matching, yielding an intent-aligned evidence set for downstream MLLMs.

3.2 Spatio-Temporal Topology Modeling

To structurally encode a long video’s spatio-temporal topology, we model it as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each node $v_k \in \mathcal{V}$ corresponds to a clip $s_k \subset S$. These clips are obtained via an event boundary detector, which preprocesses the long video into semantically coherent segments by identifying event transitions through semantic change-point detection. This process is implemented by applying PELT (Killick et al., 2012) to the frame embeddings for adaptive segmentation. For each clip, we compute an embedding $h_k = E_{v1}(s_k)$ with a video-language encoder E_{v1} . The edge set $\mathcal{E} = \{\mathcal{E}_{\text{Temp}}, \mathcal{E}_{\text{Sp}}\}$ consists of temporal edges $\mathcal{E}_{\text{Temp}} = \{(v_k, v_{k+1}) \mid w_{k,k+1} = 1\}$ and spatial edges $\mathcal{E}_{\text{Sp}} = \{(v_i, v_j) \mid w_{i,j} = \cos(h_i, h_j)\}$, where w denotes the edge weight. Here, we refer to the embedding space as the “spatial” dimension of the topology. Temporal edges connect chronologically adjacent clips, preserving the video’s temporal backbone. Spatial edges are weighted by the similarity between clip embeddings, captur-

ing inter-clip correlations when different moments share similar video-level features. By jointly modeling these edges, \mathcal{G} induces a spatio-temporal topology over clips: temporally adjacent events are linked along timeline, while semantically coherent yet temporally distant events are bridged via spatial edges. This graph re-entangles contextual relations that may be lost in a flattened representation and provides the backbone for multi-hop retrieval.

3.3 Graph-based Clip Retrieval

Given a query q and a spatio-temporal graph \mathcal{G} , we aim to coarsely retrieve clips that are semantically relevant to q and topologically supported by inter-clip correlations. This design addresses semantically sparse query-clip matching: a query may refer to only a small facet of an event (e.g., the main subject), whereas long videos often contain recurring scenes/actions/entities and temporally adjacent evidence that constitute critical spatio-temporal contexts for the underlying event, which can be missed by simply computing semantic similarity between the query and individual moments. As shown in Figure 3, with our spatio-temporal graph, these contexts can be recovered via inter-clip similarity and chronological adjacency encoded by

spatio-temporal edges.

Specifically, we first compute the query embedding $h_q = E_{vl}(q)$ using the same video-language encoder as for clips, placing queries and clip nodes in a shared embedding space. We then measure query-node similarity via $\cos(h_q, h_k)$, and select the top- N similar nodes as the anchor set \mathcal{V}_{anc} . While these anchors are matched to the query, they may not fully cover their implicitly contextual dependencies. To recover broader context, we expand \mathcal{V}_{anc} along spatio-temporal links in \mathcal{G} , traversing edges according to their weights and collecting nodes within L hops of the anchors:

$$\mathcal{V}_{\text{hop}} = \{v_j \mid d(v_j, \mathcal{V}_{\text{anc}}) \leq L, w_{ij} \geq \eta\}, \quad (1)$$

where $d(v_j, \mathcal{V}_{\text{anc}})$ is the shortest hop distance from v_j to any anchor in \mathcal{V}_{anc} , and threshold η filters out weak connections. In this way, multi-hop traversal yields a coarse spatio-temporal neighborhood around the event hinted by the query, providing contextual candidates for subsequent fine-grained intent-level filtering.

3.4 Intent-aware Frame Retrieval

Intent Relevance Reasoning and Retrieval. While multi-hop retrieval yields a set of contextually related clips, the candidates may include redundant or merely co-occurring frames that are semantically relevant but misaligned with the query’s underlying intent. We therefore introduce an intent-relevance scorer R_θ , which filters retrieved clips at frame level using the intent-reasoning capability of a lightweight MLLM backbone π_θ . Specifically, R_θ estimates frame-query intent relevance, separating frames that genuinely support the query’s intent from those with only superficial similarity.

Given the retrieved node set $\mathcal{V}_{\text{hop}} = \{v_k\}$, where each node represents a video clip, we expand them into a dense frame pool: $\mathcal{F}_{\text{ret}} = \bigcup_{v_k \in \mathcal{V}_{\text{hop}}} \{x_t \mid x_t \in v_k\}$. For a query-frame pair (q, x_t) , the scorer outputs a softmax-normalized probability distribution over discrete intent-relevance level tokens $\ell \in \{1, 2, 3, 4, 5\}$:

$$P_\theta(\ell \mid q, x_t, \mathcal{P}_{\text{intent}}) = \frac{\exp(\pi_\theta(\ell \mid q, x_t, \mathcal{P}_{\text{intent}}))}{\sum_{k=1}^5 \exp(\pi_\theta(k \mid q, x_t, \mathcal{P}_{\text{intent}}))}, \quad (2)$$

where $\pi_\theta(\cdot)$ denotes the logits produced by the MLLM backbone of R_θ , and $\mathcal{P}_{\text{intent}}$ (detailed in Appendix E) denotes the prompt that guides the MLLM backbone π_θ to reason and score intent relevance. The score of frame x_t is computed as the

weighted expectation over relevance level tokens:

$$r_t = R_\theta(q, x_t, \mathcal{P}_{\text{intent}}) = \sum_{\ell=1}^5 \ell \cdot P_\theta(\ell \mid q, x_t, \mathcal{P}_{\text{intent}}). \quad (3)$$

This converts the model’s internal reasoning into a continuous relevance score, distinguishing frames that provide intent-level support for the query from those that are merely semantically similar. After obtaining scores for \mathcal{F}_{ret} , we retain frames exceeding a relevance threshold κ_s . The resulting set $\mathcal{F}_{\text{intent}} = \{x_t \mid r_t > \kappa_s\}$ provides an intent-aligned visual context for downstream reasoning.

Distillation-based Scorer Training. To strengthen the scorer’s intent-relevance assessment, we adopt a distillation strategy to fine-tune its MLLM backbone. Specifically, given a query-frame pair (q, x_t) and the intent-relevance scoring prompt $\mathcal{P}_{\text{intent}}$, a powerful teacher MLLM \mathcal{T} is asked to rate the intent relevance between q and x_t on the discrete level (1-5). The resulting pseudo-label $y_t = \mathcal{T}(q, x_t, \mathcal{P}_{\text{intent}})$ serves as supervision for the student scorer R_θ . During fine-tuning, we inject low-rank adaptation (LoRA) parameters (Hu et al., 2022) into the MLLM backbone of R_θ , enabling efficient adaptation without substantially disrupting its pretrained priors. Given the student prediction $P_\theta(\ell \mid q, x_t, \mathcal{P}_{\text{intent}})$, the training objective minimizes the cross-entropy between the student-predicted distribution and the teacher-provided discrete label:

$$\mathcal{L}_{\text{CE}} = - \sum_{\ell=1}^5 \mathbf{1}[\ell = y_t] \log P_\theta(\ell \mid q, x_t, \mathcal{P}_{\text{intent}}). \quad (4)$$

The LoRA parameters are optimized via gradient descent, enabling the student to inherit the teacher’s relevance-assessment patterns and yield more accurate relevance scores, while retaining the computational efficiency of its lightweight backbone.

4 Experiment

4.1 Dataset and Implementation

Dataset. We curate IR-600K, a dataset designed to enhance small-scale MLLMs’ ability to judge intent relevance between video frames and queries, supporting intent-aware long-video RAG. To the best of our knowledge, IR-600K is the first dataset targeting intent-level frame-query alignment, i.e., whether a frame provides evidence that satisfies the query’s underlying information need. This differs from prior datasets such as TVSum (Song et al., 2015) and QVHighlights (Lei et al., 2021), which

Table 1: Comparison with long-video RAG baselines. ‘‘Native Input’’ indicates whether the model processes video directly without relying on external auxiliary text or tools. Bold and underlined numbers indicate the best and second-best accuracies, respectively. ‘-’ indicates not applicable.

Model	Native Input	LV-Bench (Val)		MLVU		Video-MME-Long (w/o Sub)	
		Overall \uparrow	Gain (%) \uparrow	Overall \uparrow	Gain (%) \uparrow	Overall \uparrow	Gain (%) \uparrow
LLaVA-Video (7B)	✓	56.6	-	70.8	-	-	-
+ Video-RAG	✗	58.7	3.7	72.4	2.3	-	-
+ TV-RAG	✗	58.8	2.1	72.1	1.8	-	-
+ VideoStir (Ours)	✓	60.3	6.5	73.1	3.2	-	-
mPLUG-Owl3 (8B)	✓	52.1	-	63.7	-	50.1	-
+ Video-RAG	✗	54.5	4.6	64.4	1.1	50.6	1.0
+ TV-RAG	✗	54.8	5.2	64.0	0.5	50.9	1.6
+ VideoStir (Ours)	✓	55.1	5.8	64.6	1.4	51.0	1.8
Aria (25B)	✓	64.2	-	70.6	-	58.8	-
+ Video-RAG	✗	66.4	3.4	72.1	2.1	59.6	1.4
+ TV-RAG	✗	65.5	2.0	71.2	0.8	59.0	0.3
+ VideoStir (Ours)	✓	<u>65.8</u>	<u>2.5</u>	<u>71.7</u>	<u>1.6</u>	<u>59.4</u>	<u>1.0</u>
InternVL-1.5 (26B)	✓	51.2	-	50.4	-	45.6	-
+ Video-RAG	✗	52.2	2.0	51.8	2.8	46.7	2.4
+ TV-RAG	✗	<u>52.4</u>	<u>2.3</u>	51.0	1.2	46.3	1.5
+ VideoStir (Ours)	✓	52.7	2.9	<u>51.6</u>	<u>2.4</u>	46.9	2.9
LLaVA-Video (72B)	✓	61.9	-	73.1	-	61.5	-
+ Video-RAG	✗	65.4	5.7	73.8	1.0	62.3	1.3
+ TV-RAG	✗	64.6	4.3	73.4	0.4	61.8	0.5
+ VideoStir (Ours)	✓	66.0	6.6	74.1	1.4	<u>62.1</u>	<u>1.0</u>

Table 2: System-level comparison on EgoSchema’s test set with SOTA methods built upon fixed (M)LLMs. These approaches can be categorized into two groups: one employs video-to-text tools (e.g., captioners or object detectors) to provide auxiliary textual inputs for downstream (M)LLMs, while the other leverages the native multimodal understanding capability (Native Input) of MLLMs without auxiliary text. Bold and underlined numbers indicate the best and second-best accuracies, respectively. ‘-’ indicates not applicable.

(M)LLM Backbone	Method	Overall \uparrow
<i>Methods with Video-to-Text Tools</i>		
GPT-3.5	LLoVi (Zhang et al., 2024a)	57.6
	DrVideo (Ma et al., 2025)	62.6
GPT-4	VideoAgent (Wang et al., 2024c)	60.2
	LLoVi (Zhang et al., 2024a)	61.2
	VidAgent (Fan et al., 2024)	62.8
	VideoTree (Wang et al., 2025)	66.2
	DrVideo (Ma et al., 2025)	66.4
Qwen2-VL-7B	Vgent (Shen et al., 2025a)	68.0
<i>Methods with Native Input</i>		
Qwen2-VL-7B	IG-VLM (Kim et al., 2024)	66.2
	VideoStir (Ours)	<u>67.2</u>

primarily evaluate the semantic saliency of frames with respect to descriptive statements.

Specifically, we sample 4,678 videos from three representative video QA datasets (ActivityNet-QA (Yu et al., 2019), NExT-QA (Xiao et al., 2021), and STAR (Wu et al., 2024a)). For each video, we uniformly extract frames at 3 fps to form query-frame pairs. We then distill supervision from a powerful MLLM: Qwen2.5-VL-72B-Instruct (Bai et al., 2025) serves as the teacher to produce frame-level intent-relevance labels. Guided by a carefully

designed system prompt specifying explicit intent-relevance criteria, the teacher assesses intent-level alignment between each frame and the query, yielding intent-relevance scores. IR-600K comprises 605,676 training samples and 21,791 validation samples. Further details on the intent-relevance scoring prompt, dataset statistics and annotation format are provided in Appendix D and E.

Implementation. We adopt Qwen2.5-VL-3B-Instruct (Bai et al., 2025) as the backbone of the intent-relevance scorer and fine-tune it with LoRA adapters with a rank of 16, a scaling factor of 32, and a dropout of 0.05. We optimize the LoRA using AdamW with a learning rate of 5×10^{-5} and a weight decay of 0.05, together with a cosine learning-rate schedule and a warm-up ratio of 0.05. Training is performed for 1 epoch with a batch size of 128 under bf16 precision. For the event boundary detector, we use the vision tower from the Qwen2.5-VL series as the embedding backbone. For spatio-temporal graph construction and multi-hop retrieval, we use the Perception Encoder (Bolya et al., 2025) as the video-language encoder E_{vl} and set $N = 3$, $L = 2$, $\eta = 0.4$, and $\kappa_s = 3.25$; an ablation study on these retrieval hyperparameters is provided in Appendix C. All training runs and comparative experiments are conducted on $8 \times A100$ GPUs, while the remaining experiments are conducted on a single A100 GPU. Unless otherwise specified, results are

reported as the median of three trials, and ablations are conducted on LLaVA-Video (7B) using LongVideoBench (Wu et al., 2024b).

4.2 Comparison Against Other Methods

Table 1 compares VideoStir with the native MLLM and state-of-the-art (SOTA) long-video RAG baselines on LongVideoBench (LV-Bench) (Wu et al., 2024b), MLVU (Zhou et al., 2025), and Video-MME-Long (Fu et al., 2025a), following the same experimental setup as Video-RAG (Luo et al., 2025) and TV-RAG (Cao et al., 2025). Baseline MLLMs include GPT-4o (Hurst et al., 2024), LLaVA-Video (7B/72B) (Zhang et al., 2024c), mPLUG-Owl3 (8B) (Ye et al., 2024), Aria (25B) (Li et al., 2024a), and InternVL-1.5 (26B) (Chen et al., 2024). Results are compiled from the benchmarks’ leaderboards and the original papers, with the remaining results obtained from our reproductions based on their open-sourced codes and original settings. In most cases, VideoStir outperforms both the native model and prior long-video RAG baselines, demonstrating highly competitive performance. Notably, VideoStir preserves the native MLLM input setting by operating solely on retrieved visual evidence: it introduces no auxiliary text, tool calls, or extra modalities, while still achieving SOTA performance.

Table 2 compares VideoStir on EgoSchema against specialized long-video systems built atop fixed (M)LLMs (Mangalam et al., 2023); baseline results are taken from DrVideo (Ma et al., 2025) and Vgent (Shen et al., 2025a). Relying solely on native multimodal inputs without external video-to-text tools (e.g., captioners or object detectors), VideoStir achieves performance comparable to these specialized systems that leverage such auxiliary tools.

4.3 Ablation Study

Component-level ablations. Table 3 validates the contributions of VideoStir’s key components. Replacing the intent–relevance scorer with similarity-based frame retrieval results in substantial drops in both overall performance and retrieval accuracy, even when using strong vision–language embeddings from a powerful encoder such as PE. This indicates that pure semantic similarity matching can overlook the critical evidence required for long-video understanding, underscoring the importance of intent-aware retrieval. Meanwhile, removing the probability-weighted expectation (i.e., directly us-

Table 3: Component-level ablations on the expanded 1k validation set. “Overall” denotes overall accuracy, and “Retrieval Acc.” measures the accuracy of successfully retrieving key cues. “w/o Intent-relevance Scorer” replaces the scorer with similarity-based frame retrieval; “w/o Weighted Expectation” replaces probability-weighted expectation with discrete scores; “w/o Spatio-Temporal Graph” disables structured retrieval, which bypasses the graph and directly selects the top-14 clips by PE-embedding similarity, following the best-performing setting in Table 6. Bold indicates the best results.

Method	Overall↑	Retrieval Acc.↑
w/o Intent-relevance Scorer		
+ CLIP L/14	52.8	69.2
+ Video-CLIP	56.3	75.4
+ PE	58.1	79.8
w/o Weighted Expectation	54.2	71.6
w/o Event-Boundary Detector	62.3	88.2
w/o Spatio-Temporal Graph	56.4	74.8
w/o Spatial Edges	57.2	79.3
w/o Temporal Edges	59.8	83.4
Full	64.5	92.2

Table 4: Ablations of VideoStir’s video-language embedding backbones with $N = 1000$. “Overall” denotes overall accuracy, and “Retrieval Acc.” refers to the accuracy of successfully retrieving the key evidence for answering the query. The bold numbers represent the best results.

Embedding Backbone	Overall↑	Retrieval Acc.↑
Video-CLIP (Xu et al., 2021)	62.1	88.3
X-CLIP (Ma et al., 2022)	62.8	89.4
PE (Bolya et al., 2025)	64.5	92.2

ing discrete relevance levels) significantly degrades performance, suggesting that distributional scoring provides a smoother and better-calibrated ranking signal for discriminating frames. Moreover, disabling the spatio-temporal graph (i.e., flattening cues instead of performing structured retrieval), or removing spatial/temporal edges, markedly reduces retrieval accuracy, confirming that multi-hop aggregation over structured contexts is essential for extending evidence beyond the semantically matched anchor clips. In addition, replacing the event-boundary detector with uniform sampling that produces the same number of clips slightly degrades performance, suggesting that event-boundary segmentation isn’t the main performance driver, but it still improves evidence retrieval accuracy.

Embedding backbones. Table 4 examines the impact of embedding choices on graph-based clip retrieval. Stronger video–language encoders yield consistent improvements on downstream performance. However, these gains are comparatively modest relative to the contributions from the framework’s core components, suggesting that VideoStir’s effectiveness is driven primarily by its struc-

Table 5: Ablations on the intent-relevant scorer’s distillation strategies. “IR-600K CE” and “LV-Bench CE” denote the cross-entropy loss computed on the corresponding datasets. To reduce distributional bias, the student model is distilled using teacher models from the same family. “Retrieval Acc.” refers to the accuracy of successfully retrieving the key evidence.

Base Model	Method	IR-600K CE ↓	LV-Bench CE ↓	Overall ↑	Retrieval Acc. ↑	Trainable Params ↓
<i>Teacher Models</i>						
InternVL2.5-78B	Zero-Shot	-	-	64.8	92.6	-
Qwen2.5-VL-72B	Zero-Shot	-	-	65.4	95.8	-
Qwen2-VL-72B	Zero-Shot	-	-	64.9	94.7	-
<i>Student Models</i>						
InternVL2.5-4B	Zero-Shot	8.8755	8.9423	59.3	82.1	-
	LoRA	3.4614	3.8151	62.3	87.8	3.7M
	Full Params	3.0238	3.2578	63.4	90.7	3.0B
Qwen2.5-VL-3B	Zero-Shot	8.0909	8.1742	61.2	85.8	-
	LoRA	3.3101	3.6461	64.5	92.2	3.7M
	Full Params	2.9127	3.1017	64.7	93.2	3.0B
Qwen2-VL-2B	Zero-Shot	9.3210	9.3872	58.7	80.4	-
	LoRA	3.8127	4.3093	61.3	86.8	4.4M
	Full Params	3.1846	3.5204	62.4	88.9	1.5B

tured design rather than by reliance on increasingly powerful embedding models.

Scorer training strategy. Table 5 compares zero-shot inference, LoRA tuning, and full-parameter fine-tuning for the intent-relevance scorer. We evaluate different student-teacher pairs, including InternVL2.5 (Chen et al., 2024), Qwen2.5-VL (Bai et al., 2025), and Qwen2-VL (Wang et al., 2024a). For Qwen2-VL-2B, we increase the LoRA rank to 32 (vs. 16 for others) to align trainable parameter counts for fair comparison.

In the zero-shot setting, student’s performance significantly behind teachers. Conversely, LoRA tuning markedly reduces validation cross-entropy on IR-600K and consistently improves LV-Bench performance, demonstrating robust generalization. Notably, LoRA captures the majority of full-param tuning gains using only about 4M parameters, in contrast to the billions required for full updates, with the downstream performance of LoRA-tuned Qwen2.5-VL-3B nearly matching the fully fine-tuned model. Consequently, we adopt the Qwen2.5-VL family for both teacher and student roles to maximize performance. These results confirm that IR-600K provides effective, transferable supervision, while LoRA offers an optimal accuracy–efficiency trade-off for practical long-video RAG systems.

4.4 Discussion

Tables 1 and 2 present a comprehensive comparison between VideoStir and SOTA baselines, including those that utilize external expert modules (e.g., captioners, object detectors, and OCR models) to augment video with auxiliary text. Our results indicate that VideoStir achieves superior performance using

only video frames as native inputs, effectively bypassing the need for external textual augmentation. This suggests that one of the primary bottlenecks in long-video understanding still lies in the retrieval and organization of pivotal visual evidence. Therefore, refining the selection of visual evidence is essential for scaling MLLM’s long-video understanding capabilities. VideoStir addresses this by enhancing both the relevance and spatio-temporal coherence of retrieved evidence, thereby facilitating more reliable reasoning in downstream MLLMs.

5 Conclusion

We propose VideoStir, a novel long-video RAG framework designed to address the limitations of contextual structure decoupling and intent misalignment prevalent in current paradigms. VideoStir models long videos as spatio-temporal graphs and utilizes multi-hop traversal to recover the intrinsic topology of their contexts, thereby aggregating spatio-temporally coherent evidence. Additionally, we introduce an intent-relevance scorer trained on our newly curated IR-600K dataset, which further filters visual content by capturing the query’s underlying reasoning intent beyond surface-level semantics. Both the dataset and the scorer provide a solid foundation for future research into intent-oriented long-video RAG. Extensive experiments show that VideoStir is competitive with SOTA baselines without relying on auxiliary information, underscoring the effectiveness of shifting from flattened semantic matching to structured, intent-aware reasoning, and further revealing that one of the main bottlenecks in long-video understanding lies in retrieving and organizing pivotal visual evidence.

Limitations

As VideoStir introduces an additional step to transform long videos into structured representations, it inevitably incurs system latency. Indeed, system latency remains a broad challenge for long-video understanding frameworks that rely on external components like agents or RAG pipelines. We regard further reducing end-to-end latency as a pivotal direction for future long-video RAG research.

Acknowledgments

Research reported in this publication was supported by an Amazon Research Award, Fall 2025. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of Amazon.

The work is partially supported by the U.S. National Science Foundation (NSF) Grant CRII 2451683, an NVIDIA Academic Grants Program, a U.S. Bank Academic Research Award, the University of California, Merced, and a UC Merced Faculty Research Award. The views and conclusions are those of the authors and do not necessarily reflect the official policy or position of the U.S. Government.

References

- Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Daniel Bolya, Po-Yao Huang, Peize Sun, Jang Hyun Cho, Andrea Madotto, Chen Wei, Tengyu Ma, Jiale Zhi, Jathushan Rajasegaran, Hanoona Rasheed, and 1 others. 2025. Perception encoder: The best visual embeddings are not at the output of the network. *arXiv preprint arXiv:2504.13181*.
- Zongsheng Cao, Yangfan He, Anran Liu, Jun Xie, Feng Chen, and Zhepeng Wang. 2025. Tv-rag: A temporal-aware and semantic entropy-weighted framework for long video retrieval and understanding. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 9071–9079.
- Jianing Chen, Zehao Li, Yujun Cai, Hao Jiang, Shuqin Gao, Honglong Zhao, Tianlu Mao, and Yucheng Zhang. 2025. From tokens to nodes: Semantic-guided motion control for dynamic 3d gaussian splatting. *arXiv preprint arXiv:2510.02732*.
- Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, and 1 others. 2024. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *Science China Information Sciences*, 67(12):220101.
- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Yue Fan, Xiaojian Ma, Rujie Wu, Yuntao Du, Jiaqi Li, Zhi Gao, and Qing Li. 2024. Videoagent: A memory-augmented multimodal agent for video understanding. In *European Conference on Computer Vision*, pages 75–92. Springer.
- Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, and 1 others. 2025a. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24108–24118.
- Honghao Fu, Yuan Ouyang, Kai-Wei Chang, Yiwei Wang, Zi Huang, and Yujun Cai. 2025b. Contextnav: Towards agentic multimodal in-context learning. *arXiv preprint arXiv:2510.04560*.
- Honghao Fu, Junlong Ren, Qi Chai, Deheng Ye, Yujun Cai, and Hao Wang. 2025c. Vistawise: Building cost-effective agent with cross-modal knowledge graph for minecraft. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 21895–21909.
- Honghao Fu, Hao Wang, Jing Jih Chin, and Zhiqi Shen. 2025d. Brainvis: Exploring the bridge between brain and visual signals via image reconstruction. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Honghao Fu, Yufei Wang, Wenhan Yang, Alex C Kot, and Bihan Wen. 2024. Dp-iqa: Utilizing diffusion prior for blind image quality assessment in the wild. *arXiv preprint arXiv:2405.19996*.
- Haonan Ge, Yiwei Wang, Kai-Wei Chang, Hang Wu, and Yujun Cai. 2025a. Framemind: Frame-interleaved video reasoning via reinforcement learning. *arXiv preprint arXiv:2509.24008*.
- Haonan Ge, Yiwei Wang, Ming-Hsuan Yang, and Yujun Cai. 2025b. Mrfd: Multi-region fusion decoding with self-consistency for mitigating hallucinations in vlms. *arXiv preprint arXiv:2508.10264*.

- Yuyao Ge, Shenghua Liu, Baolong Bi, Yiwei Wang, Lingrui Mei, Wenjie Feng, Lizhe Chen, and Xueqi Cheng. 2024. Can graph descriptive order affect solving graph problems with llms? *ACL 2025*.
- Yuyao Ge, Shenghua Liu, Yiwei Wang, Lingrui Mei, Baolong Bi, Xuanshan Zhou, Jiayu Yao, Jiafeng Guo, and Xueqi Cheng. 2025c. Focusing by contrastive attention: Enhancing vlms’ visual reasoning. *arXiv preprint arXiv:2509.06461*.
- Yuyao Ge, Shenghua Liu, Yiwei Wang, Lingrui Mei, Lizhe Chen, Baolong Bi, and Xueqi Cheng. 2025d. Innate reasoning is not enough: In-context learning enhances reasoning large language models with less overthinking. *arXiv preprint arXiv:2503.19602*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Soyeong Jeong, Kangsan Kim, Jinheon Baek, and Sung Ju Hwang. 2025. Videorag: Retrieval-augmented generation over video corpus. *arXiv preprint arXiv:2501.05874*.
- Hongbo Jin, Qingyuan Wang, Wenhao Zhang, Yang Liu, and Sijie Cheng. 2025. Videomem: Enhancing ultra-long video understanding via adaptive memory management. *arXiv preprint arXiv:2512.04540*.
- Rebecca Killick, Paul Fearnhead, and Idris A Eckley. 2012. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500):1590–1598.
- Wonkyun Kim, Changin Choi, Wonseok Lee, and Wonjong Rhee. 2024. An image grid can be worth a video: Zero-shot video question answering using a vlm. *IEEE Access*.
- Jie Lei, Tamara L Berg, and Mohit Bansal. 2021. Detecting moments and highlights in videos via natural language queries. *Advances in Neural Information Processing Systems*, 34:11846–11858.
- Dongxu Li, Yudong Liu, Haoning Wu, Yue Wang, Zhiqi Shen, Bowen Qu, Xinyao Niu, Fan Zhou, Chengen Huang, Yanpeng Li, and 1 others. 2024a. Aria: An open multimodal native mixture-of-experts model. *arXiv preprint arXiv:2410.05993*.
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhao Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2025a. Videochat: Chat-centric video understanding. *Science China Information Sciences*, 68(10):200102.
- Zhecheng Li, Guoxian Song, Yujun Cai, Zhen Xiong, Junsong Yuan, and Yiwei Wang. 2025b. Texture or semantics? vision-language models get lost in font recognition. In *Conference on Language Modeling COLM, 2025*.
- Zhecheng Li, Yiwei Wang, Bryan Hooi, Yujun Cai, Naifan Cheung, Nanyun Peng, and Kai-Wei Chang. 2024b. Think carefully and check again! meta-generation unlocking llms for low-resource cross-lingual summarization. *arXiv preprint arXiv:2410.20021*.
- Zhecheng Li, Yiwei Wang, Bryan Hooi, Yujun Cai, Nanyun Peng, and Kai-Wei Chang. 2024c. Drs: Deep question reformulation with structured output. In *Association for Computational Linguistics ACL, 2025*.
- Zhecheng Li, Yiwei Wang, Bryan Hooi, Yujun Cai, Zhen Xiong, Nanyun Peng, and Kai-Wei Chang. 2024d. Vulnerability of llms to vertically aligned text manipulations. In *Association for Computational Linguistics ACL, 2025*.
- Ruotong Liao, Max Erler, Huiyu Wang, Guangyao Zhai, Gengyuan Zhang, Yunpu Ma, and Volker Tresp. 2024. Videoinsta: Zero-shot long video understanding via informative spatial-temporal reasoning with llms. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 6577–6602.
- Kevin Lin, Faisal Ahmed, Linjie Li, Chung-Ching Lin, Ehsan Azarnasab, Zhengyuan Yang, Jianfeng Wang, Lin Liang, Zicheng Liu, Yumao Lu, and 1 others. 2023. Mm-vid: Advancing video understanding with gpt-4v (ision). *arXiv preprint arXiv:2310.19773*.
- Chang Liu, Hongkai Chen, Yujun Cai, Hang Wu, Qingwen Ye, Ming-Hsuan Yang, and Yiwei Wang. 2025a. Structured attention matters to multimodal llms in document understanding. *arXiv preprint arXiv:2506.21600*.
- Xin Liu, Weijia Zhang, and Min-Ling Zhang. 2025b. Hacsurv: A hierarchical copula-based approach for survival analysis with dependent competing risks. In *International Conference on Artificial Intelligence and Statistics*, pages 3079–3087. PMLR.
- Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. 2021. Clip4clip: An empirical study of clip for end to end video clip retrieval. *arXiv preprint arXiv:2104.08860*.
- Yongdong Luo, Xiawu Zheng, Guilin Li, Shukang Yin, Haojia Lin, Chaoyou Fu, Jinfeng Huang, Jiayi Ji, Fei Chao, Jiebo Luo, and 1 others. 2025. Video-rag: Visually-aligned retrieval-augmented long video comprehension. *NeurIPS 2025*.
- Yiwei Ma, Guohai Xu, Xiaoshuai Sun, Ming Yan, Ji Zhang, and Rongrong Ji. 2022. X-clip: End-to-end multi-grained contrastive learning for video-text retrieval. In *Proceedings of the 30th ACM international conference on multimedia*, pages 638–647.

- Ziyu Ma, Chenhui Gou, Hengcan Shi, Bin Sun, Shutao Li, Hamid Rezaatofighi, and Jianfei Cai. 2025. Drvideo: Document retrieval based long video understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 18936–18946.
- Kartikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. 2023. Egoschema: A diagnostic benchmark for very long-form video language understanding. *Advances in Neural Information Processing Systems*, 36:46212–46244.
- Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Yuyao Ge, Jun Wan, Yurong Wu, and Xueqi Cheng. 2025. a1: Steep test-time scaling law via environment augmented generation. *arXiv preprint arXiv:2504.14597*.
- Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Jiayi Mao, and Xueqi Cheng. 2024a. "not aligned" is not "malicious": Being careful about hallucinations of large language models' jailbreak. *COLING 2025*.
- Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Ruibin Yuan, and Xueqi Cheng. 2024b. Hid-danguard: Fine-grained safe generation with specialized representation router. *arXiv preprint arXiv:2410.02684*.
- Lingrui Mei, Shenghua Liu, Yiwei Wang, Yuyao Ge, Baolong Bi, Jiayu Yao, Jun Wan, Ziling Yin, Jiafeng Guo, and Xueqi Cheng. 2026. Gated differentiable working memory for long-context language modeling. *arXiv preprint arXiv:2601.12906*.
- Juhong Min, Shyamal Buch, Arsha Nagrani, Minsu Cho, and Cordelia Schmid. 2024. Morevqa: Exploring modular reasoning models for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13235–13245.
- Kaihang Pan, Juncheng Li, Wenjie Wang, Hao Fei, Hongye Song, Wei Ji, Jun Lin, Xiaozhong Liu, Tat-Seng Chua, and Siliang Tang. 2024. I3: Introspective retrieval conditioned on instructions. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1839–1849.
- Junlong Ren, Gangjian Zhang, Honghao Fu, Pengcheng Wu, and Hao Wang. 2025a. Wamo: Wavelet-enhanced multi-frequency trajectory analysis for fine-grained text-motion retrieval. *arXiv preprint arXiv:2508.03343*.
- Xubin Ren, Lingrui Xu, Long Xia, Shuaiqiang Wang, Dawei Yin, and Chao Huang. 2025b. Videorag: Retrieval-augmented generation with extreme long-context videos. *arXiv preprint arXiv:2502.01549*.
- Xiaoqian Shen, Wenxuan Zhang, Jun Chen, and Mohamed Elhoseiny. 2025a. Vgent: Graph-based retrieval-reasoning-augmented generation for long video understanding. *arXiv preprint arXiv:2510.14032*.
- Yunhang Shen, Chaoyou Fu, Shaoqi Dong, Xiong Wang, Yi-Fan Zhang, Peixian Chen, Mengdan Zhang, Haoyu Cao, Ke Li, Shaohui Lin, and 1 others. 2025b. Long-vita: Scaling large multi-modal models to 1 million tokens with leading short-context accuracy. *arXiv preprint arXiv:2502.05177*.
- Yan Shu, Zheng Liu, Peitian Zhang, Minghao Qin, Junjie Zhou, Zhengyang Liang, Tiejun Huang, and Bo Zhao. 2025. Video-xl: Extra-long vision language model for hour-scale video understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26160–26169.
- Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. 2015. Tvsum: Summarizing web videos using titles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5179–5187.
- Xi Tang, Jihao Qiu, Lingxi Xie, Yunjie Tian, Jianbin Jiao, and Qixiang Ye. 2025. Adaptive keyframe sampling for long video understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 29118–29128.
- Xinpeng Ti, Wentao Ye, Zhifang Zhang, Junbo Zhao, Chang Yao, Lei Feng, and Haobo Wang. 2025. Towards reverse engineering of language models: A survey. In *Findings of the Association for Computational Linguistics: EMNLP 2025*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, and 1 others. 2024a. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Shijie Wang, Qi Zhao, Minh Quan Do, Nakul Agarwal, Kwonjoon Lee, and Chen Sun. 2024b. Vamos: Versatile action models for video understanding. In *European Conference on Computer Vision*, pages 142–160. Springer.
- Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. 2024c. Videoagent: Long-form video understanding with large language model as agent. In *European Conference on Computer Vision*, pages 58–76. Springer.
- Ziyang Wang, Shoubin Yu, Elias Stengel-Eskin, Jaehong Yoon, Feng Cheng, Gedas Bertasius, and Mohit Bansal. 2025. Videotree: Adaptive tree-based video representation for llm reasoning on long videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3272–3283.
- Yuetian Weng, Mingfei Han, Haoyu He, Xiaojun Chang, and Bohan Zhuang. 2024. Longvlm: Efficient long video understanding via large language models. In *European Conference on Computer Vision*, pages 453–470. Springer.
- Bo Wu, Shoubin Yu, Zhenfang Chen, Joshua B Tenenbaum, and Chuang Gan. 2024a. Star: A benchmark

- for situated reasoning in real-world videos. *arXiv preprint arXiv:2405.09711*.
- Hang Wu, Yujun Cai, Haonan Ge, Hongkai Chen, Ming-Hsuan Yang, and Yiwei Wang. 2025a. Re-fineshot: Rethinking cinematography understanding with foundational skill evaluation. *arXiv preprint arXiv:2510.02423*.
- Hang Wu, Yujun Cai, Zehao Li, Haonan Ge, Bowen Sun, Junsong Yuan, and Yiwei Wang. 2026. Cam-reasoner: Reinforcing camera movement understanding via structured spatial reasoning. *arXiv preprint arXiv:2602.00181*.
- Hang Wu, Hongkai Chen, Yujun Cai, Chang Liu, Qingwen Ye, Ming-Hsuan Yang, and Yiwei Wang. 2025b. Dimo-gui: Advancing test-time scaling in gui grounding via modality-aware visual reasoning. In *EMNLP 2025*.
- Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. 2024b. Longvideobench: A benchmark for long-context interleaved video-language understanding. *Advances in Neural Information Processing Systems*, 37:28828–28857.
- Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. 2021. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9777–9786.
- Zhen Xiong, Yujun Cai, Zhecheng Li, and Yiwei Wang. 2025a. Unveiling the potential of diffusion large language model in controllable generation. *arXiv preprint arXiv:2507.04504*.
- Zhen Xiong, Yujun Cai, Zhecheng Li, Junsong Yuan, and Yiwei Wang. 2025b. Thinking with sound: Audio chain-of-thought enables multimodal reasoning in large audio-language models. *arXiv preprint arXiv:2509.21749*.
- Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. 2021. Video-clip: Contrastive pre-training for zero-shot video-text understanding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6787–6800.
- Zeyu Xu, Junkang Zhang, Qiang Wang, and Yi Liu. 2025. E-vrag: Enhancing long video understanding with resource-efficient retrieval augmented generation. *arXiv preprint arXiv:2508.01546*.
- Zhucun Xue, Jiangning Zhang, Xurong Xie, Yuxuan Cai, Yong Liu, Xiangtai Li, and Dacheng Tao. 2025. Omni-adavideorag: Omni-contextual adaptive retrieval-augmented for efficient long video understanding. *arXiv preprint arXiv:2506.13589*.
- Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. 2024. mplug-owl3: Towards long image-sequence understanding in multi-modal large language models. *arXiv preprint arXiv:2408.04840*.
- Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. 2019. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9127–9134.
- Nianbo Zeng, Haowen Hou, Fei Richard Yu, Si Shi, and Ying Tiffany He. 2025. Scenerag: Scene-level retrieval-augmented generation for video understanding. *arXiv preprint arXiv:2506.07600*.
- Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, and Gedas Bertasius. 2024a. A simple llm framework for long-range video question-answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21715–21737.
- Lu Zhang, Tiancheng Zhao, Heting Ying, Yibo Ma, and Kyusong Lee. 2024b. Omagent: A multi-modal agent framework for complex video understanding with task divide-and-conquer. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10031–10045.
- Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2024c. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*.
- Zhifang Zhang, Shuo He, Haobo Wang, Bingquan Shen, and Lei Feng. 2025a. Defending multimodal backdoored models by repulsive visual prompt tuning. *NeurIPS*.
- Zhifang Zhang, Yuwei Niu, Xin Liu, and Beibei Li. 2025b. Tuning vision-language models with candidate labels by prompt alignment. In *DASFAA*.
- Zhifang Zhang, Qiqi Tao, Jiaqi Lv, Na Zhao, Lei Feng, and Joey Tianyi Zhou. 2025c. Tokenswap: Backdoor attack on the compositional understanding of large vision-language models. *arXiv preprint arXiv:2509.24566*.
- Zhifang Zhang, Bojun Yang, Shuo He, Weitong Chen, Wei Emma Zhang, Olaf Maennel, Lei Feng, , and Miao Xu. 2026. Test-time attention purification for backdoored large vision language models. In *CVPR*.
- Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Zhengyang Liang, Shitao Xiao, Minghao Qin, Xi Yang, Yongping Xiong, Bo Zhang, and 1 others. 2025. Mlvu: Benchmarking multi-task long video understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13691–13701.
- Zirui Zhu, Hailun Xu, Yang Luo, Yong Liu, Kanchan Sarkar, Zhenheng Yang, and Yang You. 2025. Focus: Efficient keyframe selection for long video understanding. *arXiv preprint arXiv:2510.27280*.

A The Use of LLMs

The paper employs LLMs for language polishing.

B Extended discussion for Structured Retrieval

In Figure 3, we further illustrate and discuss the advantages of structured retrieval via visualizations.

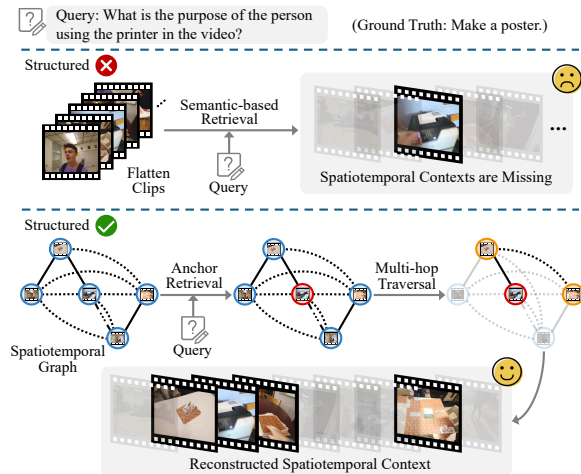


Figure 3: Advantages of structured retrieval over flat clip retrieval. Flattened retrieval mainly returns clips with explicit semantic overlap with the query, while overlooking clips that are contextually relevant but do not contain direct query-matching content. VideoStir structures long videos to reconstruct spatiotemporal context, allowing multi-hop traversal to aggregate contextually related clips around key moments. The resulting evidence is cleaner and more coherent, providing stronger support for subsequent fine-grained retrieval.

C Hyperparameter Ablations

We determine the hyperparameters for the graph-based clip retrieval procedure based on the ablation study in Table 6.

Table 6: Hyperparameter ablations on Graph-based Clip Retrieval. We investigate the impact of the number of anchors N , hop distance L , and spatial edge threshold η . ‘‘Avg. Clips’’ denotes the average number of retrieved clips per query. The bold numbers represent the best results.

Hyperparameter	Value	Avg. Clips	Retrieval Acc. \uparrow
Number of Anchors (N) (Fixed $L = 2, \eta = 0.6$)	1	7.2	88.0
	2	13.6	92.0
	3	21.9	90.5
	4	25.4	90.0
Hop Distance (L) (Fixed $N = 2, \eta = 0.6$)	1	8.4	87.0
	2	13.6	92.0
	3	16.6	92.0
	4	18.1	91.0
Edge Weight Threshold (η) (Fixed $N = 2, L = 2$)	0.2	31.4	84.0
	0.4	20.6	90.5
	0.6	13.6	92.0
	0.8	6.8	82.5

D Supplemental Details for IR-600K

D.1 Annotation Example

We provide an example annotation below; during training, the text portion of the user content will be substituted into the intent-relevance scoring prompt as {query}.

Annotation Example

```
[{"role": "system", "content": "You are a helpful assistant."}, {"role": "user", "content": [{"type": "image", "image": "8480330992/frame_000310.jpg"}, {"type": "text", "text": "why do the ladies tap the children s heads with some soft item"}]}, {"role": "assistant", "content": [{"type": "text", "text": "2"}]}
```

D.2 Dataset Distribution Statistics

We summarize the data distributions of the IR-600K training and validation sets in Figures 4.




Figure 4: Distribution of Intent Relevance Scores across the IR-600K training and validation sets. The results demonstrate that the validation set maintains a consistent class distribution with the training set across all relevance levels.

E Intent Relevance Scoring Prompt

Intent-Relevance Scoring Prompt

Infer the query’s intent and evaluate how likely it is that this frame is intent-relevant for answering: ‘{query}’.

Output only one number from 1 to 5, where:

[1] = completely irrelevant — the frame provides no visual or contextual information related to the question or its answer.

[2] = slightly relevant — the frame shows general background or context, but it is unlikely to contribute to answering.

[3] = moderately relevant — the frame includes partial clues or indirect context that might help infer the answer, but the key evidence is missing.

[4] = mostly relevant — the frame provides substantial visual or contextual information that can be used to answer the question, though not fully decisive.

[5] = highly relevant — the frame clearly contains the decisive evidence or strong contextual cues that directly or indirectly support the correct answer.