

MEMDREAMER: Decoupling Perception and Reasoning for Long Video Understanding via Agentic Hierarchical Graph Memory

Cong Chen^{1,2*}, Guo Gan^{3,*}, Kaixiang Ji^{2,*}, Zhaoyang Zhang^{2,4}, Zhen Yang⁵,
Guangming Yao², Hao Chen¹, Jingdong Chen², Yi Yuan², Chunhua Shen^{1,2,6†}

¹Zhejiang University, State Lab of CAD & CG, ²Ant Group, ³Zhejiang University

⁴Central South University, ⁵HKUST (GZ), ⁶Zhejiang University of Technology

Abstract

Current Vision-Language Models struggle with hours-long videos because processing full-length visual sequences induces prohibitive token explosion and attention dilution. To overcome this, we introduce **MEMDREAMER** to decouple perception and reasoning, shifting long-video understanding into an agentic exploration process. As a plug-and-play framework, it incrementally streams videos to construct a Hierarchical Graph Memory, a top-down three-tier architecture for semantic abstraction, anchored by a foundational graph capturing spatiotemporal and causal relations. During inference, the reasoning model employs agentic tool-augmented retrieval, navigating hierarchies, searching nodes, and traversing logical edges via an Observation-Reason-Action loop. Experiments show **MEMDREAMER** achieves SOTA results across four mainstream benchmarks, substantially narrowing the gap with human experts. It greatly reduces the reasoning context relative to full-context ingestion while delivering a substantial absolute accuracy gain. Furthermore, statistical analysis uncovers a strong positive linear correlation between an VLM’s performance on logic reasoning and long-video understanding benchmarks, establishing agentic capability scaling as a new paradigm for multimodal comprehension.

1 Introduction

Long video understanding (Jin et al., 2024; Song et al., 2024) represents a core capability for driving Vision-Language Models (VLMs) (Chen et al., 2024a; Google DeepMind, 2026; Hurst et al., 2024) toward embodied intelligence and open-world interaction. Despite breakthroughs in single image (Liu et al., 2023; Chen et al., 2024b; Lu et al., 2024) and short video analysis (Zhao et al., 2025), current VLMs (Li et al., 2024) struggle to scale to

*Equal contribution.

†Corresponding author.




Figure 1: **Decoupling perception unlocks reasoning ability for long video understanding.** We evaluate the same backbone in two regimes on LVBench (y-axis) against AIME 2025 (x-axis, reasoning ability): as an end-to-end VLM (gray) and as the reasoning model of **MEMDREAMER** over a pre-built hierarchical graph memory (red). End-to-end performance is insensitive to reasoning ability, as the model is overwhelmed by noisy long-context tokens and cannot extract task-relevant signal. In contrast, **MEMDREAMER** exhibits a strong linear trend, confirming that decoupling facilitates the direct transfer of the backbone’s agentic reasoning ability into long video understanding performance.

hours-long videos. This failure stems from the structural intricacy and extreme temporal redundancy inherent in long videos, which submerge critical reasoning signals beneath massive noise.

Current VLMs (Lin et al., 2024; Zhang et al., 2024b; Li et al., 2025) typically adopt a tightly coupled strategy, relying on ultra-long context windows to process visual perception and logical reasoning simultaneously. This coupled paradigm flattens long videos into massive token streams via brute-force frame sampling, introducing two bottlenecks. **In perception**, it incurs an intractable




Figure 2: **MEMDREAMER** organizes a long video as a three-tier coarse-to-fine hierarchy with cross-tier topological edges; at the leaf tier, each segment is a subgraph of entities and micro-events via attributive and causal edges.

token explosion: sampling a 2-hour video at 1 FPS generates over 1.6M tokens, drastically exceeding current context limits. In **Reasoning**, the influx of redundant tokens induces severe attention dilution (Xiao et al., 2024) and exacerbates the “lost in the middle” phenomenon (Liu et al., 2024), significantly degrading long-range reasoning capabilities.

To overcome these bottlenecks, we introduce **MEMDREAMER**, a novel paradigm that decouples perception from reasoning. Specifically, our framework employs a perception model (Google DeepMind, 2026; Bai et al., 2025) to process videos in a streaming fashion, incrementally constructing a persistent memory bank. During inference, a separate reasoning model executes a search loop within this memory to retrieve task-relevant cues. While this streaming-and-retrieval mechanism circumvents context limits and token redundancy, its efficacy depends on the memory’s structural organization. Existing flat (Yang et al., 2025) or chunk-based (Zhang et al., 2025a; Wang et al., 2025e) storage schemes obscure global perspectives and sever temporal-causal links, causing decoupled reasoning to degenerate into myopic, exhaustive retrieval.

To bridge this gap, **MEMDREAMER** introduces a novel Hierarchical Graph Memory. Video content is inherently hierarchical, mirroring how humans comprehend a long video coarse-to-fine, starting from the overall plot, then scene by scene, event by event, finally down to individual actions, rather than memorizing a flat, exhaustive sequence of moments (Pang and Wang, 2025; Long et al., 2025). Moreover, these fine-grained events and entities are

far from independent: they are intertwined by spatiotemporal and causal relations (Edge et al., 2024) that no purely sequential representation can faithfully capture. To realize this hierarchical and topological intuition, we construct a three-tier memory architecture shown in Figure 2. The top tier is a Video Root that summarizes global context, which is progressively decomposed into Super Events at the intermediate layer and Macro Events at the bottom. At the Macro Event tier, information is instantiated into a local subgraph characterizing entities, events, and their logical relations. This representation suppresses irrelevant detail at the appropriate granularity while preserving the long-range dependencies that downstream reasoning relies on.

To fully exploit this architecture, we propose a tool-augmented Agentic Retrieval Mechanism, superseding traditional full-context ingestion or similarity-based retrieval (Lewis et al., 2020; Gao et al., 2023). We construct a dedicated Agentic Tool Bank comprising three categories: Navigation for vertical hierarchical traversal, Search for rapid node localization, and Graph Traversal for tracking logical chains along topological edges. By engaging in an iterative Observation-Reason-Action loop (Yao et al., 2022), the model dynamically interacts with the memory. This design shifts long video understanding from passive token consumption to a multi-step active agentic exploration task.

Extensive experiments validate the superiority of **MEMDREAMER**. On challenging hours-level benchmarks (Wang et al., 2025a; Wu et al., 2024; Fu et al., 2025), **MEMDREAMER** consis-

tently achieves SOTA results (Mangalam et al., 2023), narrowing the performance gap with human experts to 8.9 points. Notably, our decoupled strategy achieves an absolute gain of 7.3 points over end-to-end coupled paradigms using only 2% of the context window on hours-long video bench LVBench (Wang et al., 2025a). Furthermore, MEMDREAMER reveals a critical observation: we empirically establish a strong positive correlation between an VLM’s intrinsic agentic capability and its long-video understanding performance, suggesting the scaling of agentic capacity as a new paradigm for future research.

In summary, our contributions are threefold:

1. MEMDREAMER decouples perception and reasoning via persistent memory and agentic retrieval, bypassing long-context bottlenecks. To our knowledge, we reveal for the first time a positive correlation between long-video performance and reasoning, suggesting a new optimization direction.
2. MEMDREAMER achieves SOTA performance across four representative long-video benchmarks. Using identical base models, we achieve substantial gains over traditional end-to-end baselines, narrowing the gap toward human expert levels.
3. MEMDREAMER establishes foundational design principles for multimodal memory systems. Empirically, hierarchical graphs and agentic tool-use outperform flat chunk-based storage and naive full-context ingestion or semantic matching, providing a cornerstone for future architectures.

2 Related Work

2.1 Long Video Understanding

Long video understanding remains a frontier for VLMs, marking the evolutionary leap to complex, real-world visual streams. To handle potentially infinite inputs, previous research has attempted to optimize visual encoders (Lin et al., 2024; Zhang et al., 2024b) and scale native context windows (Bai et al., 2025; Hurst et al., 2024; Team et al., 2023). However, even with leading models expanding context limits to 1M tokens like Gemini-2.5-Pro (Comanici et al., 2025), systems still rely on aggressive frame subsampling to accommodate lengthy videos. This brute-force paradigm leaves models vulnerable to “needle-in-a-haystack” retrieval failures (Wang et al., 2024a, 2025d) and the “lost in the middle” phenomenon (Liu et al., 2024) on benchmarks like LVBench (Wang et al., 2025a). To bypass these bottlenecks, a more scalable paradigm

structures video information within an external memory bank, transforming static perception into an agentic reasoning process.

2.2 Agentic Memory Systems

While memory systems like MemGPT (Packer et al., 2023) and MemoryBank (Zhong et al., 2024) achieve remarkable success in LLMs, adapting them for VLMs remains challenging. Early video attempts (Song et al., 2024) rely on simple frame buffers, and subsequently M3-Agent (Long et al., 2025) and WorldMM (Yeo et al., 2025) encapsulate semantics into discrete bins. Yet, these flat designs lack top-down taxonomy and topological edges. Recently, VideoARM (Yin et al., 2025) and MM-mem (Lian et al., 2026) pioneered hierarchical memory architectures. However, they maintain these hierarchies as edge-less perceptual buffers or rely on passive, uncertainty-driven vertical drill-down. MEMDREAMER bridges this gap by coupling a three-tier hierarchical abstraction with a causal graph topology, encoding multi-granular semantics and localized details to ground subsequent retrieval. CompassMem (Hu et al., 2026) similarly organizes experiences as an event graph with explicit logical relations and supports structured, goal-directed navigation for long-horizon reasoning. In contrast, MEMDREAMER targets raw video streams with a Video Root–Super Event–Macro Event hierarchy and local heterogeneous entity–event subgraphs, enabling coarse-to-fine retrieval of video evidence. Furthermore, to decouple perception and reasoning for long videos, we construct a purely textual memory. Unlike previous leading methods (Zhang et al., 2025a; Yin et al., 2025) that require additional perception models to revisit raw video frames during retrieval, our approach relies solely on interactions with text memory.

2.3 Agentic Retrieval Mechanism

Early information-seeking approaches (Luo et al., 2025; Jeong et al., 2025) rely on naive Retrieval-Augmented Generation (Lewis et al., 2020; Gao et al., 2023), using static dense embeddings to recall contexts. However, semantic similarity does not guarantee logical relevance. A visually similar clip may lack causal connection to the query, and these single-turn retrievals cannot dynamically self-correct. To address this, pioneering works like VideoAgent (Wang et al., 2024b) and DVD (Zhang et al., 2025a) introduced agentic retrieval, enabling VLMs to actively search. Yet, constrained by flat

clip databases, their exploration often degrades into blind trial-and-error. **MEMDREAMER** extends the agentic loop by coupling it with our graph memory. Through a tool-augmented Observation-Reason-Action loop, we unleash the agent’s reasoning capabilities to explore the memory across multiple dimensions and locate relevant information.

3 Method

We formulate long-video understanding as a decoupled paradigm comprising two stages: persistent memory construction and tool-augmented agentic retrieval. Given a video stream V , a perception model $\mathcal{P}$ first processes the visual inputs in a streaming fashion to construct a structured, purely textual Hierarchical Graph Memory, denoted as $\mathcal{G}$. Subsequently, upon receiving a text query Q , a reasoning model $\mathcal{R}$ equipped with a tool bank $\mathcal{T}$ actively explores $\mathcal{G}$ through a multi-step Observation-Reason-Action loop. During inference, $\mathcal{R}$ operates only on this pure text memory, never raw video information. This iterative exploration extracts a concise set of task-relevant clues C , yielding the final answer $A = \mathcal{R}(Q, C)$.

3.1 Hierarchical Graph Memory

MEMDREAMER conceptualizes a long video stream as a coarse-to-fine three-tier semantic topology comprising the *Video Root*, *Super Events* and *Macro Events* layers. Beneath the foundational *Macro Events* layer, we formulate a local subgraph to characterize entities, micro-events, and their spatiotemporal and causal dependencies. To construct this memory representation, the pipeline first adaptively **segments the streaming video** to yield semantically self-contained *Macro Events* as anchors. subsequently, the system **extracts fine-grained local subgraph** structures downward to preserve visual details, while **aggregating nodes upward** to distill macro-level hierarchical semantics.

Streaming Adaptive Segmentation. We introduce a streaming adaptive segmentation mechanism driven by semantic boundaries, rather than relying on rigid, fixed-length chunking (e.g., 30-second intervals) common in prior works (Zhang et al., 2025a). Specifically, at each iteration k , the system maintains a temporal sliding window $W_k = [t_{\text{start}}^k, t_{\text{start}}^k + \tau]$, where t_{start}^k denotes the initial timestamp of the current window and τ represents the maximum window duration. Within W_k , the perception model $\mathcal{P}$ executes temporal ground-

ing to resolve a localized set of complete *Macro Events*, denoted as $\{e_1, e_2, \dots, e_N\}$. The concluding boundary of this final event e_N directly serves as the starting point for the subsequent window. This paradigm yields two benefits: first, the adaptive semantic partitioning ensures that each *Macro Event* node remains semantically intact and self-contained, thereby avoiding the arbitrary truncation inherent in previous fixed-window strategies. Second, it caps the maximum video length of a single input to $\mathcal{P}$ within the horizon τ , bounding context window pressure during the perception phase.

Downward Subgraph Extraction. For each segmented *Macro Events* e_i , we construct a local spatiotemporal subgraph $g_i = (\mathcal{V}_i, \mathcal{E}_i)$ to capture fine-grained visual details. While conventional LLM-based graph memories (Edge et al., 2024; Guo et al., 2025b) encapsulate an entire event into a single static $\langle \text{Subject entity, Relation edge, Object entity} \rangle$ triplet, they struggle to model the highly dynamic temporal evolution and causal chains inherent in long videos. To address this limitation, we introduce a novel *Video-centric Graph Ontology* where the vertex set is formally defined as $\mathcal{V}_i = \mathcal{V}_i^E \cup \mathcal{V}_i^M$, incorporating micro-events ($\mathcal{V}_i^M$) as distinct nodes alongside standard entities ($\mathcal{V}_i^E$). These vertices are interconnected via a heterogeneous edge set $\mathcal{E}_i$ comprising three distinct topological relations: *spatial-attribute edges* ($\mathcal{V}_i^E \leftrightarrow \mathcal{V}_i^E$) delineate the relative spatial layouts and static physical attributes among entities, *subject-object edges* ($\mathcal{V}_i^E \leftrightarrow \mathcal{V}_i^M$) anchor discrete entities to the actions they initiate or undergo, and directed *temporal-causal edges* ($\mathcal{V}_i^M \rightarrow \mathcal{V}_i^M$) model the chronological progression and causal dependencies between micro-events. We use the perception model $\mathcal{P}$ to parse the video clip to extract this structured subgraph g_i alongside a textual summary of the *Macro Event*.

Upward Hierarchical Aggregation. Upon establishing the foundational graph layer, the system executes a bottom-up semantic abstraction to provide a global navigation backbone for subsequent retrieval. Specifically, the localized textual descriptions of all *Macro Events* are fed into the model $\mathcal{P}$ as the foundational leaf nodes of the three-tier topology. Guided by temporal adjacency and semantic affinity, $\mathcal{P}$ clusters and distills neighboring, highly correlated *Macro Events* into *Super Events* that span broader temporal scales, until ultimately converging at a single *Video Root*




Figure 3: **Architectural workflow of MEMDREAMER.** (Left) Memory construction comprises three phases: streaming adaptive segmentation, downward subgraph extraction and upward hierarchical aggregation. (Right) Three tool categories support an agentic retrieval mechanism driven by an Observation-Reason-Action loop.

node. This coarse-to-fine topology structurally spans from micro-level actions to the overall narrative, empowering the retrieval agent with efficient global navigation and precise target localization.

3.2 Agentic Tool-Augmented Retrieval

We introduce a tool-augmented agentic retrieval mechanism to exploit our structured, multi-granular graph memory. Equipping the reasoning model $\mathcal{R}$ with a graph-access toolbox across three functional dimensions and an Observation-Reason-Action loop, we transform static video understanding into a multi-step, proactive reasoning process.

Multi-Dimensional Tool Bank. To facilitate deep retrieval and reasoning over the hierarchical graph memory, we construct a graph-access toolbox encompassing three functional dimensions. First, the *hierarchical navigation tools*, consisting of *GetSummary*($\cdot$), *GetSuperEvent*($\cdot$), *GetMacroEvent*($\cdot$), and *GetSubgraph*($\cdot$), allow the model $\mathcal{R}$ to progressively unpack the three-tier topology from a top-down, coarse-to-fine perspective, traversing from global narratives down to super-events, macro-events and localized subgraphs at varying granularities. Second, the *precise search tools*, including *SearchNodes*($\cdot$) and *SearchByTime*($\cdot$), support vector-based semantic retrieval via text embeddings and chronological filtering within targeted temporal windows, enabling rapid localization of relevant nodes. After invoking *SearchNodes*($\cdot$) or *SearchByTime*($\cdot$), $\mathcal{R}$ autonomously determines whether the query requires access to the original video; if so, it prompts

$\mathcal{P}$ to inspect the retrieved video segment and extract query-relevant visual evidence. Third, the *local graph traversal tool* *GetRelationGraph*($\cdot$) extracts the topological neighborhood of a target node, allowing the model to execute multi-hop tracing along spatial, temporal, or causal edges to resolve complex cross-spatiotemporal logic chains. Collectively, these three tool categories formulate an expressive action space, providing the retrieval agent with the capability to shift its perspective between the macroscopic semantic hierarchy and the underlying dense graph details.

Agentic Observation-Reason-Action Loop.

Our retrieval mechanism operates within an Observation-Reason-Action loop driven by the model $\mathcal{R}$. At each step t , the model evaluates the initial user query Q alongside the historical execution trajectory $\mathcal{H}_{t-1}$ to determine the proper tool execution and corresponding parameters a_t :

$$a_t = \mathcal{R}(Q, \mathcal{H}_{t-1}) \quad (1)$$

Executing the action a_t invokes the designated tool from our toolkit to yield an environmental observation o_t from the graph memory $\mathcal{G}$. To prevent multi-turn interactions from accumulating irrelevant textual noise, which induces the “lost in the middle” phenomenon within long contexts, we eschew direct concatenation of raw observational outputs. Instead, the model $\mathcal{R}$ distills o_t to extract only the task-relevant evidence cues conditional on the query, denoted as

$$c_t = \mathcal{R}(o_t, Q). \quad (2)$$

Subsequently, the current action-clue pair is appended to the historical trajectory to transition the working memory state of the agent:

$$\mathcal{H}_t = \mathcal{H}_{t-1} \cup \{(a_t, c_t)\} \quad (3)$$

By effectively obtaining raw observations and aggregating only relevant clues, this mechanism guarantees that the reasoning model $\mathcal{R}$ exclusively attends to high-quality evidence. Through this entire Agentic Loop, the model $\mathcal{R}$ can fully explore the memory space to gather the essential evidence required for the final response.

4 Experiments

4.1 Experiment Setup

Evaluation Benchmarks. We use four challenging benchmarks: LVBench (Wang et al., 2025a) comprises 103 long videos of 30 mins to 2 hours and 1,549 QA pairs across six fine-grained dimensions. The LongVideoBench (Wu et al., 2024) validation set includes 753 videos and 1,337 questions, featuring 188 videos of 15 to 60 minutes. The long-video split of Video-MME (Fu et al., 2025) covers 300 videos (30 to 60 minutes) and 900 questions. EgoSchema (Mangalam et al., 2023) contains egocentric clips focusing on reasoning. Together, these benchmarks span diverse video durations, task granularities, and reasoning depths.

Implementation Details. In the memory construction phase, the continuous video stream is segmented via a $\tau = 10min$ sliding window, and Gemini-3.1-Pro serves as the perception model to construct the hierarchical graph memory using raw video streams without any external transcripts. In the online reasoning phase, we evaluate Gemini-2.5-Pro (Comanici et al., 2025), Gemini-3.1-Pro (Google DeepMind, 2026), and the open-source Qwen3-VL-235B-A22B-Thinking (Bai et al., 2025) as the reasoning engines. We adopt Qwen3-Embedding (Zhang et al., 2025b) for semantic vector computing and cap the agent’s maximum tool-call budget at 12 steps. All reported results are averaged across three runs.

Baselines. We benchmark MEMDREAMER against two categories of long-video systems: (1) *Vanilla Long Video VLMs*: we compare it directly against its underlying reasoning engines, covering top closed-source frontiers (Hurst et al., 2024; Google DeepMind, 2026; Guo et al., 2025a) and




Figure 4: MEMDREAMER only need a 5-6K reasoning window, 41-124x smaller than end-to-end input.

representative open-source (Bai et al., 2025; Team, 2026, 2025; Zeng et al., 2025; Chen et al., 2024a; Wang et al., 2025b,c). (2) *Memory-based Video LLMs*: we evaluate our framework against established long-form video memory baselines (Wang et al., 2025e, 2024b; Yang et al., 2025; Pang and Wang, 2025; Long et al., 2025; Yeo et al., 2025; Lian et al., 2026) alongside the latest SOTA memory-driven systems (Zhang et al., 2025a; Yin et al., 2025; Lin et al., 2026).

4.2 Main Results

SOTA Performance Across Diverse Benchmarks. Table 1 summarizes the performance of MEMDREAMER across four benchmarks. Results demonstrate that MEMDREAMER establishes new SOTA across all four benchmarks. Specifically, on LVBench, MEMDREAMER achieves a peak score of 85.5 with a 7.3-point improvement over the strongest native closed-source baseline Gemini-3.1-Pro, significantly narrowing the gap to human expert performance to 8.9 points. On LongVideoBench and Video-MME, our framework attains 87.1 and 86.1, yielding substantial improvements of 8.5 and 5.8 points. Furthermore, even on EgoSchema, which features egocentric perspectives, MEMDREAMER secures a top score of 87.1. These consistent improvements across both open-source and closed-source reasoning engines validate the generality and pronounced superiority of our plug-and-play architectural paradigm.

Mitigating Context Overload via Decoupled Reasoning. Table 2 and Figure 4 compare the context length of our decoupled paradigm against conventional native end-to-end full-context ingestion. The vanilla approach forces the vision-

Methods	Reason Model	LVBench							LongVideoBench		Video-MME	EgoSchema
		Avg.↑	ER↑	EU↑	Rea↑	KIR↑	Sum↑	TG↑	Avg.↑	Long↑	Long(w/o sub)↑	Val↑
End-to-End VLMs (Proprietary)												
GPT-4o	GPT-4o	48.9	48.9	49.5	50.3	48.1	50.0	40.9	66.7	60.9	65.3	70.4
Gemini-2.0-Flash	Gemini-2.0-Flash	48.6	47.4	48.5	44.4	56.8	41.4	39.3	—	45.7	63.0	71.2
Gemini-2.5-Pro	Gemini-2.5-Pro	72.0	71.5	71.1	67.7	80.0	63.5	69.1	71.0	68.6	75.9	72.8
Gemini-3.1-Pro	Gemini-3.1-Pro	78.2	78.1	76.7	74.1	86.3	70.7	81.4	78.6	77.0	80.3	76.4
OpenAI-o3	OpenAI-o3	57.1	57.6	56.4	50.8	62.9	67.2	46.8	66.7	60.6	64.7	63.2
Seed1.5VL	Seed1.5VL-Thinking	64.6	65.4	63.4	68.0	53.6	63.7	46.6	—	74.4	—	—
End-to-End VLMs (Open-Source)												
Qwen3-VL	Qwen3-VL-235B-A22B-Thinking	63.6	63.7	62.6	63.1	62.6	65.5	59.6	71.4	67.1	71.2	75.9
Qwen3.5	Qwen3.5-35A3B	71.3	72.2	69.4	74.4	68.9	67.5	66.2	—	—	62.6	—
GLM-4.6V	GLM-4.6V-106B-A12B	59.5	58.8	59.2	69.7	57.8	61.8	53.9	67.6	58.7	66.0	68.8
GLM-4.5V	GLM-4.5V-106B-A12B	53.4	55.3	53.1	57.5	55.3	39.6	49.9	66.4	54.1	64.7	68.4
InternVL2.5	InternVL2.5-78B	43.6	43.8	42.0	51.0	37.9	42.1	36.8	—	—	62.6	—
InternVL3.5	InternVL3.5-30BA3B	44.4	42.7	44.1	48.3	46.4	36.2	40.9	62.9	52.7	64.1	86.8
AdaRETAKE	AdaRETAKE	53.3	53.0	50.7	54.7	62.2	37.9	45.5	67.0	—	65.0	—
Agentic Video Understanding Methods												
VideoTree	VideoTree	28.8	30.3	25.1	31.9	26.5	25.5	27.7	—	—	—	67.0
VideoAgent	VideoAgent	29.3	28.0	30.3	28.0	28.0	36.4	29.3	—	—	—	63.2
VCA	VCA	41.3	43.7	40.7	46.2	37.8	27.3	38.0	—	—	—	73.6
MR-Video	MR-Video	60.8	59.8	57.4	57.7	71.4	50.0	58.8	—	61.6	61.8	73.0
M3-Agent	M3-Agent	49.3	—	—	—	—	—	—	—	—	61.8	—
WorldMM-GPT	GPT-5	61.9	—	—	—	—	—	—	—	—	76.6	—
MM-Mem	MM-Mem	—	—	—	—	—	—	—	—	—	66.1	—
DVD	OpenAI-o3	74.2	73.4	73.3	70.7	80.4	74.1	72.3	71.6	68.6	67.3	76.6
VideoARM	OpenAI-o3	79.7	—	—	—	—	—	—	78.0	76.4	81.2	76.2
VideoSeek	GPT-5	68.4	—	—	—	—	—	—	—	73.5	60.9	—
MEMDREAMER (Ours, Plug-and-Play Agentic Framework)												
MEMDREAMER	Qwen3-VL-235B-A22B-Thinking	78.9 ^(+15.3)	77.9	78.8	74.7	79.8	78.2	81.2	80.4 ^(+9.0)	77.4 ^(+10.3)	80.3 ^(+9.1)	80.6 ^(+4.7)
MEMDREAMER	Gemini-2.5-Pro	77.8 ^(+5.8)	74.7	81.0	78.5	79.6	88.2	83.5	75.9 ^(+4.9)	76.1 ^(+7.5)	82.6 ^(+6.7)	85.2 ^(+12.4)
MEMDREAMER	Gemini-3.1-Pro	85.5 ^(+7.3)	84.5	85.3	83.8	85.8	84.4	86.6	87.1 ^(+8.5)	85.0 ^(+8.0)	86.1 ^(+5.8)	87.1 ^(+10.7)
Human Expert	Human	94.4	—	—	—	—	—	—	—	—	—	—

Table 1: **Main results on four long-video understanding benchmarks.** MEMDREAMER achieves state-of-the-art across all four benchmarks, with substantial gains over end-to-end baselines (green: improvement over the strongest end-to-end baseline using the same backbone).

Role	Model	Context Windows	LVB
Vanilla Full Video			
End-to-End	Gemini-3.1-Pro	265K	78.2
	Gemini-2.5-Pro	784K	72.0
	Qwen3-VL-235B-A22B-Thinking	240K	63.6
MemDreamer (Ours)			
Perception	Gemini-3.1-Pro	40.3K	—
	Gemini-3.1-Pro	6.2K	85.5
Reasoning	Gemini-2.5-Pro	6.3K	77.8
	Qwen3-VL-235B-A22B-Thinking	5.9K	78.9

Table 2: Comparison of input context length requirements and overall performance on LVBench (LVB).

language model to ingest massive sequences ranging from 240K to 784K tokens, severely degrading reasoning accuracy due to spatial-attention dilution and the “lost in the middle” effect. Conversely, by segregating perception from reasoning, MEMDREAMER avoids stacking redundant visual tokens. During agentic reasoning, our tool-augmented navigation operates over a condensed topology, restricting the active context to 5.9K–6.3K tokens, delivering an approximate 40× reduction versus full-video ingestion. Reducing token noise and redundancy unlocks a 15.3-point leap for Qwen3-VL from 63.6 to 78.9, empirically showing that structured agentic reasoning is a superior alternative to brute-force token scaling.

Model	AIME2025	LVBench	
		End-to-end	MEMDREAMER
Claude-Opus-4.6	99.8	67.9	85.1
Gemini-3.1-Pro	98.1	78.2	85.5
Kimi-K2.5	96.1	75.9	79.2
GPT-5	94.6	60.1	84.4
Qwen3-VL-235B-A22B-Thinking	89.7	63.6	78.9
Gemini-2.5-Pro	88.0	72.0	77.8
Qwen3-VL-235B-A22B-Instruct	74.7	67.7	71.5
gpt-4o-2024-11-20	61.9	48.9	68.5
<i>Statistical Correlation with Agentic Capability</i>			
Pearson R ↑	—	0.702	0.948
p -value ↓	—	0.052	<0.001

Table 3: Correlation between Agentic reasoning capabilities (AIME 2025) and long video understanding.

4.3 Analysis

Paradigm Shift to Pure Agentic Reasoning. Table 3 uncovers the correlation between an LLM’s logical capacity and long-video understanding performance. We evaluate eight models (Team et al., 2026; Anthropic, 2026) on the widely used AIME2025 reasoning benchmark and compute the Pearson correlation (R) against their LVBench scores. Under the traditional end-to-end paradigm, the correlation is restricted to 0.702 with an insignificant p -value of 0.052, demonstrating that raw video ingestion creates perceptual barriers that hinder models from leveraging their reasoning capac-

Perception Model	Reasoning Model	LVBench
Gemini-2.5-Pro	Gemini-2.5-Pro	76.1
	Gemini-3.1-Pro	85.1
	Qwen3-VL-235B-A22B-Thinking	77.7
Gemini-3.1-Pro	Gemini-2.5-Pro	77.8
	Gemini-3.1-Pro	85.5
	Qwen3-VL-235B-A22B-Thinking	78.9

Table 4: Ablation study on different base models.

Memory Architecture	Components		LVBench
	Hierarchy	Graph	
Flat-Chunk	✗	✗	71.0
Flat-Graph	✗	✓	78.6
Hierarchical-Chunk	✓	✗	80.6
Hierarchical-Graph (Ours)	✓	✓	85.5

Table 5: Ablation study on the memory design.


ity. Conversely, integrating **MEMDREAMER** scales this correlation to **0.948** with statistical significance ($p < 0.001$). This establishes a paradigm shift: supported by structured graph memory and agentic retrieval mechanism, long-video understanding bypasses the context-window limits, scales with the model’s logical reasoning capability.

Robustness via Decoupled Base Models. Table 4 presents the ablation under varying combinations of perception and reasoning backbones. Because **MEMDREAMER** only requires the perception model to process localized short-form video clips (under 10 min), the capability gaps between different perception backbones are largely smoothed out. When driving the same downstream reasoning engine, either Gemini-3.1-Pro or Qwen3-VL, the final performance fluctuations caused by swapping the underlying perception model are merely 0.4 and 1.2 percentage points, respectively. This demonstrates that our decoupled framework exhibits high perceptual error tolerance, effectively alleviating the over-reliance on expensive, long-context perception for macroscopic video understanding.

Hierarchy and graph contribute complementary gains. Table 5 and Figure 5 isolates the independent and joint contributions of the hierarchical topology (*Hierarchy*) and cross-node relational edges (*Graph*). Removing both components *1D Flat-Chunk* yields the worst baseline performance of 71.0. Introducing topological edges alone *1D Flat-Graph* drives a prominent 7.6-point performance gain to 78.6, showing that temporal and causal graph connections effectively preserve event continuity. Similarly, incorporating the bottom-up

Retrieval Strategy	LVBench		
	Avg.	Sum.	Rea.
Vanilla Embedding Similarity	66.6	50.3	71.8
Full Memory Context	74.5	76.8	68.0
Agentic Search Only	75.5	76.2	71.6
Agentic Full Tools (Ours)	85.5	84.4	83.8

Table 6: Ablation study on memory retrieval strategies.



(a) Memory Construction (b) Retrieval Mechanism

Figure 5: Ablations of different memory designs and retrieval strategies across LVBench sub-categories.

abstraction layer alone *Hierarchical-Chunk* raises the score to 80.6, validating that a multi-granular macro-skeleton prevents the retrieval agent from getting trapped in localized trivial details. Our full design *Hierarchical-Graph* achieves a performance of 85.5. These gains demonstrate that the macroscopic hierarchy and the underlying dense graph optimize global path navigation and localized target grounding.

Tool-augmented retrieval beats embedding lookup. Table 6 and Figure 5 evaluates distinct memory access mechanisms to verify the necessity of our tool-augmented agentic retrieval. Our strategy *Agentic Full Tools* achieves a score of 85.5, outperforming the traditional static match baseline *Vanilla Embedding Similarity* by 18.9 points. This demonstrates that multi-step tool invocation captures significantly more precise cues than single-turn passive recall. While feeding the entire graph skeleton Full Memory Context benefits summarization with a score of 76.8, its performance on reasoning sharply degrades to 68.0. This confirms that passing the context window with unfiltered, graph structures introduces severe topological noise that disrupts critical logical deduction. Furthermore, restricting the agent to semantic retrieval alone Agentic Search Only yields a suboptimal accuracy of 75.5, demonstrating that the completeness of our multi-dimensional toolkit is critical for efficient

long-form video understanding.

5 Conclusion

We propose **MEMDREAMER**, a long-video understanding paradigm integrating a Hierarchical Graph Memory with tool-augmented agentic retrieval. Our framework establishes new SOTA across multiple challenging benchmarks. To our knowledge, we reveal for the first time a strong positive correlation between long-video performance and agentic reasoning capabilities. This paradigm bypasses context limits and mitigates attention dilution, offering a promising scaling direction for future multimodal comprehension.

Limitations

We discuss two boundaries of the present study.

Visual-only modality. MemDreamer constructs its hierarchical graph memory exclusively from the visual modality. Speech and ambient audio that often carry substantial semantic load in news broadcasts, lectures, interviews, and dialogue-driven films, are not incorporated into either the perception or the reasoning stage. Queries whose answers depend primarily on auditory cues therefore cannot be resolved by retrieval over the current memory. Extending the heterogeneous subgraph schema with an audio-derived node type, aligned with the existing entity and micro-event tiers, is a direction we leave to future work.

Offline video assumption. The term *streaming* in this paper refers to sliding-window perception during memory construction: the perception model $\mathcal{P}$ ingests one window of duration τ at a time rather than the full video at once. The framework nevertheless presupposes that the entire video is available before the agentic retrieval phase begins. This setting is consistent with the every long-video benchmark we use, all of which issue queries against pre-recorded, completed videos. Truly online settings, where queries arrive concurrently with an unbounded incoming stream, and the hierarchical graph must be updated, pruned, and re-summarized in real time, introduce additional challenges in latency budgeting, memory growth, and query, perception interleaving that fall outside the scope of this study. We regard real-time streaming long-video understanding as an important next step toward deployment in live broadcasting, surveillance, and embodied perception scenarios.

Both limitations delineate the operating regime of the current system rather than the decoupled paradigm itself, and we expect the hierarchical graph memory and the agentic retrieval mechanism to extend naturally along each axis.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62576315).

References

- Anthropic. 2026. [Claude Opus 4.6](#).
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, and 1 others. 2025. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, and 1 others. 2024a. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. 2024b. [Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks](#). *Preprint*, arXiv:2312.14238.
- 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*.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Chaoyou Fu, Yuhang Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, and 1 others. 2025. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24108–24118.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang,

- and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Google DeepMind. 2026. [Gemini 3.1 pro](#). Accessed: 2026-05-20.
- Dong Guo, Faming Wu, Feida Zhu, Fuxing Leng, Guang Shi, Haobin Chen, Haoqi Fan, Jian Wang, Jianyu Jiang, Jiawei Wang, and 1 others. 2025a. Seed1. 5-vl technical report. *arXiv preprint arXiv:2505.07062*.
- Zirui Guo, Lianghao Xia, Yanhua Yu, Tian Ao, and Chao Huang. 2025b. Lightrag: Simple and fast retrieval-augmented generation. In *EMNLP (Findings)*, pages 10746–10761.
- Yuyang Hu, Jiongnan Liu, Jiejun Tan, Yutao Zhu, and Zhicheng Dou. 2026. [Memory matters more: Event-centric memory as a logic map for agent searching and reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 22389–22407, San Diego, California, United States. Association for Computational Linguistics.
- 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](#). *Preprint*, arXiv:2501.05874.
- Peng Jin, Ryuichi Takanobu, Wancai Zhang, Xiaochun Cao, and Li Yuan. 2024. Chat-univi: Unified visual representation empowers large language models with image and video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13700–13710.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). *Preprint*, arXiv:2309.06180.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2025. Videochat: Chat-centric video understanding. *Science China Information Sciences*, 68(10):200102.
- Yanwei Li, Chengyao Wang, and Jiaya Jia. 2024. Llama-vid: An image is worth 2 tokens in large language models. In *European Conference on Computer Vision*, pages 323–340. Springer.
- Niu Lian, Yuting Wang, Hanshu Yao, Jinpeng Wang, Bin Chen, Yaowei Wang, Min Zhang, and Shu-Tao Xia. 2026. From verbatim to gist: Distilling pyramidal multimodal memory via semantic information bottleneck for long-horizon video agents. *arXiv preprint arXiv:2603.01455*.
- Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. 2024. Video-llava: Learning united visual representation by alignment before projection. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pages 5971–5984.
- Jingyang Lin, Jialian Wu, Jiang Liu, Ximeng Sun, Ze Wang, Xiaodong Yu, Jiebo Luo, Zicheng Liu, and Emad Barsoum. 2026. Videoseek: Long-horizon video agent with tool-guided seeking. *arXiv preprint arXiv:2603.20185*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paran-jape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the association for computational linguistics*, 12:157–173.
- Lin Long, Yichen He, Wentao Ye, Yiyuan Pan, Yuan Lin, Hang Li, Junbo Zhao, and Wei Li. 2025. Seeing, listening, remembering, and reasoning: A multimodal agent with long-term memory. *arXiv preprint arXiv:2508.09736*.
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Hao Yang, and 1 others. 2024. Deepseek-vl: towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525*.
- Yongdong Luo, Xiwu Zheng, Guilin Li, Shukang Yin, Haojia Lin, Chaoyou Fu, Jinfa Huang, Jiayi Ji, Fei Chao, Jiebo Luo, and 1 others. 2025. Video-rag: Visually-aligned retrieval-augmented long video comprehension. *Advances in Neural Information Processing Systems*, 38:168008–168033.
- Karttikeya 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.
- Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G Patil, Ion Stoica, and Joseph E Gonzalez. 2023. Memgpt: Towards llms as operating systems. *arXiv preprint arXiv:2310.08560*.

- Ziqi Pang and Yu-Xiong Wang. 2025. Mr. video: "mapreduce" is the principle for long video understanding. *arXiv preprint arXiv:2504.16082*.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2023. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*.
- Enxin Song, Wenhao Chai, Guan hong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, and 1 others. 2024. Moviechat: From dense token to sparse memory for long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18221–18232.
- Gemini Team, 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*.
- Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, SH Cai, Yuan Cao, Y Charles, HS Che, Cheng Chen, Guanduo Chen, and 1 others. 2026. Kimi k2. 5: Visual agentic intelligence. *arXiv preprint arXiv:2602.02276*.
- Qwen Team. 2026. **Qwen3.5: Towards native multimodal agents**. Accessed: 2026-02-16.
- Z.ai Team. 2025. **Glm-4.6v: Open source multimodal models with native tool use**. Accessed: 2026-02-16.
- Weihan Wang, Zehai He, Wen yi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Ming Ding, Xiaotao Gu, Shiyu Huang, Bin Xu, and 1 others. 2025a. Lvbench: An extreme long video understanding benchmark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22958–22967.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, and 1 others. 2025b. Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*.
- Weiyun Wang, Shuibo Zhang, Yiming Ren, Yuchen Duan, Tiantong Li, Shuo Liu, Mengkang Hu, Zhe Chen, Kaipeng Zhang, Lewei Lu, and 1 others. 2024a. Needle in a multimodal haystack. *Advances in Neural Information Processing Systems*, 37:20540–20565.
- Xiao Wang, Qingyi Si, Shiyu Zhu, Jianlong Wu, Li Cao, and Liqiang Nie. 2025c. Adaretake: Adaptive redundancy reduction to perceive longer for video-language understanding. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 5417–5432.
- Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. 2024b. Videoagent: Long-form video understanding with large language model as agent. In *European Conference on Computer Vision*, pages 58–76. Springer.
- Zhaowei Wang, Wenhao Yu, Xiyu Ren, Jipeng Zhang, Yu Zhao, Rohit Saxena, Liang Cheng, Ginny Wong, Simon See, Pasquale Minervini, and 1 others. 2025d. Mmlongbench: Benchmarking long-context vision-language models effectively and thoroughly. *Advances in Neural Information Processing Systems*, 38.
- Ziyang Wang, Shoubin Yu, Elias Stengel-Eskin, Jaehong Yoon, Feng Cheng, Gedas Bertasius, and Mohit Bansal. 2025e. 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.
- Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. 2024. Longvideobench: A benchmark for long-context interleaved video-language understanding. *Advances in Neural Information Processing Systems*, 37:28828–28857.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. Efficient streaming language models with attention sinks. In *International Conference on Learning Representations*, volume 2024, pages 21875–21895.
- Zeyuan Yang, Delin Chen, Xueyang Yu, Maohao Shen, and Chuang Gan. 2025. Vca: Video curious agent for long video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20168–20179.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.
- Woongyeon Yeo, Kangsan Kim, Jaehong Yoon, and Sung Ju Hwang. 2025. Worldmm: Dynamic multimodal memory agent for long video reasoning. *arXiv preprint arXiv:2512.02425*.
- Yufei Yin, Qianke Meng, Minghao Chen, Jiajun Ding, Zhenwei Shao, and Zhou Yu. 2025. Videoarm: Agentic reasoning over hierarchical memory for long-form video understanding. *arXiv preprint arXiv:2512.12360*.
- Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie, Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, and 1 others. 2025. Glm-4.5: Agentic, reasoning, and coding (arc) foundation models. *arXiv preprint arXiv:2508.06471*.
- Kaichen Zhang, Bo Li, Peiyuan Zhang, Fanyi Pu, Joshua Adrian Cahyono, Kairui Hu, Shuai Liu, Yuanhan Zhang, Jingkang Yang, Chunyuan Li, and Ziwei Liu. 2024a. **Lmms-eval: Reality check on the evaluation of large multimodal models**. *Preprint, arXiv:2407.12772*.

- Xiaoyi Zhang, Zhaoyang Jia, Zongyu Guo, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. 2025a. Deep video discovery: Agentic search with tool use for long-form video understanding. *Advances in Neural Information Processing Systems*, 38:89863–89895.
- Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, and 1 others. 2025b. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*.
- Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2024b. Llava-video: Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*.
- Yilun Zhao, Haowei Zhang, Lujing Xie, Tongyan Hu, Guo Gan, Yitao Long, Zhiyuan Hu, Weiyuan Chen, Chuhan Li, Zhijian Xu, and 1 others. 2025. Mmvu: Measuring expert-level multi-discipline video understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8475–8489.
- Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. 2024. Memorybank: Enhancing large language models with long-term memory. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 19724–19731.

A Implementation Details

A.1 Memory Construction Implementation

The persistent memory $\mathcal{G}$ is a three-tier hierarchy whose full schema is summarized in Table 7. We highlight the two design choices that define our memory: the heterogeneous subgraph at the leaf tier, and the principles behind upward aggregation.

Heterogeneous Subgraph. For each Regular Event e_i , the perception model $\mathcal{P}$ produces a local subgraph $g_i = (\mathcal{V}_i, \mathcal{E}_i)$ that exposes micro-events as nodes ($\mathcal{V}_i^M$) alongside entities ($\mathcal{V}_i^E$), rather than collapsing each interaction into a single (s, r, o) triplet as in conventional LLM graph memories. Flat triplets have no temporal extent, cannot express n -ary participation, and offer no anchor for causal reasoning between actions. Promoting micro-events to nodes lets us factor the subgraph along three orthogonal edge categories: *spatial-attribute* edges carry layout, composition, and on-screen attribute bindings, *subject-object* edges attach entities to the events they take part in, *temporal-causal* edges record chronology and causality among micro-events. This factorization is what enables multi-hop reasoning during agentic retrieval.

Upward Aggregation. The Macro Event, Super Event and Root tiers are built bottom-up by $\mathcal{P}$ rather than by clustering algorithms, since narrative coherence is poorly captured by surface signals such as entity-set Jaccard. Macro Events are partitioned into Super Events under three principles. *Temporal contiguity* restricts grouping to adjacent Macros along the timeline. *Entity consistency* requires grouped events to share the same core participants. *Goal consistency* ends the current Super whenever the narrative goal shifts. A final synthesis call then condenses the Super Events into the Video Root, giving the retrieval agent a coarse-to-fine backbone from the global narrative down to fine-grained dynamics.

A.2 Agentic Retrieval Implementation

Toolkit. The seven tools fall into hierarchical navigation, precise search, and graph traversal. Their full input and output specification is given in Table 8, and Algorithm 1 lists the corresponding Observation-Reason-Action loop. Two design choices warrant emphasis. First, before the loop starts, $\mathcal{R}$ is initialized with the Video Root summary and the summaries of all Super Events as

Algorithm 1 Agentic Retrieval loop

Require: Query Q ; memory $\mathcal{G}$; reasoning model $\mathcal{R}$; tool bank $\mathcal{T}$; round budget $T_{\max}$
Ensure: Answer A , or abstention $\perp$

```

1:  $\mathcal{H}_0 \leftarrow \emptyset$   $\triangleright$  trajectory of (action, clue) pairs
2: for  $t = 1$  to  $T_{\max}$  do
3:    $(d_t, a_t) \leftarrow \mathcal{R}(Q, \mathcal{H}_{t-1})$ 
      $\triangleright d_t \in \{\text{can\_answer, cannot\_answer, call}\}$ 
4:   if  $d_t = \text{can\_answer}$  then
5:     return  $A \leftarrow \mathcal{R}_{\text{ans}}(Q, \mathcal{H}_{t-1})$ 
6:   else if  $d_t = \text{cannot\_answer}$  then
7:     return  $\perp$ 
8:   end if
9:    $o_t \leftarrow \mathcal{T}[a_t.\text{tool}](a_t.\text{args}; \mathcal{G})$ 
      $\triangleright$  tool call over  $\mathcal{G}$  (Table 8)
10:   $c_t \leftarrow \text{EXTRACT}(o_t, Q; \mathcal{R})$   $\triangleright$  selective compression
11:   $\mathcal{H}_t \leftarrow \mathcal{H}_{t-1} \cup \{(a_t, c_t)\}$ 
12: end for
13: return  $A \leftarrow \mathcal{R}_{\text{ans}}(Q, \mathcal{H}_{T_{\max}})$   $\triangleright$  hard-cap fallback
```

a default observation, so the agent always begins with a global narrative anchor rather than from scratch. Second, SearchNodes returns each top- k hit together with its depth-1 ego-graph neighborhood, which gives $\mathcal{R}$ immediate access to the surrounding entities, micro-events, and edges of any retrieved node.

B Reproducibility

MEMDREAMER inference configuration. At retrieval time, $\mathcal{R}$ runs with a maximum output length of 16K tokens and temperature 0.7 across all backbones. We evaluate Gemini-3.1-Pro, Gemini-2.5-Pro, and Qwen3-VL under these defaults.

Closed-source baseline evaluation. Numbers in the main table are taken from each model’s technical report or prior benchmarks where reported. The remaining cells are produced in-house. Hour-scale videos in our four benchmarks exceed the multi-GB ceiling of direct API uploads. We therefore re-encode each video before sending it to the closed-source APIs. For Gemini-family models, we re-encode at FPS=1 and stage the artefact on Google Cloud Storage, which lifts the size restriction. For Claude-Opus-4.6, we uniformly sample 128 frames per video. For the GPT family, we uniformly sample 50 frames, which is the maximum the GPT API accepts. Default temperature for closed-source models is 0.7. For GPT-5, the reasoning_effort is set to high to activate extended thinking. For Claude-Opus-4.6, we set temperature to 1.0, which is the API-mandated value for activating thinking mode.

Symbol	Type	Definition
Hierarchy nodes (one per video)		
v^R	Video Root	Global root anchor: title, description, themes, key_entities, time_range.
v^S	Super Event	Narrative-phase aggregation: label, description, time_range, $\{v^{Mc}\}$ children, deduplicated key_entities.
v^{Mc}	Macro Event	Per-episode summary anchor: label, summary, time_range, child micro-event ids, top- k key_entities.
Subgraph nodes $\mathcal{V}_i = \mathcal{V}_i^E \cup \mathcal{V}_i^M$ (one subgraph per Macro Event)		
$\mathcal{V}_i^E$	Entity	Persistent participant: name, entity_type $\in \{\text{PERSON, OBJECT, LOCATION, GROUP}\}$, structured attributes, natural-language description.
$\mathcal{V}_i^M$	Micro-event	Atomic action with explicit temporal extent: time_range, subject, object, action, description.
Subgraph edges $\mathcal{E}_i$, three topological categories		
<i>Spatial-attribute</i>	$\mathcal{V}_i^E \leftrightarrow \mathcal{V}_i^E$	LOCATED_IN, NEXT_TO, ATTACHED_TO, PART_OF, IS_SAME_AS, ANNOTATES, static layout, composition, coreference, and on-screen attribute binding.
<i>Subject-object</i>	$\mathcal{V}_i^E \leftrightarrow \mathcal{V}_i^M$	PERFORMS, RECEIVES, USES_TOOL, HAPPENS_IN, CONTEXT_FOR, role assignment of entities to micro-events and contextual annotation of events.
<i>Temporal-causal</i>	$\mathcal{V}_i^M \rightarrow \mathcal{V}_i^M$	BEFORE, OVERLAP (auto-generated from disjoint/overlapping intervals), CAUSES, PREVENTS, SUBEVENT_OF (VLM-extracted), chronology and causality among micro-events.
Cross-tier edges $\mathcal{E}^H$		
<i>Hierarchical</i>	$v^{Mc} \rightarrow v^S, v^S \rightarrow v^R$	SUBEVENT_OF edges that wire the three-tier topology.
<i>Macro causal / progression</i>	$v^{Mc} \rightarrow v^{Mc}, v^S \rightarrow v^S$	Inter-episode reasoning: CAUSES, PREVENTS, ENABLES (intra-Super), LEADS_TO, RESOLVES, CONTRASTS_WITH (cross-Super). BEFORE edges between consecutive Super Events are auto-generated.

Table 7: Schema of the Hierarchical Graph Memory $\mathcal{G}$. The hierarchy spans three tiers (v^R , v^S , v^{Mc} , and the Macro Event leaves whose contents are unfolded into local subgraphs g_i). Each subgraph factorises a Macro Event into entities, micro-events, connected by three edge categories that decouple static layout, action role binding, and temporal-causal flow. Field names mirror the JSON schema produced by our pipeline.

Open-source baseline evaluation. We run open-source baselines with the lmms-eval (Zhang et al., 2024a) framework. For Qwen3-VL we follow the official recipe with max tokens 240k and the official prompt template, remaining open-source models use the framework’s default frame-sampling policy. All in-house reproduction runs use temperature 0.7 and a maximum output length of 8192 tokens. To illustrate the inference details, we report the computational costs of evaluating different models on the LVBench dataset. All tests are conducted using NVIDIA H800 GPUs: Qwen3VL-235B requires 29 hours on 32 GPUs, whereas InternVL takes 18 hours on 4 GPUs. With the vLLM framework (Kwon et al., 2023), Qwen3.5-35B (4 GPUs), GLM-4.5V (8 GPUs), and GLM-4.6V (8 GPUs) all complete the evaluation in approximately 3 hours.

Details of Token Calculation claimed in Introduction. In the Introduction, we claim that a 2-hour 720P video requires over 1.6M tokens. To

clarify, for a 2-hour video sampled at 1 FPS (7200 frames) with a standard resolution of 1280×720 pixels, we employ the Qwen3-VL visual tokenizer with a patch size of 32. Additionally, a spatial 2×2 token merging strategy is applied, which compresses every 4 adjacent tokens into 1. Consequently, the total number of visual tokens input to the LLM is calculated as: $N_{\text{tokens}} = 7200 \times \frac{1280 \times 720}{32 \times 32 \times 4} = 7200 \times 225 = 1,620,000$ (1.62M).

C Supplementary Analysis

This appendix collects supplementary experiments that probe specific aspects of our findings beyond the main-paper experiments.

C.1 Correlation with GPQA Diamond

We replicate the correlation study on **GPQA Diamond** (Rein et al., 2023), an independent benchmark targeting graduate-level scientific reasoning over textual knowledge rather than mathematical symbol manipulation. The pattern repro-

Tool	Parameters	Output	Purpose
Hierarchical Navigation			
GetSummary	–	Video-level summary: title, themes, and key_entities.	Retrieve the root-tier overview of the entire video.
GetSuperEvent	super_id	Super-event records: label, time_range, key_entities, and child-macro count.	Browse high-level narrative segments.
GetMacroEvent	super_id or macro_ids	Macro-event records: label, time_range, and key_entities.	Drill from a super-event into its constituent macro-events.
GetSubgraph	macro_id	Local subgraph of one macro-event: entities, micro-events, on-screen text, and key edges (causal / spatial / attributive).	Inspect fine-grained dynamics inside a single macro-event.
Precise Search			
SearchNodes	query, top_k, node_types	Top- k nodes ranked by embedding similarity, each annotated with its parent macro-event and one-hop ego-graph context.	Dense semantic retrieval over entities, micro-events, and on-screen text within a unified embedding space.
SearchByTime	start_sec, end_sec	Macro-events whose time intervals overlap $[t_s, t_e]$, paired with their parent super-event labels.	Time-anchored localization given a temporal range.
Graph Traversal			
GetRelationGraph	node_id	Outgoing and incoming edges of the target node, with relation types, labels, and rationales.	Multi-hop traversal along spatial-attribute, subject-object, and temporal-causal edges.

Table 8: Specification of the seven tools available to the reasoning model $\mathcal{R}$ during the Observation-Reason-Action loop. The toolkit is organized into three categories aligned with the agent’s retrieval intent: navigating the three-tier hierarchy, executing semantic or temporal search, and traversing local graph neighborhoods. All tools return formatted textual evidence consumable directly by $\mathcal{R}$.

Model	GPQA Diamond	LVBench	
		End-to-end	MEMDREAMER
Gemini-3.1-Pro	94.1	78.2	85.5
GPT-5	92.6	60.1	84.4
Kimi-2.5	87.6	75.9	79.2
Claude-Opus-4.6	87.5	67.9	85.1
Gemini-2.5-Pro	84.4	72.0	77.8
Qwen3VL-235A22B-Thinking	77.2	63.6	78.9
Qwen3VL-235A22B-Instruct	71.2	67.7	71.5
GPT-4o (2024-11-20)	54.3	48.9	68.5
<i>Statistical Correlation with GPQA Diamond</i>			
Pearson $R \uparrow$	–	0.733	0.926
p -value $\downarrow$	–	0.039	< 0.001

Table 9: Correlation between scientific-reasoning capability (GPQA Diamond) and LVBench.

duces cleanly (Table 9): GPQA Diamond correlates moderately with the end-to-end LVBench score ($R=0.733$, $p=0.039$) and substantially more strongly with the MEMDREAMER achieves score ($R=0.926$, $p<0.001$), mirroring the $0.702 \rightarrow 0.948$ lift on AIME 2025 (Table 3). The strengthening therefore generalises across mathematical and scientific reasoning, providing convergent evidence that MEMDREAMER tightens the dependence of long-video understanding on the reasoning capability of $\mathcal{R}$ by removing perception variance as a confound.

$T_{\max}$	LVBench	Avg. rounds	Tokens/round
8	83.7	2.87	5,990
10	84.9	3.00	6,141
12	85.5	3.06	6,181
15	85.3	3.07	6,222

Table 10: Ablation on the round budget $T_{\max}$ with $\mathcal{R}$ =Gemini-3.1-Pro and the full toolkit. Avg. rounds is the mean number of rounds actually taken before $\mathcal{R}$ emits an answer. Tokens/round is the mean input length per tool-call round.

C.2 Round Budget: $T_{\max}$ Sweep

We finally study the round budget $T_{\max}$ that hard-caps the loop. With $\mathcal{R}$ =Gemini-3.1-Pro and the full toolkit, we sweep $T_{\max} \in \{8, 10, 12, 15\}$ and, alongside accuracy, record the mean number of rounds actually used and the mean input length per round. The latter two probe whether a larger budget is exercised passively (cap exhaustion) or actively (more reasoning only on hard queries). Three trends are visible in Table 10. First, accuracy improves steadily from 83.7 to 85.5 as $T_{\max}$ grows from 8 to 12, then slightly regresses at 15 (85.3). A larger budget therefore does help, but with dimin-

Depth	Active tiers (top → leaf)	LVBench
1 (flat)	MACRO	80.9
2	SUPER → MACRO	83.6
3 (default)	ROOT → SUPER → MACRO	85.5
4	ROOT → L_3 → SUPER → MACRO	85.0

Table 11: Ablation on the depth of upward aggregation, holding the Macro Events layer and the retrieval pipeline fixed, with Gemini-3.1-Pro as the reasoning model $\mathcal{R}$. Depth 3 is the default used throughout the main paper. Depth 4 inserts an additional aggregation level L_3 between SUPER and ROOT.

ishing returns once the cap is no longer binding. Second, the average number of rounds actually used barely changes ($2.87 \rightarrow 3.07$), indicating that $\mathcal{R}$ self-terminates as soon as it has accumulated enough evidence and only spends extra rounds on queries that genuinely require them, rather than passively exhausting the cap. Third, the mean input tokens per round remain nearly flat ($5,990 \rightarrow 6,222$, +3.9%), showing that the selective compression of raw information keeps per-round context size bounded as $T_{\max}$ grows, scaling the round budget therefore comes at negligible incremental cost.

C.3 Depth of Hierarchical Aggregation

Our hierarchical memory stacks several aggregation tiers above the Macro Events layer. We ablate how many such tiers are actually beneficial, holding the Macro layer and the retrieval pipeline fixed and using Gemini-3.1-Pro as the reasoning model $\mathcal{R}$. Depth 1 corresponds to a flat memory (Macro Events only), depths 2, 3, and 4 progressively add aggregation levels above it. Aggregation contributes a substantial gain (+4.6 LVBench from depth 1 to depth 3), confirming that hierarchical organisation materially improves long-video understanding over a flat episode list. Beyond depth 3, however, the curve plateaus and slightly regresses (-0.5 at depth 4): three tiers already saturate the representational budget, while an additional aggregation step lengthens the retrieval path without introducing useful new abstractions. We therefore adopt depth 3 as the default in all main-paper experiments.

C.4 Search Breadth: top- k Sweep

The Precise Search tools (*SearchNodes*, *SearchByTime*) return the top- k neighbours of a query embedding. A natural follow-up question is whether returning more candidates per call yields strictly better evidence coverage. We sweep $k \in \{5, 10, 15\}$,

top- k	LVBench	Avg. rounds
5	82.4	2.91
10 (default)	83.7	2.87
15	82.6	2.90

Table 12: Ablation on the search top- k . We sweep the number of results returned by the Precise Search tools (*SearchNodes*, *SearchByTime*) per query, with $\mathcal{R}$ =Gemini-3.1-Pro and $T_{\max}=8$ held fixed.

holding the rest of the configuration ($\mathcal{R}$ =Gemini-3.1-Pro, $T_{\max}=8$, full toolkit enabled) fixed. The relationship between k and answer accuracy is non-monotonic and peaks at $k=10$ (Table 12). Both $k=5$ (82.4) and $k=15$ (82.6) underperform the default. We read the drop at $k=5$ as insufficient evidence coverage, relevant nodes occasionally fall outside the top-5 neighbour set, and the drop at $k=15$ as dilution: returning more candidates forces $\mathcal{R}$ to spend reasoning capacity disambiguating loosely related neighbours, which makes the truly informative cues harder to isolate. Crucially, the average number of rounds is essentially flat across the sweep (2.87–2.91), so k shapes the quality of each retrieval call rather than how many calls the agent chooses to issue.

D Checklist

The License For Artifacts. The artifacts in this academic work primarily consist of four open-source video understanding benchmarks (Wang et al., 2025a; Wu et al., 2024; Fu et al., 2025; Mangalam et al., 2023). Additionally, the visualization cases are from other existing open-source benchmarks (Song et al., 2024; Wang et al., 2025a). All usage strictly complies with their respective licenses.

Memory Streaming Segmentation Prompt

""You are an Expert Visual Story Segmenter.

[YOUR CORE OBJECTIVE]

Watch the video and segment it into "Major Story Episodes".

Each episode MUST be a complete, self-contained logical event, typically lasting between 3 to 8 minutes.

[STRICT SEGMENTATION RULES]

1. WHAT CONSTITUTES AN EPISODE: An episode is a full narrative arc for ONE subject or ONE specific phase. For example, a single contestant's entire segment (walking on stage -> singing -> judges' visual reactions -> judges talking to them -> leaving stage or screen fading/cutting) is ALL ONE EPISODE. Do not split it.

2. RELY ON PURE VISUAL CUES: Look for major visual resets to mark a new episode:

- A completely new contestant appearing on the main stage.
- Distinct broadcast graphics (e.g., a new name tag or lower-third graphic appearing).
- A major, permanent location change (e.g., shifting from the main stage to a long

backstage interview segment).

3. ANTI-FRAGMENTATION (DO NOT CUT):

- DO NOT cut when the camera simply switches angles (e.g., cutting from the singer to a reacting judge, or to the audience).
- DO NOT cut for flashing stage lights or brief visual effects if the core subject on stage remains the same.

4. BOUNDARY PRECISION (NO BLEEDING):

- The `end\_time` of an episode MUST NOT include the appearance of the next episode's subject.
- Stop the clock the exact second BEFORE the new contestant walks in, BEFORE the new stage is shown, or BEFORE the new graphics appear.
- Every episode must be 100% exclusive to its own subject.

[OUTPUT FORMAT]

Use your internal video timestamp perception to determine the time boundaries. Output ONLY a valid JSON array of objects.

Each object must have:

- `start\_time`: "HH:MM:SS" (based on video timestamp)
- `end\_time`: "HH:MM:SS" (based on video timestamp)
- `episode\_title`: A concise name (e.g., "Contestant_01_Full_Segment", "Host_Opening_Rules").
- `visual\_evidence`: Briefly explain the visual cue that justified this boundary, confirming no bleed-over.

Output raw JSON only. Do not wrap in markdown format.

""

Figure 6: Streaming Segmentation Prompt.

Memory Subgraph Construction Prompt

You are a video knowledge-graph annotator. Given a segment of video, output a single JSON object that describes its subgraph $G_i = (V_i, E_i)$ with three node groups (V^E entities, V^M micro-events) and the edges among them. Watch the video carefully before producing the output.

Task overview

Produce, for the given segment, an exhaustive but de-duplicated list of:

1. micro-events – every meaningful action, scoring play, transition, or salient happening, time-stamped at second precision.
2. entities – every PERSON, OBJECT, LOCATION, or GROUP that participates in or appears alongside those events, anchored to the moment they are most clearly visible.
3. edges – directed labelled relations among the three node groups, covering subject/object roles, spatial-attributive ties, and explicit causal links.

Node schema

- micro_events:

```
{event_id, event_type, time_range:[start_sec,end_sec], subject, object, action, description}
```

Rules:

- * event_type is free text (e.g. "scoring play", "foul", "timeout").
- * subject/object MUST be specific enough to match a single entity.
- * action describes the precise motion.
- * description is the retrieval payload: full sequence, with cause when applicable.
- * Do NOT create a micro_event for static on-screen graphics.

- entities:

```
{entity_id, name|null, entity_type in {PERSON,OBJECT,LOCATION,GROUP}, attributes, description, visual_grounding}
```

Rules:

- * name is the real name only if explicitly identified by broadcast graphics, jersey, or caption. Never invent.
- * attributes: PERSON{jersey_number, jersey_color, team, role}; OBJECT{type, color, location}; LOCATION/GROUP{name, type}.
- * description fuses attributes + visual cues into one sentence.
- * visual_grounding = {primary_time:"MM:SS", primary_time_sec:int, distinctive_features:[2-4 unique cues], spatial_hint}.

Rules:

- * text is the verbatim transcription.
- * If the same overlay persists, emit ONE node spanning its full time_range.
- * description interprets the text in context.

Edge schema

Each edge: {source_id, target_id, relation_label, relation_type, description}.

Allowed labels and their relation_type are:

- | | | |
|----------------------|---|----------------|
| - Entity -> Event: | PERFORMS, RECEIVES, USES_TOOL | (SEMANTIC) |
| - Event -> Entity: | HAPPENS_IN | (SEMANTIC) |
| - Event -> Event: | CAUSES, PREVENTS | (CAUSAL) |
| | SUBEVENT_OF | (HIERARCHICAL) |
| - Entity <-> Entity: | PART_OF, LOCATED_IN, NEXT_TO, ATTACHED_TO | (SPATIAL) |
| | IS_SAME_AS | (IDENTITY) |

Output format

Return ONLY one valid JSON object with the following top-level keys:

```
{
  "micro_events": [ {event_id, event_type, time_range:[start,end], subject, object, action, description}, ... ],
  "entities":      [ {entity_id, name, entity_type, attributes, description, visual_grounding}, ... ],
  "edges":         [ {source_id, target_id, relation_label, relation_type, description}, ... ]
}
```

Output the final JSON now.

Figure 7: Subgraph Construction Prompt.

Memory Hierarchical Aggregation Prompt

You are a video narrative analyst. Given a chronologically ordered list of Macro events { macro_id, label, time_range, summary, key_entities, event_types, ocr_texts}, fold them upward into a hierarchy $V^{Mc} \rightarrow V^S \rightarrow v^R$ and emit a single JSON object.

Task overview

Produce four things in one pass:

1. super_events – 10-20 Super Events that partition all Macro events; each captures one narrative arc.
2. macro_relations – strong logical edges between Macro events within the same Super Event.
3. super_relations – high-level edges between Super Events (not restricted to adjacent ones).
4. root – a video-level summary v^R synthesised from all Super Events.

Clustering rules (Macro $\rightarrow$ Super)

A Super Event is a contiguous run of Macro events that share (i) temporal adjacency, (ii) scene/topic, and (iii) a common goal. A boundary fires when any of the three breaks. Use as cluster signals:

- key_entities overlap between adjacent Macros (strong group signal);
- event_types continuity (e.g. a sustained run of "shot_made" macros);
- ocr_texts cues (scoreboard increments stay inside a phase; broadcast graphic changes mark boundaries).

Every Macro MUST be assigned to exactly one Super Event; each Super Event holds 3-8 Macros on average; a singleton is allowed for a self-contained segment.

Edge schema

- Macro $\rightarrow$ Macro: CAUSES (A directly triggers B), PREVENTS (A blocks B), ENABLES (A supplies a prerequisite for B). Stay within the same Super Event. Each edge needs a one-sentence reason. Omit when uncertain – never use BEFORE/AFTER.
- Super $\rightarrow$ Super: LEADS_TO (phase A naturally progresses to B), RESOLVES (B resolves a tension introduced in A), CONTRASTS_WITH (sharp reversal in outcome/tone). Cross-distance edges are allowed.

Naming and description

- super_event.label: noun phrase under 10 words; include the central entity when one exists ("LeBron's Third-Quarter Run").
- super_event.description: 100-200 words covering the arc, key entities (woven in, not listed), and any salient OCR milestones.
- key_entities are deduplicated within a Super Event AND canonicalised across the whole video: variants of the same real-world entity ("LeBron", "King James", "LeBron James") collapse to one canonical name with a consistent type. When in doubt, keep them separate.
- root.title: under 15 words, specific. root.description: 3-5 sentences synthesising ALL Super Events. root.themes: 3-5 short tags. root.key_entities: 5-10 canonical names. root.emotional_tone: 2-3 adjectives.

Output format

Return ONLY one valid JSON object, no markdown fences, no commentary:

```
{
  "super_events": [ {super_id, label, description, sub_macro_ids:[macro_id,...], time_range:[
start,end], key_entities:[{name,type},...], ... ],
  "macro_relations": [ {source, target, type, reason}, ... ],
  "super_relations": [ {source, target, type, reason}, ... ],
  "root": {title, description, themes:[...], key_entities:[...], emotional_tone}
}
```

Hard constraints: every Macro id appears exactly once across all sub_macro_ids; super_event.time_range = [min(start), max(end)] of its Macros; every relation references ids that exist in the output; emit only confident causal/structural edges.

Output the final JSON now.

Figure 8: Consolidated hierarchical aggregation prompt used by the reasoning model during memory construction. The prompt fuses the clustering, refinement, entity-unification, and global-synthesis stages of Phase 3 into a single end-to-end specification that lifts Macro Events $\mathcal{V}^{Mc}$ to Super Events $\mathcal{V}^S$ and the Video Root v^R , with the resulting hierarchy mirroring Table 7.

Agentic Retrieval Prompt

You are the reasoning model $\mathcal{R}$ in an Observation-Reason-Action loop over a hierarchical video memory. At each round you read the current observation, decide whether you can answer, and either emit a final letter A/B/C/D or call exactly one tool.

```
# Memory hierarchy
The memory is a 4-tier graph:
- Root v^R          - video-level title, description, themes, key entities.
- SuperEvent V^S     - 10-20 narrative phases with time_range, key entities, description.
- MacroEvent V^Mc    - detailed events inside a Super, with time_range and key entities.
- Subgraph G_i       - per-Macro detail: Entity / micro-Event and their edges.
Initial context contains ONLY Root + the Super Event list. Macro descriptions, subgraph contents, and edges are NOT pre-loaded - you MUST drill down with tools.

# Toolkit (call at most one per round)
Hierarchical Navigation:
- get_summary()
- get_super_events()          - list Supers
- get_macro_events(super_id)  - list Macros under a Super.
- get_subgraph(macro_id)      - full Entity/Event/edges of a Macro. Information-richest tool; never call it without
                                first identifying the macro_id.

Precise Search:
- search_nodes(query, top_k, node_types) - semantic search over Entity/Event embeddings.
- search_by_time(start_sec, end_sec)      - Macros covering a time window.

Graph Traversal:
- get_relations(node_id, direction) - outgoing/incoming/both edges of any node, typed PROGRESSION / CAUSAL / TEMPORAL /
  HIERARCHICAL / SEMANTIC / SPATIAL.

# Strategy
Default workflow: search -> identify macro_id -> get_subgraph for full detail. Macro labels are summaries; do NOT answer
who-did-what or exact-score questions from them alone. Pick the tool from the question pattern: time references ->
search_by_time; person + action -> search_nodes; causal chain -> get_relations.

# Query construction (for search_nodes and search_semantic)
- Write a DESCRIPTIVE STATEMENT, not a question - embeddings match event descriptions, not interrogatives.
- Use information SHARED across answer options as the retrieval signal; stay neutral on details where options disagree.
- Good: "A foul is committed, leading to free throws." Bad: "Australia #7 fouled" (commits to one option) or "Who
  fouled?" (interrogative).

# Round memory (Extract operator)
Previous tool observations are NOT carried forward - only your `useful_info` field persists. Distil into it every fact,
ID, and reasoning thread you will need next round (super_ids, macro_ids found, time ranges, score state, what you have
already searched, what to try next). The Exploration History block lists prior queries and which Macros remain unexplored
- use it to avoid repeating searches and to drill into unexplored macros instead of re-searching.

# Output (strict JSON, no markdown fences, no commentary)
{
  "reasoning": "What does the question ask? What do I already have? What is missing?",
  "can_answer": true | false,
  "answer": "A" | "B" | "C" | "D" | null,
  "useful_info": "All facts / IDs / reasoning threads - the only state carried to the next round.",
  "information_need": null OR {
    "type": "browse_macros | get_macro_details | search_nodes | search_time | get_relations",
    "target": "...",
    "search_hints": [...],
    "reason": "..."
  }
}
Termination: when can_answer=true emit `answer` and set information_need=null; otherwise emit one tool call via
information_need; the loop is hard-capped at T_max rounds.
```

Figure 9: Consolidated agentic retrieval prompt used by the reasoning model $\mathcal{R}$ inside the Observation-Reason-Action loop.

System Prompt for End-to-end Methods

```
""Please watch the video carefully and answer the following multiple choice question.

{question}

Instructions:
1. First, provide your step-by-step reasoning about the question based on the video content.
2. Then, choose the correct answer from the options.
3. Output your response as a JSON object with two fields:
  - "reasoning": your step-by-step analysis
  - "answer": the letter of your chosen answer (e.g., "A", "B", "C", "D", or "E")

Example:
{{
  "reasoning": "The video shows a person cooking pasta. At 2:30, they add tomatoes to the
sauce, which corresponds to option C.",
  "answer": "C"
}}

Output ONLY the JSON object, no additional text.""
```

Figure 10: System prompt used for Evaluating End-to-end methods with thinking mode.