

VIDEOEXPLORER: BOOSTING LONG VIDEO UNDERSTANDING WITH DYNAMIC TEMPORAL GROUNDING

Anonymous authors

Paper under double-blind review

ABSTRACT

Long-video understanding (LVU) is a challenging problem in computer vision. Existing methods either downsample frames for single-pass reasoning, sacrificing fine-grained details, or depend on textual reasoning over task-agnostic representations, hindering task-specific perception and exploration. In this paper, we propose **VideoExplorer**, a framework grounded in the principle of “thinking with video”, which naturally intertwines planning, temporal grounding, and scalable perception into a coherent reasoning process. Rather than reasoning over a static context, VideoExplorer iteratively formulates sub-questions, locates relevant moments, and performs task-oriented, temporally scalable video understanding until reaching the final answer, enabling faithful, efficient, and interpretable reasoning. To address the lack of LVU training resources, we construct a long-video reasoning dataset using difficulty-adaptive sampling to ensure high-quality trajectories on complex tasks. Building on this dataset, we design a two-stage training pipeline: supervised trajectory initialization followed by trajectory-level preference optimization, encouraging adaptive temporal grounding and iterative information integration guided by downstream rewards. Extensive evaluations on popular long-video understanding and reasoning benchmarks demonstrate VideoExplorer’s significant advantage over existing baselines, highlighting its robustness, adaptability, and efficiency. Our code is available in this repository ¹.

1 INTRODUCTION

Long video understanding is crucial for real-world applications such as autonomous driving, embodied agents, and large-scale surveillance. However, it poses significant challenges for current visual-language models (VLMs) Wang et al. (2024a); Bai et al. (2025); Zhu and etc (2025); Lin et al. (2023) due to the extensive input video frames. Modern VLMs Bai et al. (2025); Lin et al. (2023) address this by heavily downsampling frames and applying single-pass reasoning, which inevitably sacrifices fine-grained details and compromises reasoning accuracy. To mitigate the limitations from brute-force downsampling, recent agentic frameworks Wang et al. (2024b; 2025); Tian et al. (2025); Luo et al. (2024); Zhi et al. (2025) augment the reasoning process with retrieval over preprocessed textual assets (e.g., dense captions, object trajectories, audio transitions, etc.), allowing the model to retrieve and incorporate task-specific textual evidence during iterative inference. However, as this preprocessing is task-agnostic, the reasoning process remains constrained to static and lossy textual representations, which are often computationally expensive, thereby precluding direct, task-adaptive perception of the raw video stream.

To address to the above problems, we introduce the principle of “**thinking with video**”. Analogue to “thinking with images”, where visual models can zoom in, zoom out, or focus on specific regions adaptively, “thinking with video” treats reasoning as a dynamic process of temporally grounded exploration. Instead of reasoning over a static context, “thinking with video” allows the model to interact directly with the raw video. At each step, it decides *what to look for*, *where to watch*, and *at what temporal scale*, flexibly combining on-demand visual perception, including both fine-grained inspection and coarse-grained holistic browsing. Through this iterative process, the model builds a coherent reasoning chain that integrates localized evidence into a task-driven, temporally grounded understanding, enabling faithful, scalable, and interpretable long-video reasoning.

¹<https://anonymous.4open.science/r/Video-DeepResearch-CCEF>

Building on this principle, we propose **VideoExplorer**, a framework that unifies reasoning with dynamic temporal grounding and task-aware perceptual exploration. Given a task, VideoExplorer incrementally decomposes it into sequential sub-tasks, iteratively identifies relevant temporal segments, and attends to video content at the granularity required for each step. A planner generates intermediate information needs and coordinates reasoning, while a temporal grounder retrieves and verifies candidate segments from long videos, ensuring accurate and concise evidence. The system then performs temporal-scalable video understanding, dynamically adjusting frame sampling according to task demands, before producing the final answer. By integrating these components into a cohesive reasoning loop, VideoExplorer enables efficient, interpretable, and fine-grained long-video understanding that scales to complex tasks without unnecessary computation.

To train VideoExplorer for effective long-video reasoning, we first construct a reasoning-centric dataset with difficulty-adaptive sampling, generating multi-step reasoning trajectories and emphasizing challenging cases to ensure rich supervision beyond simple QA pairs. Based on this dataset, we adopt a two-stage optimization strategy. First, structured imitation via supervised fine-tuning (SFT) teaches the planner and temporal grounder to reproduce expert reasoning traces in addition to final answers. Then, trajectory-level direct preference optimization (TDPO) evaluates complete reasoning trajectories including intermediate steps and rewards faithful reasoning while penalizing flawed paths. This combination enables VideoExplorer to learn interpretable, task-aligned multi-step reasoning for robust performance on complex long-video tasks. We conduct extensive experiments on widely-used long video datasets, such as MLVU and LVBench for general understanding tasks, as well as VideoMMU and MH-NIAH for complex reasoning tasks. The experimental results demonstrate the effectiveness of our method. In summary, our contributions are threefold:

- (1). We introduce the concept of “thinking with video” and propose **VideoExplorer**, which integrates planning, temporal grounding, and scalable video perception into a unified reasoning paradigm.
- (2). We construct a **reasoning-centric dataset** and design a **two-stage optimization pipeline** combining structured imitation and trajectory-level preference alignment.
- (3). We conduct extensive experiments across diverse long-video benchmarks, demonstrating the scalability, accuracy, and interpretability of our approach compared to existing methods.

2 RELATED WORK

2.1 MULTI-MODAL LARGE LANGUAGE MODELS

Advancements in large language models (LLMs) (Brown et al., 2020) have significantly advanced visual understanding by incorporating visual features, leading to the emergence of Multi-modal Large Language Models (MLLMs) (Wang et al., 2024a; Chen et al., 2024a; Liu et al., 2023; OpenAI, 2024; Bai et al., 2025; Zhu and etc, 2025; Yuanhan Zhang et al., 2024; Lin et al., 2023; Zhang et al., 2024; Chen et al., 2024b). These models are designed for general-purpose vision tasks and have also demonstrated strong performance in video understanding. However, most existing MLLMs primarily focus on short-context comprehension, typically handling only a few images or short video clips at a time, which limits their effectiveness on long video understanding tasks (Fu et al., 2024; Zhou et al., 2024). To address this limitation, recent efforts have explored more efficient architectures to extend the input context (Xue et al., 2024; Shu et al., 2024; Liu et al., 2025), enabling the processing of longer videos. Nonetheless, the length of input context that these models can handle remains restricted, and understanding videos of hour-long durations remains a significant challenge. In this work, we propose VideoExplorer, which adaptively retrieves task-relevant visual context in the reasoning chain, avoiding accuracy loss from brute-force frame sampling.

2.2 LARGE REASONING MODELS

Recent advancements have significantly enhanced the reasoning capabilities of large language models. OpenAI’s o1 (OpenAI, 2025) demonstrated test-time scaling effects, enabling free-form reasoning prior to generating answers and achieving substantial performance gains. DeepSeek-R1 Team (2025) further incentivizes reasoning through a reinforcement learning (RL) algorithm called group robust preference optimization, which automates the generation of high-quality reasoning trajectories via trial-and-error search without extensive human oversight. Building on these insights, several

108 studies Feng et al. (2025); Chen et al. (2025); Zhang et al. (2025) have explored reasoning in video
 109 understanding. These works train models to generate reasoning processes before producing final
 110 answers and optimize them with RL algorithms. However, they operate solely on static down-sampled
 111 video frames, leading to significant information loss in long-video understanding. Motivated by these
 112 findings, we introduce the principle of “thinking with video”, enabling models to not only reason over
 113 contextual information but also interact with video streams through on-demand temporal grounding.

114 2.3 AGENTIC FRAMEWORKS FOR LONG VIDEO UNDERSTANDING

115 Leveraging the powerful capabilities of large language models OpenAI (2023); Brown et al. (2020),
 116 recent studies have explored agentic approaches to tackle the long video understanding chal-
 117 lenges Wang et al. (2024b); Zhi et al. (2025); Kugo et al. (2025); Cao et al. (2025). However,
 118 these approaches transit videos into task-agnostic textual representations (e.g. object locations,
 119 dense captions, audio transitions, etc.), and then leverage powerful LLMs to reasoning over the
 120 textual representations. This paradigm will inevitably lose the rich visual information in original
 121 long videos, leading to sub-optimal performance. Additionally, many of these methods are designed
 122 for limited scenarios such as egocentric videos Wang et al. (2024b); Tian et al. (2025), deducing
 123 their applicability to more general real-world settings. Building on this perspective, we present
 124 VideoExplorer, a new framework grounded in “thinking with video” that supports efficient, scalable,
 125 and adaptive long-video understanding.

126 3 METHODOLOGY

127 3.1 PRELIMINARY OF LONG VIDEO UNDERSTANDING

128 Long-video reasoning requires a model to comprehend an extended visual sequence v in order to
 129 answer a query q . This process is formally defined as:

$$130 \mathcal{Y} = \Phi(v, q), \tag{1}$$

131 where \mathcal{Y} denotes the generated answer. Ideally, the model could ingest the entire video at once
 132 and answer questions based on the complete content. However, in practice, due to the input length
 133 constraints of models, it is often difficult to process the entire video in one go. To address this issue,
 134 modern visual language models often perform brute-force downsampling of v into a shorter sequence
 135 $v' \ll v$, and then reason over the sampled v' : $\mathcal{Y} = \Phi(v', q)$. While this approach is straightforward
 136 and sufficient for short videos, it becomes problematic for hour-long videos, which typically contain
 137 tens of thousands of raw frames. Severe temporal compression inevitably sacrifices fine-grained
 138 details, leading to reasoning based on insufficient context and degraded accuracy.

139 **Agentic frameworks with retrieval augmentation** present a promising alternative. Rather than
 140 downsampling the raw video video and reasoning in a single pass, these frameworks iteratively
 141 expand the available context through retrieval augmentation. Specifically, these approaches typically
 142 preprocess the video into a fixed set of textual representations (e.g. dense captions, object locations,
 143 etc.) and construct a retrieval index based on them, denoted as \mathcal{D}_v . During inference, a textual
 144 reasoning model Θ analyzes the task, generates sub-queries, retrieves relevant information from the
 145 index, and continues reasoning until an answer is obtained. This process can be formulated as:

$$146 \mathcal{D}_v = \Gamma_{\text{preprocess}}(v), \tag{2}$$

$$147 \mathcal{Y} = \Theta(q | \{\mathcal{K}_t\}_{t=1}^T), \quad \mathcal{K}_t = \Gamma_{\text{retrieve}}(q_t^{\text{sub}} | \mathcal{D}_v), \tag{3}$$

148 where q_t^{sub} represents the sub-queries generated by the reasoning model. This iterative retrieval-
 149 augmented approach offers a more scalable solution for long video reasoning compared to single-pass
 150 methods, as it alleviates the problem of insufficient context.

151 This iterative retrieval alleviates context limits and scales better than single-pass methods. However,
 152 task-agnostic preprocessing compresses rich video signals into generic text, causing information loss.
 153 Moreover, such preprocessing is computationally heavy, limiting scalability in practice.

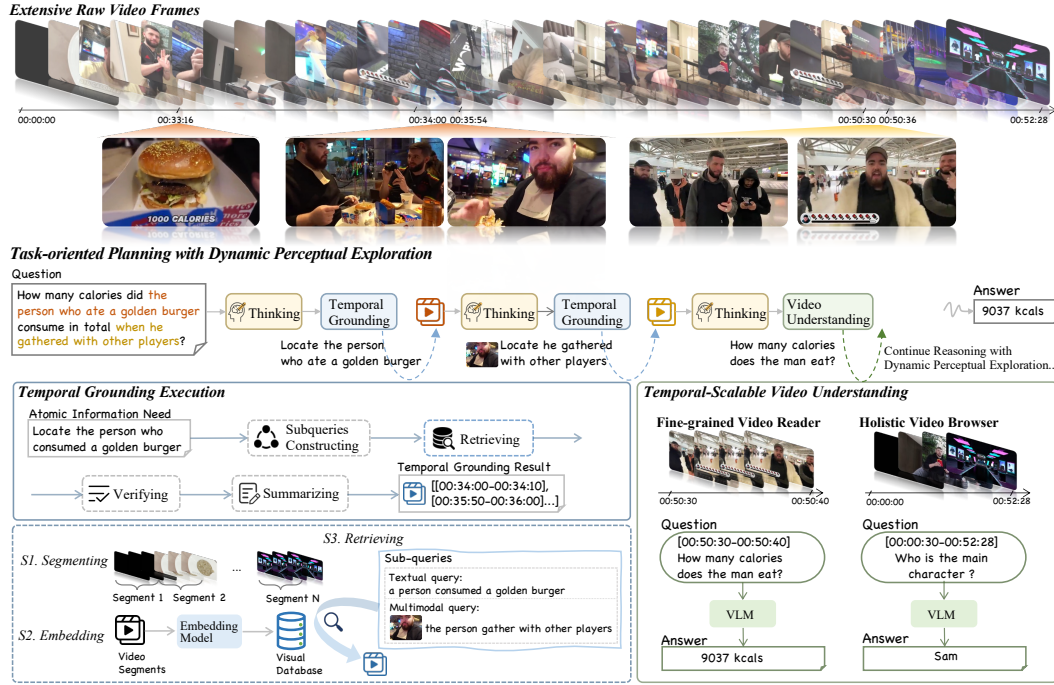


Figure 1: Illustration of VideoExplorer framework. Instead of reasoning over brute-force downsam-pled frames or static preprocessed database, VideoExplorer integrates planning, temporal grounding, and scalable video perception into a unified reasoning paradigm, where the planner decomposes complex tasks into sub-questions, the temporal grounder adaptively localizes relevant temporal spans, and the perception module dynamically adjusts granularity to meet task demands, enabling faithful, efficient and interpretable long-video understanding.

3.2 THINKING WITH VIDEO

Inspired by the idea of “thinking with images”, where models can adaptively zoom in, zoom out, or focus on specific regions for a task, we introduce “thinking with video”. This paradigm treats reasoning as a dynamic process over on-demand temporal context, where the model raises temporally scalable perceptual queries tailored to sub-tasks, ultimately arriving at the final answer. Rather than relying on brute-force downsampling or task-agnostic preprocessing, “thinking with video” retains direct access to the raw video and adaptively gathers task-relevant information through task-aware temporal grounding and scalable video understanding. Formally, the process is defined as:

$$\mathcal{Y} = \Theta(q | \{\mathcal{P}_t\}_{t=1}^T), \quad \mathcal{P}_t = \Phi_{\text{perc}}(\tau_t), \quad \tau_t = \Gamma_{\text{ground}}(q_t^{\text{sub}} | v), \quad (4)$$

where sub-questions q_t^{sub} are generated by a planner, Γ_{ground} maps them to temporal spans τ_t , and Φ_{perc} performs on-demand perception within these spans. This iterative mechanism enables step-by-step reasoning by selectively accessing and processing task-relevant segments of the raw video, rather than compressing or statically pre-processing it.

3.3 OUR PROPOSED METHOD: VIDEOEXPLORER

Building on the principle of “thinking with video”, we introduce the **VideoExplorer** framework, which achieves scalable long-video reasoning by unifying the reasoning process with dynamic temporal grounding. As illustrated in Figure 1, the generation process unfolds as a sequential composition of cognitive blocks: a thinking block captures intermediate reasoning, a temporal grounding block localizes relevant temporal spans, a video understanding block performs temporal scalable video understanding and return results, and finally an answer block delivers the final response.

Task-oriented Planning with Dynamic Perceptual Exploration. The planner, analogous to the human central nervous system, is responsible for analyzing the problem, raising iterative information

216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269

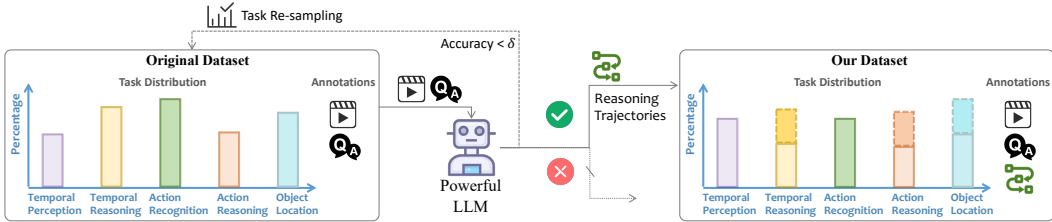


Figure 2: Difficulty-adaptive dataset generation. Tasks are uniformly sampled, reasoning trajectories generated by VideoExplorer, and hard cases re-sampled by first-round accuracy. Only correct-answer trajectories are retained, yielding faithful reasoning trajectories and challenging training data.

needs, reasoning over the current context, and generating the final answer. Instead of passively receiving a fixed set of video frames, the planner can decompose complex tasks, sequentially raise sub-questions, and attach the necessary video segments at controllable temporal granularity, serving as a higher-level controller that guides the overall reasoning trajectory.

Decoupled Temporal Grounding Execution. While the planner specifies required information, precise retrieval from long videos remains difficult due to query ambiguity and redundant searches. To reduce reasoning complexity, we introduce a temporal grounder which decouples the retrieval process from the main reasoning chain. In the offline stage, videos are segmented into clips and stored as multimodal embeddings. At inference, the grounder matches queries against candidate segments, supporting both text-only and multimodal queries (as shown in the Figure 1). It transforms abstract sub-queries into concrete actions such as construction, retrieval, verification, and summarization, and returns validated segments to the planner. For verification, the temporal grounder leverages VLMs as accurate verifiers to assess whether retrieved segments are relevant. Irrelevant results are discarded, while validated segments are summarized and passed back to the planner. This decoupling enables lightweight reasoning, ensures accurate evidence localization, and mitigates hallucinations from irrelevant segments, allowing the planner to focus on high-level tasks.

Temporal-Scalable Video Understanding. Different tasks require different levels of visual granularity. The planner adapts by specifying temporal intervals for perception. For fine-grained queries, it focuses on task-oriented segments located by temporal grounder or specified by the task (e.g., “What happens between 01:20-01:40?”), during which the VLM densely samples frames to obtain accurate answers; for global tasks (e.g., “Who is the main character?”), it covers broader intervals with coarse sampling to support holistic reasoning. This dynamic adjustment balances scalability and fidelity, enabling efficient context management within limited window length while preserving access to high-resolution details when needed.

In summary, VideoExplorer unifies reasoning with dynamic temporal grounding to enable scalable long-video understanding. It features a planner that decomposes complex tasks into iterative sub-questions, a decoupled temporal grounder that localizes and verifies relevant video spans, and a temporal-scalable video understanding module that adapts perception granularity to task demands. Together, these components operate in a coordinated reasoning loop, allowing VideoExplorer to progressively decompose tasks, ground relevant video segments, and adapt perceptual focus, resulting in faithful, efficient, and interpretable long-video reasoning.

3.4 MULTI-STAGE OPTIMIZATION

To enable effective long-video reasoning, we construct a reasoning-centric dataset with difficulty-adaptive sampling, ensuring soft supervising beyond simple outcome supervision and coverage of complex cases. Building on this foundation, we propose a two-stage training framework: (1) structured imitation via supervised fine-tuning (SFT) to instill expert reasoning patterns, and (2) trajectory-level direct preference optimization (TDPO) to refine multi-step reasoning by rewarding faithful trajectories while penalizing flawed ones.

Constructing a Reasoning-Centric Dataset for Long Video Understanding. A major challenge in long-video reasoning is the lack of high-quality, reasoning-oriented training data. To address this, we first select a diverse set of long videos from existing datasets, considering video length, source,




Figure 3: Data statistics of VideoExplorer dataset.

and task type. However, these datasets only provide QA labels, which are insufficient for effective training. To enrich them, we leverage a powerful LLM Team (2025) with VideoExplorer to generate additional reasoning data. Since many of the generated tasks are relatively simple, we introduce difficulty-adaptive sampling to emphasize complex reasoning. As illustrated in Figure 2, tasks are uniformly sampled, reasoning trajectories are generated by the LLM with VideoExplorer, and hard cases are re-sampled based on first-round accuracy to construct a more challenging dataset. Only trajectories with correct answers are retained to ensure validity. In total, we curate 11.1k planner trajectories and 10.8k grounding trajectories, with detailed dataset statistics provided in Figure 3.

Stage I: Structured Imitation via SFT. We initialize both the planner and the temporal grounding agent through supervised fine-tuning on expert trajectories. Unlike conventional direct QA supervision, this stage emphasizes structural imitation, where the model learns to reproduce expert reasoning traces in addition to final answers. The objective is the standard negative log-likelihood:

$$\mathcal{L}_{\text{SFT}}(\theta) = -\mathbb{E}(v, q, \mathcal{Y}) \sim \mathcal{D}_{\text{SFT}} [\log \pi_{\theta}(\mathcal{Y} | v, q)], \quad (5)$$

where \mathcal{D}_{SFT} includes both answer labels and their associated expert reasoning trajectories. This step enables fast initialization by grounding the model’s behavior in high-quality, interpretable reasoning.

Stage II: Trajectory-Level Preference Optimization. While SFT provides a fast warmup, it is inherently limited by teacher-forced demonstrations. To move beyond imitation, we introduce trajectory-based direct preference optimization (TDPO), which adapts the DPO framework to the multi-turn and multi-step nature of video reasoning. Instead of evaluating individual responses, TDPO considers entire reasoning trajectories, capturing both intermediate steps and final outcomes. By comparing chosen and rejected trajectories, it encourages the policy to align faithful reasoning with correct task results. Formally, for trajectories \mathbf{y}_w and \mathbf{y}_l , the TDPO objective is:

$$\mathcal{L}_{\text{TDPO}}(\pi_{\theta}; \pi_{\text{ref}}) = -\mathbb{E}_{(\mathbf{y}_w, \mathbf{y}_l) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left(\log \frac{\pi_{\theta}(\mathbf{y}_w)}{\pi_{\text{ref}}(\mathbf{y}_w)} - \log \frac{\pi_{\theta}(\mathbf{y}_l)}{\pi_{\text{ref}}(\mathbf{y}_l)} \right) \right) \right] \quad (6)$$

Here, π_{θ} is the learnable policy, π_{ref} is reference policy, β controls the preference sharpness, and the expectation is over trajectory pairs $(\mathbf{y}_w, \mathbf{y}_l)$ from dataset \mathcal{D} . TDPO optimizes model performance by enlarging the reward gap between preferred and rejected outputs while penalizing divergence from a reference model, guiding the planner toward adaptive multi-step visual exploration.

Together, the two stages define an efficient and effective training methodology for long-video reasoning, which offers a general recipe for teaching models to reason with lightweight traces, where imitation provides structure and preference optimization provides alignment.

4 EXPERIMENT

4.1 SETTINGS

Baselines. We compare VideoExplorer against both state-of-the-art VLMs and agentic frameworks with retrieval augmentation. For VLMs methods, we compare **GPT-4o** OpenAI (2024) and **Gemini-1.5-pro** Team and etc (2024). For agentic frameworks, we consider: (1). **Vanilla**, which feeds

Table 1: Main experimental results on long video benchmarks. To ensure a fair comparison of reasoning frameworks, all baselines utilize the same visual encoder, compatible with both the Qwen-2.5-VL-7B and 32B models. We report results using each method’s fine-tuned LLM where available; otherwise, we use the zero-shot Qwen2.5-32B model.

Model	LLM Params	VLM Params	Long Video Understanding		Reasoning	Avg
			LVBench	MLVU	MH-NIAH	
<i>Proprietary VLMs</i>						
GPT-4o	-	-	48.9	54.9	-	-
Gemini-1.5-pro	-	-	33.1	53.8	-	-
<i>Direct Reasoning VLMs</i>						
Video-R1	-	7B	38.3	46.2	42.0	42.2
<i>Agentic Frameworks (Qwen2.5VL-7B)</i>						
Vanilla	-	7B	34.5	36.1	39.0	36.5
VideoAgent	32B	7B	34.9	38.6	<u>43.0</u>	38.8
VideoTree	32B	7B	30.3	38.9	41.1	36.8
VideoRAG	-	7B	36.2	37.8	39.4	37.8
Ego-R1	3B	7B	39.6	40.9	42.3	41.0
VideoExplorer	3B	7B	<u>42.0</u>	<u>46.3</u>	39.9	<u>42.7</u>
VideoExplorer	7B	7B	50.6	55.4	49.1	51.7
<i>Agentic Frameworks (Qwen2.5VL-32B)</i>						
Vanilla	-	32B	34.9	41.8	39.5	38.7
VideoAgent	32B	32B	35.8	42.7	<u>43.7</u>	40.7
VideoTree	32B	32B	32.4	39.7	41.4	37.8
VideoRAG	-	32B	37.1	42.7	39.8	39.9
Ego-R1	3B	32B	41.6	43.0	42.6	42.4
VideoExplorer	3B	32B	<u>43.1</u>	<u>49.0</u>	40.7	<u>44.3</u>
VideoExplorer	7B	32B	51.4	58.6	53.4	54.5

densely sampled frames at 1 fps to the VLM to directly generate answers. (2). **VideoAgent** Wang et al. (2024b), which iteratively locates the informative segments with pre-calculated captions. (3). **VideoTree** Wang et al. (2025), which builds a query-adaptive and hierarchical video representation for LLM reasoning. (4). **VideoRAG** Luo et al. (2024), which employs visually-aligned auxiliary texts to help facilitate cross-modality alignment while providing additional information beyond the visual content. (5). **Ego-R1** Tian et al. (2025) employs a reinforcement learning-trained agent that navigates a hierarchical video index to execute a structured chain-of-thought-for-tool process. We evaluated all models with two different visual modules: Qwen2.5VL-7B and 32B. We preferentially used each model’s official weights. For models without publicly available weights, we substituted Qwen2.5-32B and evaluated it in a zero-shot manner. We conducted all analyses using the 7B parameter versions of the LLM and the visual module. Full implementation details for all baselines and VideoExplorer are provided in Appendix B.

Dataset. We evaluate our method on two categories of benchmarks: (1) *Long Video Understanding* on LVBench Wang et al. (2024c) and MLVU (test set) Zhou et al. (2024), (2) *Multi-hop Reasoning* on MH-NIAH Li et al. (2025) (max frames is set to 10,000). We report accuracy for all benchmarks. Besides, we assessed temporal grounding effectiveness using the IoU@0.1 metric, which compares the framework-grounded temporal segments to the golden temporal intervals provided by LVBench.

4.2 MAIN EXPERIMENT

In Table 1 and Table 2, we present the main experimental results and temporal grounding results on the LVBench, from which we have several key findings: (1) Search-enhanced reasoning methods (e.g. VideoAgent, Ego-R1 and VideoExplorer) show particular strength on multi-hop reasoning tasks such as MH-NIAH. Their ability to iteratively decompose queries, formulate subgoals, and retrieve contextually relevant information mirrors human-like problem-solving behavior and leads to consistently stronger results compared to static, one-shot retrieval pipelines such as VideoRAG. VideoExplorer further benefits from reward-guided reasoning, which encourages the generation of effective sub-queries and efficient trajectories. (2) VideoExplorer significantly outperforms all base-

Table 2: Evaluating Temporal Grounding Accuracy and Question Answering Accuracy on the LVBench Dataset.

Model	VLM	IoU@0.1	Accuracy
VideoAgent	7B	16.7	35.9
VideoRAG	7B	14.5	36.2
Ego-R1	7B	19.6	39.6
VideoExplorer	7B	27.8	50.6

Table 3: Ablation Study. Orange, purple and green rows indicate ablations on framework effectiveness, training strategy, and dataset effectiveness.

Settings	MLVU	MH-NIAH
VideoExplorer	55.4	49.1
- w/o Decoupled TG	53.0	47.3
- w/o TDPO	52.1	47.0
- w/o Difficulty Sampling	53.3	47.6




Figure 4: Visual Token Usage Comparison.

lines across all the benchmarks, demonstrating strong robustness and generalizability. This reveals the effectiveness of reasoning with on-demand temporal grounding and visual perception, leading to robust performance across both general long video understanding benchmarks and complex reasoning tasks such as multi-hop reasoning and tutorial reasoning. (3) Compared with retrieval-augmented agentic frameworks, VideoExplorer achieves higher temporal grounding accuracy. Existing frameworks struggle to precisely localize relevant segments, leading to lower QA accuracy. In contrast, VideoExplorer decomposes questions into reasoning steps, aligns them with relevant video segments, and answers based on accurately grounded evidence. As shown in Table 2, prior agentic frameworks fail to reliably identify key segments, resulting in lower overall accuracy, while our method significantly outperforms baselines in IoU@0.1 and QA accuracy. This confirms our hypothesis that adaptive, reasoning-aware temporal grounding is crucial for long-video understanding. VideoExplorer’s principled integration of temporal grounding into the reasoning loop, guided by downstream rewards, allows it to dynamically navigate complex visual spaces more effectively than single-turn or static strategies.

4.3 ABLATION STUDY

We evaluate the impact of VideoExplorer’s components through ablation experiments (Table 3). (1) **Training Strategies:** Removing TDPO causes a notable drop of 3.3% on MLVU and 2.1% on MH-NIAH, showing that it is crucial for stable optimization. By enlarging the reward gap between preferred and rejected outputs based on environment feedback, TDPO encourages the planner to avoid shallow reasoning and instead pursue adaptive, multi-step exploration, ultimately improving the reliability of generated trajectories. (2) **Framework Design:** Decoupling temporal grounding from reasoning reduces the planner’s burden and prevents premature answering. Without the grounding agent, performance decreases by 2.4% and 1.8%, as the planner struggles to gather sufficient evidence before producing an answer. This highlights that explicit grounding enables the planner to focus on higher-level reasoning while relying on the grounder to supply relevant temporal spans. (3) **Dataset Effectiveness:** Using only uniformly sampled data leads to 2.1 and 1.5 point drops, confirming that difficulty-adaptive sampling plays an important role. By exposing the model to harder and more diverse cases, it strengthens the reasoning ability and prevents overfitting to trivial patterns, resulting in better generalization on evaluation benchmarks. (4) **Model Scaling:** Expanding the planner from 3B to 7B yields consistent improvements across both datasets, indicating that VideoExplorer can effectively scale with larger model capacity. This suggests that the modular design aligns well with stronger backbones and provides a clear path for further enhancement. Overall, the ablations demonstrate that each component, including training strategies, framework design, dataset construction, and model scaling, contributes meaningfully to the final performance, and their integration is key to the effectiveness of VideoExplorer.

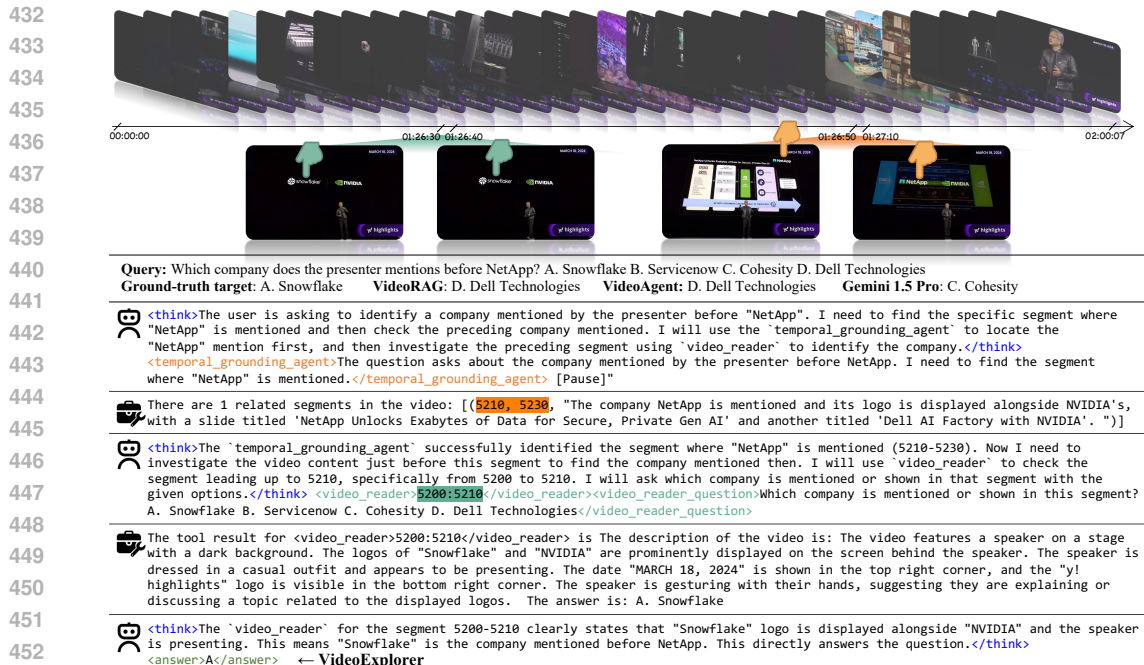


Figure 5: Case Study of VideoExplorer on LVBench. VideoExplorer correctly decomposes the task, and progressively identifies the relevant video segment through iterative reasoning with temporal grounding. It then performs fine-grained dense perception on this segment to deliver an accurate answer. In contrast, baseline methods all fail on this type of multi-hop fine-grained reasoning task.

4.4 TOKEN EFFICIENCY

We compare total visual token usage of different methods in Figure 4. Unlike methods such as VideoAgent and Ego-R1, which indiscriminately process all video segments and thus incur heavy token consumption, VideoExplorer performs dynamic reasoning by selectively focusing only on on-demand task-relevant segments. Despite using far fewer visual tokens, VideoExplorer achieves significantly better performance. This demonstrates its ability to progressively reason and precisely localize meaningful video content, leading to both higher efficiency and stronger effectiveness.

4.5 CASE STUDY

We illustrate VideoExplorer’s reasoning on an LVBench example in Figure 5. The question, “Which company does the presenter mention before NetApp?”, requires multi-hop reasoning: locating NetApp and then tracing backward. Gemini-1.5-pro fails due to downsampling, and traditional RAG methods miss evidence since one-pass retrieval cannot support multi-hop reasoning. In contrast, VideoExplorer decomposes the task, explores iteratively, and finds the correct answer, demonstrating reasoning through dynamic exploration. Additional cases and failure analyses are in Appendix C.

5 CONCLUSION

In this work, we introduce VideoExplorer, a novel agentic framework for long-video understanding that follows the principle of “thinking with video”. VideoExplorer decomposes complex tasks into sub-questions, adaptively grounds relevant temporal spans, and dynamically adjusts perceptual granularity to enable faithful and interpretable reasoning. This design overcomes the limitations of approaches that rely solely on downsampled frames or static video representations. To fully unlock its capabilities, we curate a specialized dataset and develop a two-stage training pipeline tailored for reasoning-intensive video tasks. Extensive experiments on popular long-video understanding and reasoning benchmarks show that VideoExplorer significantly outperforms existing methods, demonstrating its robustness, adaptability, and efficiency in agentic video understanding.

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

ETHICS STATEMENT

Our dataset builds on the open-source VideoMarathon Lin et al. (2025), which has undergone ethical review by its creators and is released under the Apache License 2.0. We did not modify the original videos; instead, we added textual annotations (reasoning trajectories) to construct the dataset. Our models are initialized from the open-source Qwen-2.5-3B and Qwen-2.5-7B, trained solely on this dataset without harmful content. If any ethical concerns arise, please contact us for resolution.

REPRODUCIBILITY STATEMENT

To comply with anonymity requirements, we provide the two-stage training and evaluation code in an anonymous repository ². Upon acceptance, we will release the full codebase, model checkpoints, and training dataset publicly.

²<https://anonymous.4open.science/r/Video-DeepResearch-CCEF>

REFERENCES

- 540
541
542 Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu,
543 Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the
544 world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024a.
- 545 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang,
546 Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan,
547 Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng,
548 Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025. URL
549 <https://arxiv.org/abs/2502.13923>.
- 550
551 Jinguo Zhu and etc. Internvl3: Exploring advanced training and test-time recipes for open-source
552 multimodal models, 2025. URL <https://arxiv.org/abs/2504.10479>.
- 553 Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united
554 visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023.
- 555
556 Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. Videoagent: Long-form video
557 understanding with large language model as agent, 2024b. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2403.10517)
558 [2403.10517](https://arxiv.org/abs/2403.10517).
- 559
560 Ziyang Wang, Shoubin Yu, Elias Stengel-Eskin, Jaehong Yoon, Feng Cheng, Gedas Bertasius, and
561 Mohit Bansal. Videotree: Adaptive tree-based video representation for llm reasoning on long
562 videos, 2025. URL <https://arxiv.org/abs/2405.19209>.
- 563 Shulin Tian, Ruiqi Wang, Hongming Guo, Penghao Wu, Yuhao Dong, Xiuying Wang, Jingkan
564 Yang, Hao Zhang, Hongyuan Zhu, and Ziwei Liu. Ego-r1: Chain-of-tool-thought for ultra-long
565 egocentric video reasoning, 2025. URL <https://arxiv.org/abs/2506.13654>.
- 566
567 Yongdong Luo, Xiawu Zheng, Xiao Yang, Guilin Li, Haojia Lin, Jinfa Huang, Jiayi Ji, Fei Chao,
568 Jiebo Luo, and Rongrong Ji. Video-rag: Visually-aligned retrieval-augmented long video com-
569 prehension, 2024. URL <https://arxiv.org/abs/2411.13093>.
- 570 Zhuo Zhi, Qiangqiang Wu, Minghe shen, Wenbo Li, Yinchuan Li, Kun Shao, and Kaiwen Zhou.
571 Videoagent2: Enhancing the llm-based agent system for long-form video understanding by
572 uncertainty-aware cot, 2025. URL <https://arxiv.org/abs/2504.04471>.
- 573
574 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
575 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
576 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 577 Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong,
578 Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to
579 commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*, 2024a.
- 580
581 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv*
582 *preprint arXiv:2304.08485*, 2023.
- 583
584 OpenAI. Gpt-4o. <https://openai.com/index/hello-gpt-4o/>, May 2024.
- 585
586 Yuanhan Zhang, Bo Li, Haotian Liu, Yong Jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu,
587 and Chunyuan Li. Llava-next-video: Advancing video understanding with llava-next. April 2024.
588 URL <https://llava-vl.github.io/blog/2024-04-30-llava-next-video/>.
Accessed: 2024-05-15.
- 589
590 Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video in-
591 struction tuning with synthetic data, 2024. URL <https://arxiv.org/abs/2410.02713>.
- 592
593 Lin Chen, Xilin Wei, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong
Duan, Bin Lin, Zhenyu Tang, et al. Sharegpt4video: Improving video understanding and genera-
tion with better captions. *arXiv preprint arXiv:2406.04325*, 2024b.

- 594 Chaoyou Fu, Yuhan Dai, Yondong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu
595 Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation
596 benchmark of multi-modal llms in video analysis. *arXiv preprint arXiv:2405.21075*, 2024.
597
- 598 Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang,
599 Tiejun Huang, and Zheng Liu. Mlvu: A comprehensive benchmark for multi-task long video
600 understanding. *arXiv preprint arXiv:2406.04264*, 2024.
- 601 Fuzhao Xue, Yukang Chen, Dacheng Li, Qinghao Hu, Ligeng Zhu, Xiuyu Li, Yunhao Fang, Haotian
602 Tang, Shang Yang, Zhijian Liu, et al. Longvila: Scaling long-context visual language models for
603 long videos. *arXiv preprint arXiv:2408.10188*, 2024.
604
- 605 Yan Shu, Peitian Zhang, Zheng Liu, Minghao Qin, Junjie Zhou, Tiejun Huang, and Bo Zhao.
606 Video-xl: Extra-long vision language model for hour-scale video understanding. *arXiv preprint*
607 *arXiv:2409.14485*, 2024.
- 608 Zhijian Liu, Ligeng Zhu, Baifeng Shi, Zhuoyang Zhang, Yuming Lou, Shang Yang, Haocheng Xi,
609 Shiyi Cao, Yuxian Gu, Dacheng Li, et al. Nvila: Efficient frontier visual language models. In
610 *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 4122–4134, 2025.
611
- 612 OpenAI. Learning to reason with LLMs, 2025. URL [https://openai.com/index/
613 learning-to-reason-with-llms/](https://openai.com/index/learning-to-reason-with-llms/).
- 614 DeepSeek-AI Team. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement
615 learning, 2025. URL <https://arxiv.org/abs/2501.12948>.
- 616 Kaituo Feng, Kaixiong Gong, Bohao Li, Zonghao Guo, Yibing Wang, Tianshuo Peng, Junfei Wu,
617 Xiaoying Zhang, Benyou Wang, and Xiangyu Yue. Video-r1: Reinforcing video reasoning in
618 mllms, 2025. URL <https://arxiv.org/abs/2503.21776>.
- 619 Yukang Chen, Wei Huang, Baifeng Shi, Qinghao Hu, Hanrong Ye, Ligeng Zhu, Zhijian Liu, Pavlo
620 Molchanov, Jan Kautz, Xiaojuan Qi, Sifei Liu, Hongxu Yin, Yao Lu, and Song Han. Scaling rl to
621 long videos, 2025. URL <https://arxiv.org/abs/2507.07966>.
- 622 Xingjian Zhang, Siwei Wen, Wenjun Wu, and Lei Huang. Tinyllava-video-r1: Towards smaller
623 llms for video reasoning, 2025. URL <https://arxiv.org/abs/2504.09641>.
- 624 OpenAI. Gpt-4 technical report, 2023.
625
- 626 Noriyuki Kugo, Xiang Li, Zixin Li, Ashish Gupta, Arpandeeep Khatua, Nidhish Jain, Chaitanya
627 Patel, Yuta Kyuragi, Yasunori Ishii, Masamoto Tanabiki, Kazuki Kozuka, and Ehsan Adeli.
628 Videomultiagents: A multi-agent framework for video question answering, 2025. URL <https://arxiv.org/abs/2504.20091>.
- 629 Shiwen Cao, Zhaoxing Zhang, Junming Jiao, Juyi Qiao, Guowen Song, Rong Shen, and Xiangbing
630 Meng. Masr: Self-reflective reasoning through multimodal hierarchical attention focusing for
631 agent-based video understanding, 2025. URL <https://arxiv.org/abs/2504.17213>.
- 632 Gemini Team and etc. Gemini 1.5: Unlocking multimodal understanding across millions of tokens
633 of context, 2024. URL <https://arxiv.org/abs/2403.05530>.
- 634 Weihang Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Xiaotao Gu, Shiyu
635 Huang, Bin Xu, Yuxiao Dong, Ming Ding, and Jie Tang. Lvbench: An extreme long video
636 understanding benchmark, 2024c. URL <https://arxiv.org/abs/2406.08035>.
- 637 Xinhao Li, Yi Wang, Jiashuo Yu, Xiangyu Zeng, Yuhan Zhu, Haiyan Huang, Jianfei Gao, Kun-
638 chang Li, Yanan He, Chenting Wang, Yu Qiao, Yali Wang, and Limin Wang. Videochat-
639 flash: Hierarchical compression for long-context video modeling, 2025. URL <https://arxiv.org/abs/2501.00574>.
- 640 Jingyang Lin, Jialian Wu, Ximeng Sun, Ze Wang, Jiang Liu, Yusheng Su, Xiaodong Yu, Hao Chen,
641 Jiebo Luo, Zicheng Liu, et al. Unleashing hour-scale video training for long video-language
642 understanding. *arXiv preprint arXiv:2506.05332*, 2025.
643
644
645
646
647

648 Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiayi Cui, HongFa Wang, Yatian Pang, Wenhao Jiang,
649 Junwu Zhang, Zongwei Li, et al. Languagebind: Extending video-language pretraining to n-
650 modality by language-based semantic alignment. *arXiv preprint arXiv:2310.01852*, 2023.
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

A USE OF LLMs

The entire research project, including the formulation of the core idea, all experimental work was carried out without LLM assistance.


B IMPLEMENTATION DETAILS

Our experiments are based on the open-source implementations of most baselines. As VideoAgent’s code is not publicly available, we replicated the method based on its paper. To ensure the fairness and rationality of our evaluation, we implement the following controls: (1). The maximum input length of the VLM is set to 32 frames. For Vanilla and VideoRAG, each video is uniformly downsampled to 32 frames as input to the VLM; for Ego-R1 and VideoExplorer, each VLM call to the VLM does not exceed 32 frames in length. (2). For all benchmarks, all baselines rely solely on visual information to answer questions; audio is not provided to prevent unfairness due to information asymmetry.

For VideoAgent and VideoTree, raw videos are segmented into 4s video segments and dense captioned by the VLM. For the Ego-R1 baseline, we adopted the implementation from its official code repository, which builds a hierarchical index using both 10-second and 600-second segments. Both Qwen2.5-VL-7B and Qwen2.5-VL-32B are deployed using the vLLM inference framework to accelerate the inference process. All experiments are conducted on 1 node of 8*80G H100 gpu.

For VideoExplorer, we initialize both planner and temporal grounder with Qwen2.5-7B . We use LanguageBind Zhu et al. (2023) as the video clip embedder. For textual query, we use LanguageBind Zhu et al. (2023) to calculate the text embedding, and calculate similarity with the video clip embeddings. For multimodal query, we average the textual embedding and query image embedding to get the multimodal embedding, and calculate similarity with the video clip embeddings. We set top-k to 10 to balance the accuracy and efficiency. The maximum turns of planner is 20.

C SUPPLEMENTARY CASE STUDIES AND FAILURE ANALYSIS



00:00:00 00:06:18

Query: In what setting does the entire video take place? A. Beach B. Library C. Desert D. City E. Forest F. Campus
Ground-truth target: D **VideoAgent:** F **Gemini 1.5 Pro:** D

🤖 `<think>`The question asks about the overall setting of the entire video. This is a holistic question that the `'video_browser'` agent is best suited to answer. I will use `'video_browser'` to get a general idea of the video's setting, and then use `'temporal_grounding_agent'` and `'video_reader'` to confirm it if necessary.`</think>``<video_browser>`In what setting does the entire video take place? A. Beach B. Library C. Desert D. City E. Forest F. Campus`</video_browser>`

🤖 The tool results for `<video_browser>`In what setting does the entire video take place? A. Beach B. Library C. Desert D. City E. Forest F. Campus`</video_browser>` is:The video features scenes that include a car interior, a street with a yellow bus, and a lively street market, which are all typical of a city environment. Therefore, the correct answer is D. City.

🤖 `<think>`The `'video_browser'` agent has provided a clear answer, stating that the video's setting is a city based on the scenes of a car interior, a street with a yellow bus, and a lively street market. This seems to be a sufficient answer for the holistic question. I will proceed with this answer.`</think>``<answer>`D`</answer>`" ← **VideoExplorer**

Figure 6: Case Study of Holistic Task on MLVU.

756
757
758
759
760
761
762
763

764 **Query:** Which of the following options correctly matches the sequence of actions as they actually appear in the video?
 765 A. jetskiing --> shredding paper --> paragliding --> abseiling
 766 B. abseiling --> jetskiing --> paragliding --> shredding paper
 767 C. paragliding --> shredding paper --> abseiling --> jetskiing
 768 D. abseiling --> shredding paper --> jetskiing --> paragliding
 769 E. jetskiing --> paragliding --> abseiling --> shredding paper
 770 F. shredding paper --> jetskiing --> abseiling --> paragliding
 771 **Ground-truth target:** F **VideoAgent:** F. **Gemini 1.5 Pro:** C.

772 **Q:** `<think>` The user is asking to identify the correct sequence of four actions as they appear in the video. I need to find the temporal order of "jetskiing", "shredding paper", "paragliding", and "abseiling". I will use the temporal grounding agent to locate each of these actions and then determine their order.`</think>` `<temporal_grounding_agent>` Locate the segments for "jetskiing", "shredding paper", "paragliding", and "abseiling". `</temporal_grounding_agent>`

773 **A:** There are 4 related segments in the video: [(200, 220, 'Jetskiing'), (40, 50, 'Shredding paper'), (600, 610, 'Abseiling'), (630, 640, 'Paragliding')]

774 **Q:** `<think>` I have received the timestamps for all four actions: - Jetskiing: 200-220 - Shredding paper: 40-50 - Abseiling: 600-610 - Paragliding: 630-640 Now I need to arrange them in chronological order based on their start times. 1. Shredding paper (40-50) 2. Jetskiing (200-220) 3. Abseiling (600-610) 4. Paragliding (630-640) Let's check the options. A. jetskiing --> shredding paper --> paragliding --> shredding paper (Incorrect) B. abseiling --> jetskiing --> paragliding --> shredding paper (Incorrect) C. paragliding --> shredding paper --> abseiling --> jetskiing (Incorrect) D. abseiling --> shredding paper --> jetskiing --> paragliding (Incorrect) E. jetskiing --> paragliding --> abseiling --> shredding paper (Incorrect) F. shredding paper --> jetskiing --> abseiling --> paragliding (Correct based on 40-50, 200-220, 600-610, 630-640) The sequence derived from the timestamps matches option F.`</think>` `<answer>`F`</answer>` ← **VideoExplorer**

Figure 7: Case Study of Multi-detail Task on MLVU.

781
782
783
784
785
786
787
788

789 **Query:** What did the yellow sponge use to fly into the sky? A. Wings B. Spoon C. Boat D. Pot Lid E. Bamboo Dragonfly F. Shovel
 790 **Ground-truth target:** F **VideoAgent:** A **Gemini 1.5 Pro:** A

791 **Q:** `<think>` The user is asking a question about what the yellow sponge used to fly. I need to first locate the segment where the yellow sponge flies into the sky. Then, I will examine that segment to determine what it used for lift-off. I will use the temporal grounding agent to find the relevant segment.`</think>` `<temporal_grounding_agent>` The yellow sponge uses something to fly into the sky. What did it use? A. Wings B. Spoon C. Boat D. Pot Lid E. Bamboo Dragonfly F. Shovel`</temporal_grounding_agent>`

792 **A:** There are 4 related segments in the video: [(190, 200, 'The yellow sponge flies using a spatula and megaphone. '), (210, 220, 'The yellow sponge flies using a propeller. '), (390, 400, 'The yellow sponge flies using a large, white, rectangular device with a red propeller at the back. '), (420, 430, 'The yellow sponge flies using a modified boat and a mechanical device attached to it. ')]

793 **Q:** `<think>` The temporal grounding agent found four segments where the yellow sponge is flying and being propelled by different means. The question asks what it used to "fly into the sky", implying the initial means of ascent. I need to check each segment to confirm if it's the "initial" use. I will start by examining the earliest segment provided by the temporal grounding agent that mentions flight. The earliest relevant segment is (190, 200). I will ask a specific question about what is used for flying in that segment, including all the options to help the video reader. I will then proceed to other segments if needed.`</think>` `<video_reader>`190:200`</video_reader>``<video_reader_question>`What did the yellow sponge use to fly or ascend in this segment? A. Wings B. Spoon C. Boat D. Pot Lid E. Bamboo Dragonfly F. Shovel`</video_reader_question>`

794 **A:** The tool result for `<video_reader>`190:200`</video_reader>` is The description of the video is: In this segment from an animated show, a yellow sponge character is seen holding a spoon and a pot lid. The sponge character is standing on a wooden surface, and there are other characters in the background. The sponge character then uses the pot lid to fly or ascend into the air, while the other characters watch. The environment appears to be indoors, possibly on a ship or a boat, given the wooden planks and the presence of other characters in the background. The sponge character's expression is one of excitement and determination as it uses the pot lid to fly. The answer is: D. Pot Lid

795 **Q:** `<think>` The video reader for the segment 190:200 explicitly states that "The sponge character then uses the pot lid to fly or ascend into the air". This directly answers the question "What did the yellow sponge use to fly into the sky?". Option D, "Pot Lid", matches this description. This seems to be the correct answer. I should confirm if this is the *initial* use if the question implies that. However, since the question is about "What did the yellow sponge use to fly into the sky?", and the video reader clearly states the pot lid is used for flight, this seems to be the correct answer.`</think>` `<answer>`D`</answer>` ← **VideoExplorer**

Figure 8: Case study of a failure on MLVU. The planner correctly identifies the temporal interval, but the VLM fails due to semantic hallucination, leading to wrong answer.