

# Video-Browser: Towards Agentic Open-web Video Browsing

Anonymous ACL submission

## Abstract

The evolution of autonomous agents is redefining information seeking, transitioning from passive retrieval to proactive, open-ended web research. However, a significant modality gap remains in processing the web’s most dynamic and information-dense modality: video. In this paper, we first formalize the task of Agentic Video Browsing and introduce **Video-BrowseComp**, a benchmark evaluating open-ended agentic browsing tasks that enforce a mandatory dependency on videos. We observe that current paradigms struggle to reconcile the scale of open-ended video exploration with the need for fine-grained visual verification. Direct visual inference (e.g., RAG) maximizes perception but incurs prohibitive context costs, while text-centric summarization optimizes efficiency but often misses critical visual details required for accurate grounding. To address this, we propose **Video-Browser**, a novel agent leveraging *Pyramidal Perception*, filtering with cheap metadata and zooming in with expensive visual perception only when necessary. Experiments demonstrate that our approach achieves a 37.5% relative improvement while reducing token consumption by 58.3% compared to Direct visual inference, establishing a foundation for verifiable open-web video research. We open-source all codes, benchmark at <https://anonymous.4open.science/r/VideoBrowser>.

## 1 Introduction

The rapid evolution of Large Language Models (LLMs) has catalyzed a paradigm shift from static question-answering to autonomous agents capable of actively navigating the web to solve complex problems (Nakano et al., 2021b; He et al., 2024). This transition towards agentic web browsing has become a dominant trend in AI research. Pioneering benchmarks such as GAIA (Mialon et al., 2023), BrowseComp (Wei et al., 2025), and MM-BrowseComp (Li et al., 2025a) have established

standards for agents operating within textual and static multimodal environments. These works reflect an emerging consensus: real-world agents must proactively seek, retrieve, and reason over information rather than passively receiving it.

Despite these advancements, the current landscape of deep research agents faces a fundamental limitation: a modality blind spot. While the community has focused heavily on static web information like text and images (Jin et al., 2025; Geng et al., 2025), it has largely overlooked the most dynamic and information-dense modality on the web: video, which houses everything from product reviews that demonstrate dynamic usage to complex tutorial procedures. Current AI research, video benchmarks like (Fu et al., 2024; Zhou et al., 2025a; Li et al., 2024) and video agents like (Fan et al., 2024; Wang et al., 2024, 2025b), treat video primarily as a static file to be perceived rather than a dynamic source to be explored. They typically input a curated clip to a model and query its content without requiring external information retrieval. This setting fails to reflect the agentic nature of real-world research. Users do not simply ask video a question, they search for videos through problem-driven then summarize the answers.

We argue that the field should move towards Agentic Video Browsing. To address these problems, we present the following three contributions.

First, we formally define this task not as video QA, but as an iterative decision-making process. To operationalize this task, we introduce **Video-BrowseComp**, a challenging benchmark designed to evaluate open-ended agentic browsing tasks that enforce a mandatory dependency on videos. To systematically assess agentic limits, we stratify the benchmark into three difficulty levels (Level 1 to Level 3), scaling from explicit retrieval to complex, cross-source reasoning that demands aggregating fragmentary evidence across disparate videos.

Second, to overcome the limitations of current

paradigms, we propose **Video-Browser**, a novel agent leveraging *Pyramidal Perception*. Existing approaches face a dilemma: direct video inference (e.g. RAG) incurs prohibitive context costs, while text-centric summarization suffers from a modality gap that misses visual details. Our architecture resolves this by treating the web videos as a pyramid. Utilizing lightweight video metadatas to prune irrelevant search results and zooming in with high-fidelity perception only when necessary.

Third, comprehensive experiments validate the effectiveness of our approach. Results on Video-BrowseComp demonstrate that Video-Browser achieves a 37.5% relative improvement in accuracy while reducing token consumption by 58.3% compared to direct video inference baselines. This confirms that our method successfully balances the scalability required for open-web video research with precision needed for fine-grained verification.

## 2 Related Works

### 2.1 Video Understanding

The field of video understanding has witnessed rapid advancements with the emergence of Multimodal Large Language Models (MLLMs) (Comanici et al., 2025; OpenAI et al., 2024; Liu et al., 2023; Zhu et al., 2025). Video LLMs (Shu et al., 2025; Qin et al., 2025; Li et al., 2025d; Liu et al., 2025; Chen et al., 2024; Shen et al., 2024) and agents (Wang et al., 2024; Fan et al., 2024; Wang et al., 2025b) have demonstrated impressive capabilities in processing long-context visual inputs. Concurrently, benchmarks (Fu et al., 2024; Zhou et al., 2025a; Li et al., 2024; Wang et al., 2025a) have been established to rigorously evaluate these models on traditional video QA tasks.

However, these existing works predominantly operate under a paradigm of passive perception. In typical evaluation settings, a model is fed a curated video clip and queried about its internal content in a closed-world manner. This approach remains at the perception level, falling short of research-oriented investigation. In real-world scenarios, answering complex queries often requires more than just “watching” a single video (Peng et al., 2025; He et al., 2025; Jang et al., 2025; Fu et al., 2025). It necessitates combining internal video evidence with external knowledge or actively searching across multiple videos to triangulate facts. Current benchmarks’ inability to measure external verification and cross-source reasoning underscores

the need for agents with active search capabilities.

### 2.2 Web Browsing Agents

To address the limitations of passive models, the community has shifted focus towards autonomous browsing agents. This evolution began with text-based agents, where pioneers (Nakano et al., 2021a; Jin et al., 2025; Li et al., 2025c; Wu et al., 2025a; Li et al., 2025b) demonstrated that LLMs could effectively utilize search engine tools to solve complex text-based questions (Wei et al., 2025; Zhou et al., 2025b; Mialon et al., 2023; Phan et al., 2025; Chen et al., 2025). Recently, this paradigm has expanded into the multimodal domain. Models and benchmarks (Li et al., 2025a; Wu et al., 2025b; Jiang et al., 2024) have pushed the boundary by requiring agents to process static web information, including text and images, marking a significant step towards multimodal information seeking.

Despite these advancements, current agents face a dilemma in balancing context costs with perceptual granularity when confronting dynamic video content. First, regarding visual granularity, direct visual inference strategies (e.g., standard RAG) theoretically offer the highest recall by streaming raw video directly to the model. However, this approach incurs prohibitive context costs, causing a “context explosion” that renders it unscalable for open-ended video browsing tasks where agents must verify evidence across multiple long videos (Yuan et al., 2025; Zhang et al., 2025; Ren et al., 2025). Second, to mitigate this computational burden, existing frameworks typically resort to text-centric summarization or rely on sparse metadata (Jiang et al., 2024; Li et al., 2025e). While this significantly reduces the context load, it introduces a critical modality gap. By compressing dynamic video streams into static textual proxies, these methods lose the fine-grained visual details essential for deep research, leaving agents unable to verify complex visual queries that lack explicit textual descriptions.

## 3 Benchmark

### 3.1 Task Definition

We formulate the video browsing task as an open-ended information seeking problem constrained by multimodal evidence verification. Given a natural language query  $Q$ , an agent is provided access to the open web  $\mathcal{W}$ . The agent must autonomously plan a sequence of actions  $S = \{a_1, a_2, \dots, a_t\}$ ,

184 such as searching the web, retrieving specific video  
185 timestamps, or reading subtitles, to generate a final  
186 result.

### 187 3.2 Principles and Scope

188 To ensure Video-BrowseComp serves as a rigorous  
189 benchmark for agentic video research, we adhere  
190 to three foundational design principles:

191 **1) Mandatory Video Dependency.** A core ob-  
192 jective of this benchmark is to evaluate video per-  
193 ception and reasoning, not merely text-based web  
194 search. To prevent models from solving questions  
195 via textual shortcuts (e.g., retrieving a movie plot  
196 from Wikipedia without watching the video). **2)**  
197 **Hard-to-Find, Easy-to-Verify.** The questions are  
198 intentionally complex, often requiring multi-step  
199 reasoning, cross-referencing, or long-context scan-  
200 ning. However, the output format is constrained  
201 to short, objective strings, such as specific entity  
202 names, counts, or colors, allowing for automated  
203 and unambiguous verification against the ground  
204 truth. **3) Answer Uniqueness.** To guarantee evalu-  
205 ation reliability, we ensure that every question has a  
206 single answer. Questions are grounded in objective  
207 spatio-temporal attributes rather than subjective in-  
208 terpretations.

### 209 3.3 Difficulty Stratification

210 To systematically assess the limits of agentic capa-  
211 bilities, we stratify Video-BrowseComp into three  
212 distinct difficulty levels based on the complexity  
213 of the search space and the depth of reasoning re-  
214 quired. We provide detailed benchmark cases in  
215 Appendix A.1.

216 **Level 1: Explicit Retrieval.** Questions in this  
217 level provide explicit constraints regarding the  
218 search domain, such as specific publication dates,  
219 platform names, or title keywords. The challenge  
220 lies in instruction following to locate the target  
221 video and temporal grounding to extract visual de-  
222 tails. **Level 2: Implicit Retrieval.** Removing  
223 explicit search keywords, the target video is refer-  
224 enced only through indirect descriptions. Agents  
225 must first bridge the gap between visual cues and  
226 textual search queries, and then filter through mul-  
227 tiple candidate videos to find the one matching the  
228 event description. **Level 3: Multi-Source Rea-**  
229 **soning.** Representing the highest tier of difficulty,  
230 these questions cannot be answered by viewing a  
231 single video. They require cross-source reasoning  
232 where information retrieved from one source serves  
233 as a prerequisite for locating the next.

### 234 3.4 Dataset Construction

235 **Annotation Pipeline.** We adopt a reverse construc-  
236 tion strategy that begins with video exploration and  
237 culminates in complex query formulation. First,  
238 each annotator is assigned 2 video categories and  
239 encouraged to explore relevant topics through web  
240 browsing. After identifying candidate videos, an-  
241 notators watch the complete video and design an  
242 initial question-answer pair focused on a specific  
243 clip. These simple questions are then expanded  
244 into complex queries following two principles: *In-*  
245 *direct & Compositional Queries:* Questions are  
246 augmented with compositional conditions to ensure  
247 the target video cannot be easily retrieved through  
248 a simple search. Annotators use indirect descrip-  
249 tions to reference video content rather than explicit  
250 keywords. *Multi-Source Reasoning:* Level 1 and  
251 Level 2 samples can be derived from the above  
252 steps with single-video evidence. For Level 3 sam-  
253 ples, annotators first identify a set of related videos  
254 covering the same topic or event, then design ques-  
255 tions that require aggregating complementary in-  
256 formation across these videos, such as comparing  
257 statistics from different matches, verifying conflict-  
258 ing claims across multiple sources, or tracing an  
259 event’s progression through several recordings. Fi-  
260 nally, annotators document the complete evidence  
261 chain or authoritative source video URLs, ensuring  
262 full traceability of the ground-truth answers.

263 **Quality Control.** We implement a two-stage  
264 quality control protocol to ensure the quality of  
265 Video-BrowseComp. *1.Easy Question Filtering.*  
266 Although annotators are required to design ques-  
267 tions based on the visual content of videos, some  
268 samples may still be answerable through common-  
269 sense knowledge without watching the video. To  
270 identify and remove such cases, we use the Google  
271 text search api<sup>1</sup> to return 5 relevant results, input  
272 it to GPT-5 to filter out questions that can be  
273 answered correctly. *2.Answer Uniqueness Valida-*  
274 *tion.* This step ensures that each question has only  
275 one correct and unambiguous answer. We employ  
276 a cross-validation strategy among the annotators:  
277 each annotator answers questions annotated by oth-  
278 ers following the evidence chain, and we compare  
279 their responses against the original ground-truth  
280 answers. If an alternative answer satisfies all task  
281 constraints but differs from the original, the sample  
282 is considered ambiguous and discarded.

<sup>1</sup><https://serper.dev>

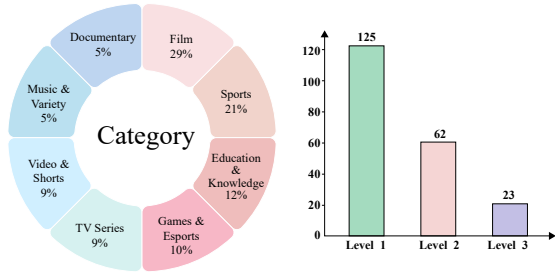


Figure 1: Statistical overview of our Benchmark, including the category (Left) and difficulty distribution (Right).

### 3.5 Dataset Statistics

The final benchmark consists of 210 validated questions spanning 8 diverse genres. As illustrated in Figure 1, the distribution prioritizes complex visual narratives, with Film (29%) and Sports (21%) forming the core, ensuring agents are evaluated on dynamic events rather than just static metadata. The difficulty is stratified into a pyramidal structure: while Level 1 (125 samples) tests explicit retrieval, Level 2 (62) and Level 3 (23) specifically target advanced capabilities in implicit retrieval and cross-source reasoning. We provide some analysis of the benchmark in Appendix A.5.

## 4 Method

### 4.1 Framework Formulation

We formalize Video-Browser as an agentic system  $\mathcal{S} = \langle \mathcal{P}, \mathcal{W}, \mathcal{A}, \mathcal{M} \rangle$ , comprising three specialized modules: Planner  $\mathcal{P}$ , Watcher  $\mathcal{W}$ , and Analyst  $\mathcal{A}$ , sharing a global memory  $\mathcal{M}$ . Given a user query  $Q$ , the system operates iteratively over time steps  $t$  to navigate the open video corpus  $\mathcal{V}_{web}$  and generate an answer  $A$ . The overview of the framework is shown in Figure 2.

**Memory  $\mathcal{M}$ .** Memory stores the interaction history, denoted as  $H_t = \{Q, (q_1, o_1), \dots, (q_{t-1}, o_{t-1})\}$ , where  $q_i$  represents search queries and  $o_i$  represents the structured observations returned by the Watcher.

**Planner  $\mathcal{P}$ .** Acting as the cognitive controller, the Planner orchestrates the process as follows: Upon receiving the query  $Q$ , the Planner first decomposes it into sequential sub-tasks:

$$\mathcal{S}_{plan} = \{s_1, s_2, \dots, s_m\} = \mathcal{P}(Q) \quad (1)$$

Then at each time step  $t$ , the Planner focuses on the current sub-task  $s_t$  and generates specific search keywords  $q_t$  tailored to this sub-task based on the

interaction history  $H_t$ :

$$q_t = \mathcal{P}(s_t, H_t) \quad (2)$$

These search queries  $q_t$  are then used to retrieve candidate videos  $\mathcal{V}_{ret}$  from the open video corpus. Once all sub-tasks in  $\mathcal{S}_{plan}$  are resolved or the maximum steps  $\mathcal{T}_{max}$  are reached, the Planner triggers the Analyst.

**Watcher  $\mathcal{W}$ .** Watcher addresses the Perception-Context Trade-off by functioning as a high-efficiency filter and localizer. It takes the raw candidate videos  $\mathcal{V}_{ret}$  retrieved by the search engine and the current query context  $q_t$  as input to produce a focused evidence set  $E_t$ . It employs a Pyramidal Perception mechanism  $\Psi$ :

$$E_t, o_t = \mathcal{W}(\mathcal{V}_{ret}, q_t; \Psi) \quad (3)$$

Here,  $E_t = \{(v_i, [t_{start}, t_{end}])\}$  denotes the visual evidence set consisting of prioritized temporal windows. Simultaneously, the Watcher generates a structured observation  $o_t$  (containing metadata and relevance summaries of  $E_t$ ) to provide feedback to the Planner, updating  $H_{t+1} = H_t \cup \{(q_{t+1}, o_{t+1})\}$ . **Analyst  $\mathcal{A}$ .** Upon termination at step  $T$ , the Analyst performs fine-grained reasoning. It takes the accumulated evidence  $E_{total} = \bigcup_{t=1}^T E_t$  to synthesize the final answer:

$$A = \mathcal{A}(Q, E_{total}) \quad (4)$$

**Iterative Reasoning Loop.** The system operates through continuous feedback between modules. At each iteration  $t$ , the Planner evaluates the Watcher’s observations  $o_t$  to assess evidence quality. If the collected evidence is insufficient or ambiguous for the current sub-task  $s_t$ , the Planner adaptively refines its search strategy—either generating new queries to explore alternative perspectives or adjusting keywords for improved precision. This reasoning loop continues until all sub-tasks in  $\mathcal{S}_{plan}$  are adequately resolved or the maximum iteration limit  $\mathcal{T}_{max}$  is reached, whereupon the Planner triggers the Analyst for final answer synthesis.

### 4.2 Pyramidal Perception

Open-ended research requires retrieving and synthesizing information from multiple video sources. To address the inherent scale of open-ended video browsing, the Watcher employs *Pyramidal Perception*.

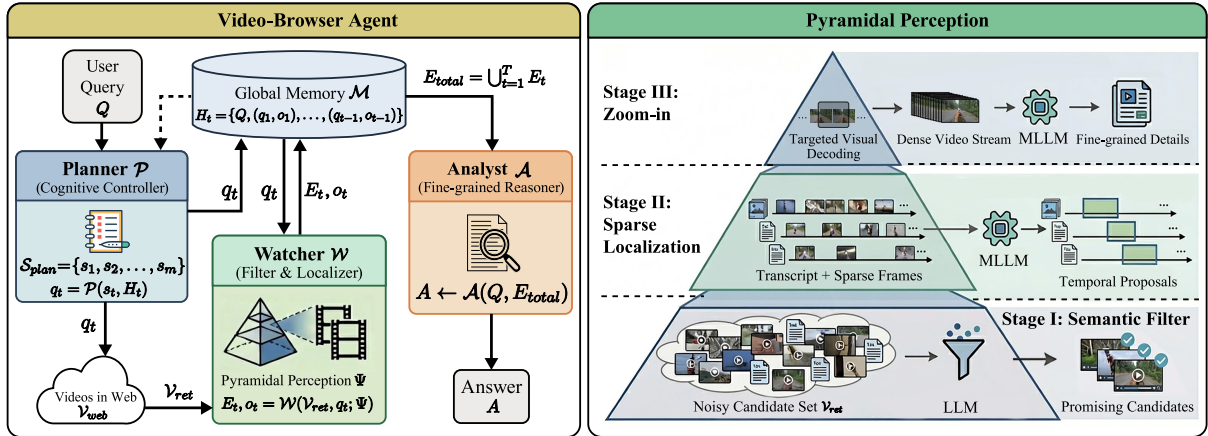


Figure 2: The overview of Video-Browser framework (Left), which consists of Planner, Watcher and Analyst. The Watcher employs Pyramidal Perception mechanism (Right), effectively mitigating the perception-context trade-off.

**Stage I: Semantic Filter.** At the base of the pyramid, the agent confronts a noisy candidate set  $V_{\text{ret}}$  retrieved from the open web. Since performing visual inference on this raw distribution is computationally prohibitive, the Watcher functions as a metadata-driven filter. By prompting an LLM to analyze solely textual metadata (e.g., titles, snippets), it assesses the relevance of each candidate to the sub-task at zero visual inference cost. This step aggressively prunes the search space before data loading occurs, ensuring that expensive perceptual resources are reserved strictly for promising candidates.

**Stage II: Sparse Localization.** For the retained videos, the challenge shifts from selection to localization, finding where the relevant information resides without incurring the cost of full decoding. The Watcher employs a sparse scanning strategy that balances context with efficiency. It ingests the full textual transcript to grasp the narrative structure, while simultaneously sampling a fixed set of sparse frames to capture visual states. An MLLM leverages this combined audio-visual context to identify temporal proposals, specific time windows  $[t_{\text{start}}, t_{\text{end}}]$  that potentially contain the visual evidence. This process compresses hour-long videos into focused “glimpses”, providing the Planner with sufficient feedback to make informed routing decisions.

**Stage III: Zoom-in.** At the tip of the pyramid, the system performs Targeted Visual Decoding on the prioritized temporal windows. This stage represents the transition from scanning to scrutinizing. The Analyst re-decodes the video strictly within the identified windows at high frequency to con-

struct a dense visual context. By applying a MLLM to this high-fidelity stream, the agent can resolve fine-grained visual details that were invisible during the sparse stage. This ensures that the most expensive reasoning power is allocated exclusively to the most information-dense moments, effectively resolving the perception-context trade-off.

## 5 Experiments

In this section, we conduct comprehensive experimental studies to evaluate the effectiveness of Video-Browser on the Video-BrowseComp benchmark. We structure our analysis around four core research questions: **RQ1: (Performance)** How does the Video-Browser compare to the SOTA models? **RQ2: (Efficiency)** Does the *Pyramidal Perception* architecture effectively minimize token consumption while maintaining good visual perception? **RQ3: (Test-time scaling)** Can Video-Browser effectively leverage increased test-time compute to yield continuous performance gains? **RQ4: (Ablation)** What is the contribution of each architectural module to the overall system effectiveness? See each experiment setting in Appendix B.1.

### 5.1 Metrics

**Accuracy (OA).** We employ Overall Accuracy as the primary metric. Following the established evaluation protocol of BrowseComp (Wei et al., 2025), we utilize LLM as a judge<sup>2</sup> (Zheng et al., 2023) to assess the semantic equivalence between the model’s prediction and the ground truth, rather than relying on rigid string matching.

<sup>2</sup>Evaluation prompts are in Appendix A.4.

Model	OA (%)	Difficulty			CE (%)	Genre Accuracy (%)							
		L1	L2	L3		Doc	Edu	Film	Game	Music	Sport	TV	Vid
<i>Tool-Free Models</i>													
Qwen3-VL-8B-Think	7.14	12.00	0.00	0.00	52.49	10.00	0.00	13.11	0.00	<b>20.00</b>	0.00	5.26	16.67
Qwen3-VL-235B-Ins	13.33	22.40	0.00	0.00	77.64	0.00	12.00	19.67	13.64	0.00	2.22	21.05	27.78
GLM-4.6V	10.95	16.80	3.23	0.00	44.40	0.00	12.00	11.48	13.64	10.00	2.22	26.32	16.67
gpt-4o-2024	17.62	28.00	3.23	0.00	58.81	10.00	16.00	26.23	13.64	0.00	4.44	31.58	27.78
gpt-4o-mini	9.52	16.00	0.00	0.00	63.55	0.00	12.00	18.03	9.09	0.00	0.00	15.79	5.56
gpt-5-mini	15.71	26.40	0.00	0.00	37.47	0.00	28.00	22.95	13.64	10.00	0.00	21.05	22.22
gemini-2.5-flash	16.67	27.20	1.61	0.00	77.79	<b>20.00</b>	20.00	24.59	4.55	0.00	2.22	36.84	22.22
gemini-2.5-pro	19.52	31.20	3.23	0.00	79.18	10.00	28.00	29.51	13.64	0.00	2.22	36.84	22.22
<i>Search Models</i>													
gemini-2.5-flash (S)	20.95	32.80	4.84	0.00	35.98	0.00	36.00	29.51	4.55	10.00	11.11	36.84	16.67
gemini-2.5-pro (S)	23.81	<b>37.60</b>	4.84	0.00	31.45	<b>20.00</b>	32.00	<b>31.15</b>	9.09	10.00	8.89	<b>57.89</b>	16.67
gpt-5.1 (S)	15.24	21.60	6.45	4.35	<b>30.20</b>	0.00	8.00	21.31	18.18	10.00	8.89	21.05	22.22
o4-mini-deep-research	22.86	30.40	<b>12.90</b>	<b>8.70</b>	42.55	10.00	28.00	29.51	<b>27.27</b>	10.00	17.78	21.05	16.67
<i>Video Browser</i>													
VideoBrowser (Qwen3)	19.05	25.60	9.68	<b>8.70</b>	69.81	0.00	36.00	22.95	4.55	10.00	11.11	15.79	38.89
VideoBrowser (GPT-5.1)	<b>26.19</b>	<b>37.60</b>	11.29	4.35	44.60	10.00	<b>40.00</b>	<b>31.15</b>	4.55	<b>20.00</b>	<b>20.00</b>	26.32	<b>44.44</b>

Table 1: Comprehensive results with a heatmap on Genre Accuracy. **Bold** indicates best performance. (S) means model has web search ability.

**Calibration Error (CE).** To quantify the reliability of the agent’s uncertainty estimation, we report the Calibration Error (CE). This metric measures the alignment between the model’s self-assigned confidence scores and its actual empirical accuracy, where a lower CE indicates a less overconfident and more trustworthy agent. Detailed formulations for CE are provided in Appendix A.3.

## 5.2 Baselines

To systematically evaluate the agentic capabilities for video-based research, we categorize our baselines into two groups: Tool-Free Models and Search-Augmented Models.

**Tool-Free Models.** This category evaluates the capability of state-of-the-art Multimodal Large Language Models (MLLMs), including Qwen3-VL-8B-Thinking, Qwen3-VL-235B-A22B-Instruct (Yang et al., 2025), GLM-4.6V (Team et al., 2025), gpt-4o-2024-11-20, gpt-4o-mini-2025-0807 (OpenAI et al., 2024), gemini-2.5-flash-2025-06, gemini-2.5-pro-2025-06 (Comanici et al., 2025), to answer questions relying solely on their internal parametric knowledge and the provided context.

**Search-Augmented Models.** This category evaluates the official, search-augmented model services provided by major vendors, including gemini-2.5-flash-2025-06 (w/ Search), gemini-2.5-pro-2025-06 (w/ Search), gpt-5.1-2025-11-13 (w/ Search) (OpenAI, 2025b), o4-mini-deep-research-2025-06-26 (OpenAI, 2025a). These systems possess integrated web search capabilities that enable them to

query the internet for up-to-date information autonomously during inference.

## 5.3 Performance

**Accuracy.** As shown in Table 1, Tool-Free models struggle significantly, plateauing at 19.52% (Gemini-2.5-pro) due to reliance on parametric memory. Integrating web search yields immediate gains in explicit retrieval tasks (Level 1), but performance collapses in Level 2 and Level 3 tasks, exposing that search agents can locate metadata but cannot verify visual details. Even the advanced o4-mini-res, is capped at 22.86% overall due to this visual blindness. Qualitative analysis of SOTA models is in Appendix A.6.

By equipping the agent with *Pyramidal Perception*, Video-Browser (GPT-5.1) achieves a state-of-the-art accuracy of 26.19%. This represents a decisive 71.8% relative improvement over the standard Search-Augmented GPT-5.1 (15.24%). Notably, our approach outperforms the Deep Research baseline (o4-mini) without requiring expensive reinforcement learning, demonstrating that complex planning cannot fully compensate for the lack of direct perception, and that visual grounding serves as a prerequisite for effective reasoning in video tasks.

**Calibration Error.** Tool-Free models exhibit severe overconfidence, with generally high Calibration Errors stemming from parametric hallucination. Search augmentation acts as a critical grounding mechanism, yielding a substantial reduction in CE across all baselines. Our Video-Browser main-

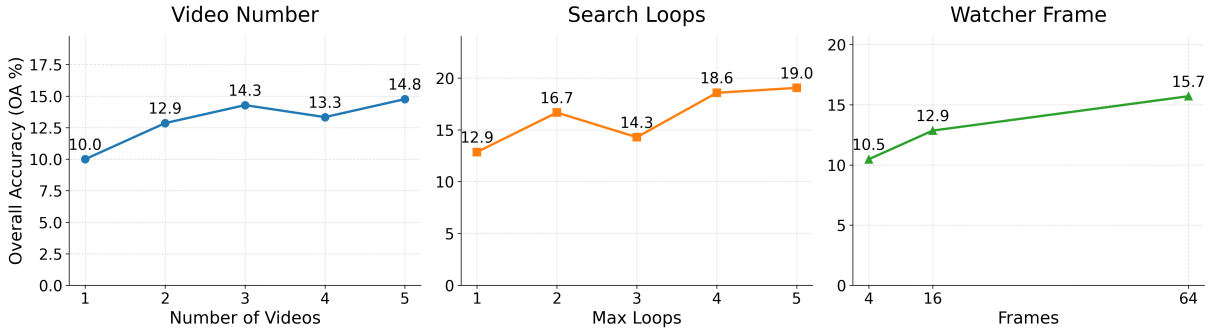


Figure 3: Test-time Scaling comparison including Search Breadth, Reasoning Depth and Perceptual Granularity.

tains a comparable calibration profile to search-augmented models, effectively balancing confidence with capability even while attempting significantly more challenging visual queries.

#### 5.4 Efficiency

In Table 2, we report both *Tokens* (the cumulative cost of scanning, filtering, and reasoning) and *Context* (the final prompt size fed to the Analyst).

Model	OA	CE	Tokens	Context
<i>Direct Visual Inference</i>				
w/ Qwen3	13.81	70.83	55,047	53,764
w/ GPT-5.1	19.05	50.80	78,229	76,955
<i>Summarization</i>				
w/ Qwen3	9.52	82.10	55,971	<b>1,260</b>
w/ GPT-5.1	18.57	51.34	75,728	<b>919</b>
<i>Pyramidal Perception</i>				
w/ Qwen3	<b>14.29</b>	<b>65.74</b>	<b>26,525</b>	7,678
w/ GPT-5.1	<b>26.19</b>	<b>38.10</b>	<b>32,627</b>	11,198

Table 2: Comparison of Three Paradigms on Video BrowseComp. **Bold** denotes the best result.

**Direct Visual Inference.** Input raw video frames directly to the model yields a baseline accuracy of 19.05% (GPT-5.1). However, this approach is prohibitively expensive, consuming 78,229 total tokens and filling the context window with 76,955 tokens. This confirms the context explosion bottleneck: the agent is forced to watch every frame, leaving little room for multi-step reasoning and limiting scalability.

**Summarization.** The Summarization paradigm effectively compresses the video input to text evidence, reducing the context load to a negligible 919 tokens. However, this reveals two critical flaws: 1). High Pre-processing Cost: The total tokens remain high (75,728) because the model must still process the full video to generate the summary. 2). Lossy Compression: Accuracy drops to 18.57%, and Calibration Error spikes (51.34%). Validating

the modality gap: textual proxies fail to capture the fine-grained visual details.

**Pyramidal Perception.** Our Pyramidal Perception reduces the Total Tokens to 32,627, a 58.3% (GPT-5.1) reduction compared to Direct Visual Inference. Crucially, unlike Summarization, this efficiency does not compromise perception. We achieve the highest accuracy of 26.19%. This demonstrates that our architecture successfully identifies where to spend compute, minimizing consumption while maximizing visual perception.

#### 5.5 Test-time Scaling

Figure 3 visualizes the performance of Video-Browser under varying computational budgets. We isolate three critical scaling dimensions: search breadth, reasoning depth, and perceptual granularity.

**Search Breadth (Video Number).** As the number of processed candidate videos increases from 1 to 5, accuracy improves from 10.0% to 14.8%. This suggests that for open-ended queries, by aggregating evidence from a broader set of videos, the agent can mitigate single-source hallucinations and improves recall.

**Reasoning Depth (Search Loops).** Increasing the maximum search loops allows the agent to iteratively refine its queries and explore new paths. We observe a strong positive correlation, with accuracy peaking at 19.0% with 5 loops. This confirms that the iterative feedback loop is a driver of complex reasoning, rather than a redundant overhead.

**Perceptual Granularity (Watcher Frames).** Increasing the number of scanning frames (from 4 to 64) leads to a consistent performance gain (10.5%  $\rightarrow$  15.7%). This indicates that while the Watcher operates on sparse signals to save tokens, increasing its sampling density significantly enhances its ability to localize correct temporal windows.

**Question:** In a film starring and directed by Ben Stiller, the ordinary photo technician named Walter Mitty (Walter Mitty) has a pen that runs throughout the entire film in his shirt pocket; the pen's cap appearing hints that Walter's inner longing is for a passionate life. What color is this pen?

**Ground Truth:** red

**Direct Visual Inference:**

Across the frames from the analyzed clips, Walter Mitty's shirt pocket and chest area are visible multiple times, **but no pen or pen cap can be seen** in any of them. The transcripts likewise never mention the pen or its color. Therefore, the videos provided do not contain the information needed to determine the color of Walter Mitty's pen.

**Answer:** The color of **the pen cannot be determined** from the provided videos.

**Summarization:**

None of the provided video summaries mention or show Walter Mitty's shirt pocket or the pen, **nor do they discuss its color or symbolism**. Therefore, the color of the pen cannot be determined from these sources.

**Answer:** The color of the pen **is not specified in the provided videos**.

**Pyramidal Perception:**

In the early close-up frames of Walter Mitty, his shirt pocket area **is visible and the only pen present shows a distinctly red cap** protruding from the pocket. This recurring red-capped pen is used symbolically in the film to hint at his inner desire for a more passionate life.

**Answer:** Red

Figure 4: Qualitative comparison of three paradigms.

Setting	OA	CE	Token	Context
<i>Component</i>				
w/o stage 1	9.52	66.64	16,846	6,068
w/o stage 2	12.38	69.50	44,961	42,401
w/o stage 3	10.48	67.52	18,330	4,198
<i>Modality</i>				
w/ Transcript	7.14	52.49	15,606	5,155
w/ Visual	12.38	37.31	8,353	3,901
Full	12.86	68.23	19,505	5,649

Table 3: Detailed analysis of Ablation and Modality.

## 5.6 Ablation Study

Table 3 dissects the contribution of each architectural component and modality.

**Component.** Stage 1 serves as the foundation for relevance. Removing it yields the lowest performance among all component variants (9.52%), significantly below the full model (12.86%). This drop indicates that without effective initial filtering, the agent wastes its perceptual budget on irrelevant content. Removing Stage 2 causes a catastrophic spike in computational cost: total token usage more than doubles (19,505  $\rightarrow$  44,961) and the final context load explodes by  $\sim 7.5\times$  (5,649  $\rightarrow$  42,401). This confirms that Stage 2 is essential for preventing context explosion. Removing Stage 3 degrades accuracy to 10.48%, highlighting that sparse sampling alone lacks the spatial-temporal resolution to verify fine-grained details. In summary, the architecture follows a logical progression: Stage 1 ensures relevance, Stage 2 ensures efficiency, and Stage 3 ensures precision.

**Modality.** The *w/ Transcript* setting (only input the video transcript) yields the lowest accuracy of 7.14%. This 44% relative performance drop compared to the full model (12.86%) serves as definitive empirical evidence of the modality gap. Explicit visual perception is mandatory for high-fidelity grounding.

## 5.7 Qualitative comparison

To intuitively understand the advantages of our approach, we present a qualitative case study in Figure 4. More case studies are in the Appendix B.3.

## 5.8 Reproducibility

We provide comprehensive implementation details in Appendix B.1 & B.2. Including the prompts for each module, and configurations etc.

## 6 Conclusion

We introduced **Video-BrowseComp**, a benchmark enforcing mandatory video dependency to rigorously evaluate agentic browsing. To overcome the limitations of existing paradigms, specifically the modality gap in textual summarization and the prohibitive costs of direct retrieval. We proposed **Video-Browser**. Leveraging *Pyramidal Perception*, our agent hierarchically filters content to resolve the perception-context dilemma. Experiments demonstrate that Video-Browser achieves a 37.5% accuracy improvement while reducing token consumption by 58.3% compared to direct visual inference, establishing a solid foundation for verifiable and efficient video agents.

## 610 Limitations

### 611 **Benchmark Scale and Computational Accessi-**

612 **bility.** We acknowledge that Video-BrowseComp  
613 operates at a modest scale (210 samples) compared  
614 to traditional, single-turn QA datasets. This scale  
615 is a deliberate design choice driven by the trade-off  
616 between annotation rigor and computational acces-  
617 sibility. First, unlike scalable synthetic generation,  
618 Video-Browsecomp necessitates expensive expert  
619 validation to guarantee answer uniqueness in an  
620 open-ended search space. Second, and more criti-  
621 cally, agentic video browsing is computationally  
622 intensive. Unlike standard VQA, a single query  
623 involves multi-step reasoning loops, web search  
624 retrieval, and high-frequency video decoding, often  
625 incurring substantial token costs. A massive-scale  
626 benchmark would impose a prohibitive computa-  
627 tional barrier for the broader research community.  
628 By prioritizing sample difficulty and quality over  
629 magnitude, we position Video-BrowseComp as a  
630 high-fidelity “Golden Test Set” that enables rigor-  
631 ous yet affordable comparisons of expensive agen-  
632 tic pipelines.

633 **Use of AI Assistants** We utilized Large Lan-  
634 guage Models (e.g., GPT-5.1) in two capacities:  
635 1) as the foundational backbone for the proposed  
636 Video-Browser agent, and 2) to assist with minor  
637 text polishing during the drafting process. All scien-  
638 tific claims and experimental results were verified  
639 by human authors.

## 640 Ethical Considerations

641 This work studies agents that browse and reason  
642 over open-web videos. To respect copyright and  
643 platform policies, we do not redistribute or host  
644 video content; any benchmark release should pro-  
645 vide only question–answer pairs and pointers (e.g.,  
646 URLs/timestamps) needed to locate publicly avail-  
647 able sources, and we will honor takedown requests  
648 and removal of problematic items when notified.  
649 Because videos may contain personal data (faces,  
650 names, incidental background information), we  
651 avoid collecting non-public sources and do not de-  
652 sign tasks that require identifying private individu-  
653 als; the benchmark is intended for factual, publicly  
654 verifiable information seeking rather than surveil-  
655 lance or profiling.

## References 656

- 657 Yukang Chen, Fuzhao Xue, Dacheng Li, Qinghao Hu,  
658 Ligeng Zhu, Xiuyu Li, Yunhao Fang, Haotian Tang,  
659 Shang Yang, Zhijian Liu, Ethan He, Hongxu Yin,  
660 Pavlo Molchanov, Jan Kautz, Linxi Fan, Yuke Zhu,  
661 Yao Lu, and Song Han. 2024. [Longvila: Scaling  
662 long-context visual language models for long videos.](#)  
663 *Preprint*, arXiv:2408.10188.
- 664 Zijian Chen, Xueguang Ma, Shengyao Zhuang, Ping  
665 Nie, Kai Zou, Andrew Liu, Joshua Green, Kshama  
666 Patel, Ruoxi Meng, Mingyi Su, Sahel Shari-  
667 fymoghaddam, Yanxi Li, Haoran Hong, Xinyu  
668 Shi, Xuye Liu, Nandan Thakur, Crystina Zhang,  
669 Luyu Gao, Wenhu Chen, and Jimmy Lin. 2025.  
670 [Browsecomp-plus: A more fair and transparent eval-  
671 uation benchmark of deep-research agent.](#) *Preprint*,  
672 arXiv:2508.06600.
- 673 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann,  
674 Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Mar-  
675 cel Blistein, and etc. Ori Ram. 2025. [Gemini 2.5:  
676 Pushing the frontier with advanced reasoning, multi-  
677 modality, long context, and next generation agentic  
678 capabilities.](#) *Preprint*, arXiv:2507.06261.
- 679 Yue Fan, Xiaojian Ma, Rujie Wu, Yuntao Du, Jiaqi  
680 Li, Zhi Gao, and Qing Li. 2024. [Videoagent: A  
681 memory-augmented multimodal agent for video un-  
682 derstanding.](#) *Preprint*, arXiv:2403.11481.
- 683 Chaoyou Fu, Yuhan Dai, Yondong Luo, Lei Li, Shuhuai  
684 Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yun-  
685 hang Shen, Mengdan Zhang, and 1 others. 2024.  
686 [Video-mme: The first-ever comprehensive evaluation  
687 benchmark of multi-modal llms in video analysis.](#)  
688 *arXiv preprint arXiv:2405.21075*.
- 689 Mingyang Fu, Yuyang Peng, Dongping Chen, Ze-  
690 tong Zhou, Benlin Liu, Yao Wan, Zhou Zhao,  
691 Philip S. Yu, and Ranjay Krishna. 2025. [Seeking  
692 and updating with live visual knowledge.](#) *Preprint*,  
693 arXiv:2504.05288.
- 694 Xinyu Geng, Peng Xia, Zhen Zhang, Xinyu Wang, Qi-  
695 uchen Wang, Ruixue Ding, Chenxi Wang, Jialong  
696 Wu, Yida Zhao, Kuan Li, Yong Jiang, Pengjun Xie,  
697 Fei Huang, and Jingren Zhou. 2025. [Webwatcher:  
698 Breaking new frontier of vision-language deep re-  
699 search agent.](#) *Preprint*, arXiv:2508.05748.
- 700 Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu,  
701 Yong Dai, Hongming Zhang, Zhenzhong Lan, and  
702 Dong Yu. 2024. [Webvoyager: Building an end-to-  
703 end web agent with large multimodal models.](#) *arXiv  
704 preprint arXiv:2401.13919*.
- 705 Zhihao He, Tianyao He, Yun Xu, Tiejuan Chen, Huabin  
706 Liu, Chaofan Gan, Zuxuan Wu, and Weiyao Lin.  
707 2025. [Enhancing video large language models  
708 with structured multi-video collaborative reasoning.](#)  
709 *Preprint*, arXiv:2509.13161.
- 710 Lawrence Keunho Jang, Yinheng Li, Dan Zhao, Charles  
711 Ding, Justin Lin, Paul Pu Liang, Rogerio Bonatti,

712	and Kazuhito Koishida. 2025. <a href="#">Videowebarena: Evaluating long context multimodal agents with video understanding web tasks</a> . In <i>The Thirteenth International Conference on Learning Representations</i> .	
713		
714		
715		
716	Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanmin Wu, Jiayi Lei, Pengshuo Qiu, Pan Lu, Zehui Chen, Chaoyou Fu, Guanglu Song, Peng Gao, Yu Liu, Chunyuan Li, and Hongsheng Li. 2024. <a href="#">Mmsearch: Benchmarking the potential of large models as multimodal search engines</a> . <i>Preprint</i> , arXiv:2409.12959.	
717		
718		
719		
720		
721		
722	Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. <a href="#">Search-r1: Training llms to reason and leverage search engines with reinforcement learning</a> . <i>Preprint</i> , arXiv:2503.09516.	
723		
724		
725		
726		
727	Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. 2024. <a href="#">Mvbench: A comprehensive multi-modal video understanding benchmark</a> . In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)</i> , pages 22195–22206.	
728		
729		
730		
731		
732		
733		
734	Shilong Li, Xingyuan Bu, Wenjie Wang, Jiaheng Liu, Jun Dong, Haoyang He, Hao Lu, Haozhe Zhang, Chenchen Jing, Zhen Li, Chuanhao Li, Jiayi Tian, Chenchen Zhang, Tianhao Peng, Yancheng He, Jiahao Gu, Yuanxing Zhang, Jian Yang, Ge Zhang, and 5 others. 2025a. <a href="#">Mm-browsecomp: A comprehensive benchmark for multimodal browsing agents</a> . <i>Preprint</i> , arXiv:2508.13186.	
735		
736		
737		
738		
739		
740		
741		
742	Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. 2025b. <a href="#">Search-o1: Agentic search-enhanced large reasoning models</a> .	
743		
744		
745		
746	Xiaoxi Li, Jiajie Jin, Guanting Dong, Hongjin Qian, Yongkang Wu, Ji-Rong Wen, Yutao Zhu, and Zhicheng Dou. 2025c. <a href="#">Webthinker: Empowering large reasoning models with deep research capability</a> . <i>arXiv preprint arXiv:2504.21776</i> .	
747		
748		
749		
750		
751	Xinhao Li, Yi Wang, Jiashuo Yu, Xiangyu Zeng, Yuhan Zhu, Haiyan Huang, Jianfei Gao, Kunchang Li, Yinan He, Chenting Wang, Yu Qiao, Yali Wang, and Limin Wang. 2025d. <a href="#">Videochat-flash: Hierarchical compression for long-context video modeling</a> . <i>Preprint</i> , arXiv:2501.00574.	
752		
753		
754		
755		
756		
757	Zijian Li, Xin Guan, Bo Zhang, Shen Huang, Houquan Zhou, Shaopeng Lai, Ming Yan, Yong Jiang, Pengjun Xie, Fei Huang, Jun Zhang, and Jingren Zhou. 2025e. <a href="#">Webweaver: Structuring web-scale evidence with dynamic outlines for open-ended deep research</a> . <i>Preprint</i> , arXiv:2509.13312.	
758		
759		
760		
761		
762		
763	Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. <a href="#">Visual instruction tuning</a> . <i>Preprint</i> , arXiv:2304.08485.	
764		
765		
766	Zhijian Liu, Ligeng Zhu, Baifeng Shi, Zhuoyang Zhang, Yuming Lou, Shang Yang, Haocheng Xi, Shiyi Cao,	
767		
	Yuxian Gu, Dacheng Li, Xiuyu Li, Haotian Tang, Yunhao Fang, Yukang Chen, Cheng-Yu Hsieh, De-An Huang, An-Chieh Cheng, Jinyi Hu, Sifei Liu, and 6 others. 2025. <a href="#">Nvlla: Efficient frontier visual language models</a> . In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)</i> , pages 4122–4134.	768 769 770 771 772 773 774
	Grégoire Mialon, Clémentine Fourier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2023. <a href="#">Gaia: a benchmark for general ai assistants</a> . <i>Preprint</i> , arXiv:2311.12983.	775 776 777 778
	Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2021a. <a href="#">Webgpt: Browser-assisted question-answering with human feedback</a> . <i>CoRR</i> , abs/2112.09332.	779 780 781 782 783 784 785 786
	Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, and 1 others. 2021b. <a href="#">Webgpt: Browser-assisted question-answering with human feedback</a> . <i>arXiv preprint arXiv:2112.09332</i> .	787 788 789 790 791 792
	OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, and etc. Alex Beutel. 2024. <a href="#">Gpt-4o system card</a> . <i>Preprint</i> , arXiv:2410.21276.	793 794 795 796 797 798
	OpenAI. 2025a. <a href="#">Deep Research System Card</a> .	799
	OpenAI. 2025b. <a href="#">GPT-5.1 Instant and GPT-5.1 Thinking System Card Addendum</a> .	800 801
	Tianhao Peng, Haochen Wang, Yuanxing Zhang, Zekun Wang, Zili Wang, Gavin Chang, Jian Yang, Shihao Li, Yanghai Wang, Xintao Wang, Houyi Li, Wei Ji, Pengfei Wan, Steven Huang, Zhaoxiang Zhang, and Jiaheng Liu. 2025. <a href="#">Mvu-eval: Towards multi-video understanding evaluation for multimodal llms</a> . <i>Preprint</i> , arXiv:2511.07250.	802 803 804 805 806 807 808
	Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Sean Shi, Michael Choi, Anish Agrawal, Arnav Chopra, Adam Khoja, and etc Richard Ren. 2025. <a href="#">Humanity’s last exam</a> .	809 810 811 812
	Minghao Qin, Xiangrui Liu, Zhengyang Liang, Yan Shu, Huaying Yuan, Juenjie Zhou, Shitao Xiao, Bo Zhao, and Zheng Liu. 2025. <a href="#">Video-xl-2: Towards very long-video understanding through task-aware kv sparsification</a> . <i>Preprint</i> , arXiv:2506.19225.	813 814 815 816 817
	Xubin Ren, Lingrui Xu, Long Xia, Shuaiqiang Wang, Dawei Yin, and Chao Huang. 2025. <a href="#">Videorag: Retrieval-augmented generation with extreme long-context videos</a> . <i>Preprint</i> , arXiv:2502.01549.	818 819 820 821

822	Xiaoqian Shen, Yunyang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, Chenchen Zhu, Zechun Liu, Fanyi Xiao, Balakrishnan Varadarajan, Florian Bordes, and 1 others. 2024. Longvu: Spatiotemporal adaptive compression for long video-language understanding. <i>arXiv preprint arXiv:2410.17434</i> .	878
823		879
824		880
825		881
826		882
827		
828	Yan Shu, Zheng Liu, Peitian Zhang, Minghao Qin, Junjie Zhou, Zhengyang Liang, Tiejun Huang, and Bo Zhao. 2025. Video-xl: Extra-long vision language model for hour-scale video understanding. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)</i> , pages 26160–26169.	883
829		884
830		885
831		886
832		
833		
834		
835	V Team, Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, Shuaiqi Duan, Weihan Wang, Yan Wang, Yean Cheng, Zehai He, Zhe Su, Zhen Yang, Ziyang Pan, and 69 others. 2025. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. <i>Preprint</i> , arXiv:2507.01006.	887
836		888
837		889
838		890
839		891
840		892
841		
842		
843	Weihan Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Xiaotao Gu, Shiyu Huang, Bin Xu, Yuxiao Dong, Ming Ding, and Jie Tang. 2025a. Lvbench: An extreme long video understanding benchmark. <i>Preprint</i> , arXiv:2406.08035.	893
844		894
845		895
846		896
847		897
848	Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. 2024. Videoagent: Long-form video understanding with large language model as agent. <i>Preprint</i> , arXiv:2403.10517.	898
849		899
850		
851		
852	Ziyang Wang, Shoubin Yu, Elias Stengel-Eskin, Jaehong Yoon, Feng Cheng, Gedas Bertasius, and Mohit Bansal. 2025b. Videotree: Adaptive tree-based video representation for llm reasoning on long videos. <i>Preprint</i> , arXiv:2405.19209.	900
853		901
854		902
855		903
856		904
857	Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. 2025. Browsecomp: A simple yet challenging benchmark for browsing agents. <i>Preprint</i> , arXiv:2504.12516.	905
858		
859		
860		
861		
862		
863	Jialong Wu, Baixuan Li, Runnan Fang, Wenbiao Yin, Liwen Zhang, Zhengwei Tao, Dingchu Zhang, Zekun Xi, Gang Fu, Yong Jiang, and 1 others. 2025a. Webdancer: Towards autonomous information seeking agency. <i>arXiv preprint arXiv:2505.22648</i> .	906
864		907
865		908
866		909
867		910
868	Jinming Wu, Zihao Deng, Wei Li, Yiding Liu, Bo You, Bo Li, Zejun Ma, and Ziwei Liu. 2025b. Mmsearchr1: Incentivizing llms to search.	911
869		
870		
871	An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. <i>Preprint</i> , arXiv:2505.09388.	
872		
873		
874		
875		
876		
877		

## Appendix

### Contents

- §A Video-BrowseComp (Benchmark Details)
  - §A.1 Benchmark Samples
  - §A.2 Annotators
  - §A.3 Evaluation Metrics Details
  - §A.4 Evaluation Prompts
  - §A.5 Benchmark Analysis
  - §A.6 Benchmark Case Studies
- §B Video-Browser (Agent Details)
  - §B.1 Experiment Settings
  - §B.2 Prompts
  - §B.3 Case Studies

## A Video-BrowseComp

### A.1 Benchmark Samples

We provide some benchmark examples in Figure 5 for reference.

### A.2 Annotators

To curate Video-BrowseComp, we recruited 8 student volunteers (holding Master’s or PhD degrees) with extensive experience in LLM usage, web search, and video understanding. The construction process follows a rigorous annotation pipeline with two-stage quality control in Sec 3.4.

Informed consent was obtained from all participants prior to the study. They were informed that the collected data would be used for research purposes and released publicly.

### A.3 Evaluation Metrics Details

#### A.3.1 LLM-based Accuracy Judge

Since the reference answers in Video-BrowseComp are designed to be short and verifiable (e.g., specific entities, colors, or counts), rigid string matching often fails due to minor formatting differences. We utilize gpt-5-mini-2025-08-07 as the automated judge. The judge is prompted to verify if the prediction conveys the same factual information as the ground truth. The specific prompt used for evaluation is shown in Figure 6.

#### A.3.2 Calibration Error (CE)

To calculate CE, we first prompt the model to provide a confidence score  $p \in [0, 1]$  alongside its final answer. We partition the predictions into  $B = 5$  equally spaced bins (i.e.,  $[0, 0.2), \dots, [0.8, 1.0]$ ). The CE is calculated as the weighted average of the absolute difference between the accuracy and the average confidence within each bin:

$$CE = \sum_{i=1}^B \frac{n_i}{N} |\text{acc}(i) - \text{conf}(i)| \quad (5)$$

Where  $N$  is the total number of samples,  $n_i$  is the number of samples in the  $i$ -th bin,  $\text{acc}(i)$  is the empirical accuracy of the samples in bin  $i$ , and  $\text{conf}(i)$  is the average predicted confidence of the samples in that bin.

### A.4 Evaluation Prompts

We provide the prompts used in evaluation and judge process at Figure 6.

### A.5 Benchmark Analysis

#### A.5.1 Category-Wise Performance

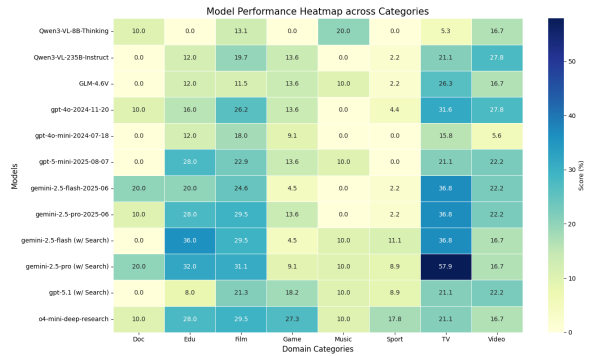


Figure 8: Performance heatmap.

Figure 8 provides a granular breakdown of model accuracy across the eight video genres. A clear performance dichotomy is observable based on the availability of external textual metadata. In categories like TV Series and Education, where transcripts, wikis, and plot summaries are abundant on the open web, search-augmented models achieve their highest scores. For example, Gemini-2.5-Pro (w/ Search) reaches 57.9% accuracy on TV Series. Conversely, performance collapses in dynamic, stochastic environments like Games and Sports. These categories require temporal grounding in specific visual moments (e.g., a specific foul or gameplay sequence) that are rarely indexed

by text search engines. Consequently, the same model’s accuracy drops to 9.1% in Games and 8.9% in Sports , highlighting the "Modality Gap" where agents fail to process visual information when text shortcuts are unavailable.

### A.5.2 Cost-Efficiency Analysis (Token Usage)

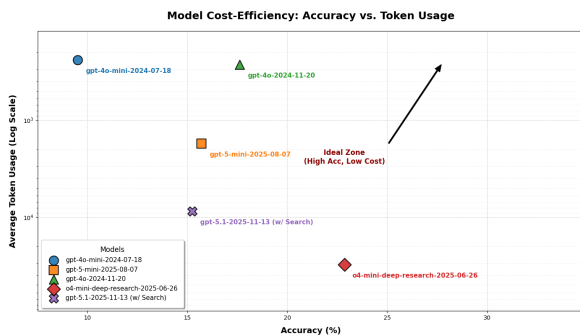


Figure 9: Token usage.

Figure 9 visualizes the trade-off between computational cost (measured in average token usage per query) and performance (accuracy). As indicated by the directional arrow, the top-right quadrant represents the ideal balance of High Accuracy and Low Token Usage. This is the target frontier for future agentic development. The plot reveals that current high-performing agents like o1-mini-deep-research achieve superior accuracy but at the expense of significantly higher token consumption, likely due to extensive iterative browsing and multi-step reasoning. In contrast, lightweight models like gpt-4o-mini are token-efficient but lack the deep reasoning capabilities required for complex video tasks. The arrow suggests that the goal of future research is to push these agents "upwards" towards the ideal zone, optimizing the reasoning process to be both accurate and token-efficient.

### A.6 Benchmark Case Studies

Our evaluation reveals a critical reliance on textual metadata. In dynamic scenarios like sports, where specific gameplay moments (e.g., a specific foul sequence or a buzzer-beater in overtime) are not indexed in search engine results, models fail to answer even when the visual evidence is clear. Consider the example of identifying two NBA teams based on a double-overtime sequence in Figure 10. The visual evidence explicitly shows the Houston Rockets and Oklahoma City Thunder. However, because this specific game narrative does not appear in textual game logs, both GPT-4o and Gemini-2.5-

Pro refuse to answer, stating that the information is impossible to determine.

Similarly, in Figure 11 involving technical fouls and baseline altercations , models like GPT-5.1 (w/ Search) default to Unknown. This demonstrates that current search-augmented models often function merely as text search, incapable of using temporal grounding to fill information gaps when the open web falls silent.

In Figure 12, the agent is asked to read the temperature on a digital display inside a space module. The visual ground truth is 85. However, o1-mini-deep-search retrieves an external news report stating the oven’s design specification is 190°C and confidently outputs that incorrect value. The model successfully retrieved relevant text but failed to verify it against the visual reality.

In Figure 13, a complex query requiring the identification of a comic villain (Loki), a related film (Iron Man), and a specific visual detail (a food stain), models struggle to maintain coherence. While GPT-5.1 correctly identifies the “Cheeseburger” by linking the text clues to the correct movie scenes, Gemini-2.5-Pro fails the reasoning chain. It incorrectly identifies the stain as an “Acai popsicle” by hallucinating a scene from Avengers: Infinity War rather than the target scene from Iron Man . This highlights the difficulty agents face in maintaining context across disparate video sources.

## B Video-Browser

### B.1 Experiment Settings

To ensure reproducibility, we detail the specific hyperparameter configurations used for the Video-Browser agent across different experimental sections. The settings for the Planner (search loops), Watcher (sparse sampling), and Analyst (dense sampling) are adjusted to verify specific capabilities as follows:

**Main Results (Table 1).** For the comprehensive performance evaluation reported in Table 1, we configure the agent to maximize reasoning accuracy. Search Strategy: We set the maximum search loops to  $T_{max} = 5$ . In each iteration, the top  $K = 3$  most relevant videos are retrieved and processed. Perception: The Watcher employs a sparse sampling rate of 16 frames per video for ROI localization. For the Analyst’s Zoom-in stage, we employ a dense sampling rate of 1 FPS within the identified temporal windows, with a maximum cap of 32 frames per window to balance context limits.

1061	<b>Efficiency Analysis (Table 2).</b>	To strictly evaluate the token efficiency of Pyramidal Perception against baselines, we standardize the interaction rounds. Search Strategy: We fix the search process to a single loop ( $T_{max} = 1$ ) with $K = 3$ retrieved videos. Baselines: For <i>Direct Visual Inference</i> and <i>Summarization</i> baselines, we adopt a uniform sampling strategy of 128 frames per video. Video-Browser (Ours): We use 16 sparse frames for the Watcher. For the Analyst, we maintain the dense sampling setting of 1 FPS with a maximum window size of 32 frames.	1108
1062			1109
1063			1110
1064			
1065			
1066			
1067			
1068			
1069			
1070			
1071			
1072			
1073	<b>Test-time Scaling Analysis (Figure 3).</b>	We investigate the scaling laws by varying one parameter while fixing others. The specific configurations for each scaling dimension are: Scaling Video Number: We fix the search loop to $T = 1$ and sparse sampling to 16 frames, while varying the number of retrieved videos. Scaling Search Loops: We fix the sparse sampling to 16 frames. To prevent context explosion during multi-step reasoning, we fix the number of retrieved videos per round to $K = 2$ , while varying the max search loops. Scaling Watcher Frames: We fix the video number to $K = 2$ and search loop to $T = 1$ . The dense sampling is kept at 1 FPS (max 32 frames), while we vary the sparse sampling frames in the Watcher.	1119
1074			1120
1075			1121
1076			1122
1077			1123
1078			1124
1079			1125
1080			1126
1081			1127
1082			1128
1083			1129
1084			1130
1085			1131
1086			1132
1087			1133
1088	<b>Ablation Experiments (Table 3).</b>	Search Strategy: We fix the search process to a single loop ( $T_{max} = 1$ ) with $K = 2$ retrieved videos. We use 16 sparse frames for the Watcher. For the Analyst, we maintain the dense sampling setting of 1 FPS with a maximum window size of 32 frames.	1134
1089			1135
1090			1136
1091			1137
1092			1138
1093			1139
1094			1140
1095	<b>B.2 Prompts</b>	We provide all prompts used in the Video-Browser.	1141
1096	<b>B.2.1 Direct Visual Perception Prompts</b>	Check at Figure 15.	1142
1097			1143
1098	<b>B.2.2 Text-centric Summarization Prompts</b>	Check at Figure 16 and Figure 17.	1144
1099			1145
1100	<b>B.2.3 Pyramidal Perception Prompts</b>	Check at Figure 18, Figure 19, Figure 20 and Figure 21.	1146
1101			1147
1102			1148
1103	<b>B.3 Case Studies</b>		1149
1104			1150
1105			1151
1106			1152
1107			1153
			1154
			1155
			1156
			1157
			1158
			1159
			1160
			1161
			1162
			1163
			1164
			1165
			1166
			1167
			1168
			1169
			1170
			1171
			1172
			1173
			1174
			1175
			1176
			1177
			1178
			1179
			1180
			1181
			1182
			1183
			1184
			1185
			1186
			1187
			1188
			1189
			1190
			1191
			1192
			1193
			1194
			1195
			1196
			1197
			1198
			1199
			1200
			1201
			1202
			1203
			1204
			1205
			1206
			1207
			1208
			1209
			1210
			1211
			1212
			1213
			1214
			1215
			1216
			1217
			1218
			1219
			1220
			1221
			1222
			1223
			1224
			1225
			1226
			1227
			1228
			1229
			1230
			1231
			1232
			1233
			1234
			1235
			1236
			1237
			1238
			1239
			1240
			1241
			1242
			1243
			1244
			1245
			1246
			1247
			1248
			1249
			1250
			1251
			1252
			1253
			1254
			1255
			1256
			1257
			1258
			1259
			1260
			1261
			1262
			1263
			1264
			1265
			1266
			1267
			1268
			1269
			1270
			1271
			1272
			1273
			1274
			1275
			1276
			1277
			1278
			1279
			1280
			1281
			1282
			1283
			1284
			1285
			1286
			1287
			1288
			1289
			1290
			1291
			1292
			1293
			1294
			1295
			1296
			1297
			1298
			1299
			1300
			1301
			1302
			1303
			1304
			1305
			1306
			1307
			1308
			1309
			1310
			1311
			1312
			1313
			1314
			1315
			1316
			1317
			1318
			1319
			1320
			1321
			1322
			1323
			1324
			1325
			1326
			1327
			1328
			1329
			1330
			1331
			1332
			1333
			1334
			1335
			1336
			1337
			1338
			1339
			1340
			1341
			1342
			1343
			1344
			1345
			1346
			1347
			1348
			1349
			1350
			1351
			1352
			1353
			1354
			1355
			1356
			1357
			1358
			1359
			1360
			1361
			1362
			1363
			1364
			1365
			1366
			1367
			1368
			1369
			1370
			1371
			1372
			1373
			1374
			1375
			1376
			1377
			1378
			1379
			1380
			1381
			1382
			1383
			1384
			1385
			1386
			1387
			1388
			1389
			1390
			1391
			1392
			1393
			1394
			1395
			1396
			1397
			1398
			1399
			1400

the noise of the full video context, concluding that no pen or pen cap can be seen”. **Video-Browser** succeeds by leveraging its Zoom-in mechanism. It identifies the relevant close-up shots of the character’s pocket and performs targeted decoding, allowing it to clearly distinguish the distinctly red cap” that serves as a symbolic plot device.

**Case 2: Complex Dynamic Reasoning.** In Figure 22, the user query describes a highly specific, multi-stage basketball sequence from the 2024 NBA play-offs” (alley-oop, block, trip, slam floor). **Direct Visual Inference** fails because it retrieves irrelevant videos (e.g., Giannis’s 2021 block or Bronny James clips) and cannot semantically match the complex visual sequence to the query within the raw pixel stream. **Summarization** also fails, as the textual summaries describe unrelated games (e.g., Wizards vs. Pelicans” or 2016 NBA Finals”), exhibiting a severe hallucination issue where the model claims no relevant sequence exists. In contrast, **Video-Browser** correctly filters the candidate videos and localizes the exact moment in the 2024 NBA Finals, Mavericks vs. Celtics.” By accurately interpreting the visual commentary and action (“Tatum step back, stopped, blocked”), it correctly identifies the player as Jayson Tatum.

### B.3.3 Failure Case Analysis

Despite the significant improvements, Video-Browser still exhibits limitations in specific complex scenarios. We analyze typical failure modes to provide insights for future improvements.

**1. Semantic Distractors (Fig. 25).** In open-ended browsing, multiple videos may share similar semantic attributes. In the Icelandic food case, the agent correctly identified the context (uncomfortable food in Iceland) but was misled by a distractor candidate, fermented shark (hákarl), which shares the attribute of being smelly and hard to eat. The agent failed to distinguish the specific target (“sheep’s butt”) from the semantically similar distractor, highlighting the need for stricter entity verification.

**2. Fine-grained Visual Hallucination (Fig. 26).** For extremely small objects, the model may hallucinate specific attributes. In the friendship test case, while the agent correctly located the scene and the object (a beverage can), it misidentified the brand as Dr Pepper instead of the ground truth “Coca-Cola.” This suggests that even with Zoom-in, current MLLMs still struggle with zero-shot OCR or logo recognition on low-resolution objects.

**3. Information Deficit & Entity Mismatch** The case 23 (Kyle Lowry) demonstrates an Information Deficit: the agent successfully retrieved the correct highlight clip but failed to answer because the specific statistic (final score) was neither spoken in the transcript nor explicitly displayed in the visual summary. The agent honestly reported not stated, failing to infer or retrieve external box scores. The case 24 shows an Action-to-Entity Mismatch. The agent observed the correct basketball plays (block, steal, save) but misattributed them to LeBron James instead of the target player (who played for the Lakers and Thunder), leading to a factually correct answer for the wrong entity.

1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220  
1221  
1222

Level 1



In an NBA Christmas promo video from a certain year, a lanky small forward nailed the first three-pointer that rang with a bell, and a player wearing jersey number 6 later dunked. Who was the passer to the player wearing number 6, and which NBA team did he play for at the time?



Los Angeles Lakers

- 1: Confirm that this promotional video is the 2013 NBA Christmas promo.
- 2: Confirm that the player wearing jersey number 6 who dunked is LeBron James.
- 3: Confirm that the passer to James was Steve Nash, who was playing for the Los Angeles Lakers at the time.



Level 2



The first Sequoia Chair Professor at a renowned comprehensive university taught a course on “Deep Learning for Computer Vision” at a world-top-ranked university in 2025. In one lecture devoted to “A core task in Computer Vision,” the professor introduced the concept of the semantic gap. When the slides asked, “What does the computer see?”, the answer given in the lecture was:



An image is a tensor of integers

- 1: Confirm the instructor is Fei-Fei Li
- 2: Confirm the institution is Stanford and the course is CS231N | Spring 2025
- 3: Identify the lecture as Lecture 2: Image Classification with Linear Classifiers
- 4: Locate the segment “What the computer sees”; the answer is “An image is a tensor of integers”



Level 3



An actor starred in a film adapted from a famous novel, playing the lead who goes to war to earn a better future for the woman he loves, makes money overseas at all costs, returns to find she has already married, yet still tries every means to see her and rekindle their romance, and ultimately dies alone. In another film the same actor plays an African smuggler, befriends a Black man, escapes alive, kills the colonel, and—mortally wounded—hands the crucial item to his Black friend before dying. What is it?



Diamond

- 1: Identify the first film as the 2013 version of “The Great Gatsby”
- 2: Confirm the actor is Leonardo DiCaprio as Gatsby
- 3: Identify the second film as “Blood Diamond”
- 4: Locate the scene where, dying, he gives the priceless diamond to Solomon to improve his family’s life; the answer is the diamond



Figure 5: Benchmark samples.

#### Evaluation Prompt

Your response should be in the following JSON format:

```
{  
  "Explanation": "your explanation for your final answer",  
  "Answer": "your succinct, final answer",  
  "Confidence": "your confidence score between 0% and 100% for your answer"  
}
```

#### Judge Prompt

Question: {question}

Ground Truth: {ground\_truth}

Model Prediction: {prediction}

Evaluate if the Prediction matches the Ground Truth.

Return JSON: {"is\_correct": true} or {"is\_correct": false}

Figure 6: Full prompts used for evaluation and judging.

**[User Query]**

A legendary power forward, after switching careers to become a commentator, once bet with his co-host on a popular American basketball analysis show that a No. 1 draft pick center from Asia could not score 19 points in a single game. Subsequently, the center proved himself in a game, forcing the commentator to fulfill the bet — kissing a donkey’s butt on a subsequent live broadcast. What was the center’s final score in that game?

**[Planner]**

Thought:

Decomposition:

- 1) Identify the legendary power forward turned commentator on a popular American basketball analysis show.
- 2) Identify the No. 1 draft pick center from Asia involved in the bet.
- 3) Find the specific game where he needed at least 19 points.
- 4) Determine his final point total in that game.

Gap Analysis: We have no prior videos or leads, so everything is missing. Strategy: Start broadly but targeted to the well-known anecdote about a commentator betting that an Asian No. 1 pick center wouldn’t score 19 points and then having to kiss a donkey’s butt on air. The key is to surface the exact game and box score. One precise query that includes all distinctive elements (bet, donkey, kiss, Asian No. 1 pick center, points) should quickly reveal the names, game, and final scoring line.

**[Planner]**

Queries: [’NBA bet kiss donkey butt Asian No. 1 pick center 19 points game what was his final score’]

**[Searcher]**

executing 1 queries: [’NBA bet kiss donkey butt Asian No. 1 pick center 19 points game what was his final score’] -> YouTube found 0 results for ’NBA bet kiss donkey butt Asian No. 1 pick center 19 points game what was his final score’

**[Searcher]**

Found 0 unique video candidates.

**[Watcher]**

Starting sparse sampling and window identification...

**[Planner]**

Step 1/3 -> Knowledge Accumulation: 0 videos in store, 1 queries tried. -> Round complete. Continuing search loop (strict mode)...

**[Planner]**

Thought: Decomposition:

1. Identify the personalities and show: legendary power forward turned commentator, co-host, popular American NBA analysis show, bet about an Asian No. 1 pick center (clearly Yao Ming) scoring at least 19 points, and the loser kissing a donkey’s butt on live TV.
2. Then find the specific game where Yao Ming ‘proved himself’ and exceeded 19 points, and extract his final point total in that game.

Figure 7: Workflow of the Video-Browser

## Workflow of the Video-Browser

Gap Analysis: - We have no video or textual confirmation yet; everything is missing. - The prior search was very general and mixed many keywords, which may dilute results. Strategy: - Directly target the well-known anecdote about Charles Barkley (legendary PF) and Yao Ming on TNT's Inside the NBA, and then focus on the game box score from the night that resolved the bet. The key is to find a video or segment that explicitly states Yao Ming's point total in that proving game. - Use a concise, entity-focused query with both names and mention of the bet and points.

### [Planner]

Queries: ['Charles Barkley bet Yao Ming 19 points game final score donkey Inside the NBA']

### [Searcher]

executing 1 queries: ['Charles Barkley bet Yao Ming 19 points game final score donkey Inside the NBA'] -> YouTube found 10 results for 'Charles Barkley bet Yao Ming 19 points game final score donkey Inside the NBA' [Searcher]

Found 10 unique video candidates.

### [Watcher]

Starting sparse sampling and window identification... -> Selecting top videos using LLM... -> LLM selected video indices: [0, 1, 3] -> Selected 3 videos for processing. -> Processing Video 1/3: Inside the NBA: If Yao Ming scores 19 points...

Cache hit for video file:

data/cache/videos/1ARR0XqanM.mp4 -> Extracted 16 frames with timestamps.

Cache hit for transcript with timestamps.....

Cache hit for video file:

data/cache/videos/SOK0mVRPfdk.mp4 -> Extracted 16 frames with timestamps.

Cache hit for transcript with timestamps.....

Cache hit for video file:

data/cache/videos/I5BqF2Dn5bk.mp4 -> Extracted 16 frames with timestamps.

Cache hit for transcript with timestamps.....

### [Planner]

Step 2/3 -> Knowledge Accumulation: 3 videos in store, 2 queries tried. -> Round complete. Continuing search loop (strict mode)...

### [Planner]

Thought: The user already has the key details confirmed: the show is Inside the NBA, the commentator is Charles Barkley, the center is Yao Ming, and multiple verified clips explicitly state that in the game triggering the bet payoff Yao Ming went 9-for-9 in L.A. and scored 20 points. This directly answers the user's question about the final score in that game. No remaining sub-question is unsolved, and any further search would be redundant verification of the same specific stat line.

## Workflow of the Video-Browser

### [Searcher]

No queries provided. Skipping.

### [Watcher]

Starting sparse sampling and window identification...

— Step: watcher —

### [Checker]

Step 3/3 -> Knowledge Accumulation: 3 videos in store, 2 queries tried. -> Max loops reached. Proceeding to Analyst.

— Step: checker —

### [Analyst]

Extracting relevant clips and synthesizing final answer... -> Processing Video 1: Inside the NBA: If Yao Ming scores 19 points... (3 windows)

Cache hit for video file:

data/cache/videos/1ARROXqanM.mp4 -> Extracting clip 7.0s - 22.0s (1 FPS)... -> Added ['[0.0s - 7.0s] Next Thursday at Yao Ming, the Rockets taking on Dallas and also Seattle and Portland.', '[7.0s - 11.0s] Will this be the week to have to kiss Kenny's ass?', '[11.0s - 13.0s] Why would you do that, Chuck. . . . . [51.0s - 54.0s] You say he's going to get 19 points in the game.", "[54.0s - 56.0s] I'm telling you, he has some offensive skills.", '[56.0s - 59.0s] Why would you go without eating until he gets 19 points?' . . . . . '[142.0s - 145.0s] Man, I hope he gets to 19 next week.', '[145.0s - 147.0s] Anyway, in our game.')] transcript lines. -> Processing Video 2: Charles Barkley kisses Kenny Smith's Donkey (1 windows)

Cache hit for video file:

data/cache/videos/SOK0mVRPfdk.mp4 -> Extracting clip 20.2s - 31.5s (1 FPS)... -> Added ['[19.2s - 20.2s] And', "[20.2s - 23.4s] So Sunday night y'all mean went nine for nine", '[24.4s - 26.4s] in LA', '[26.8s - 31.5s] 20 points and and everybody who heard those words was saying well', "[31.5s - 36.0s] I guess that means Charles are gonna have to pay up. I wonder how that's gonna happen"] transcript lines. -> Processing Video 3: Inside the NBA on TNT Throwback to when Charles Barkley kiss kenny ass if Yao Ming scores 19 points (1 windows)

Cache hit for video file:

data/cache/videos/I5BqF2Dn5bk.mp4 -> Extracting clip 172.0s - 188.0s (1 FPS)... -> Added ["[169.0s - 172.0s] I'm going to kiss his ass right here in front of everybody.", '[172.0s - 179.0s] And so Sunday night, Yao Ming went 9 for 9 in L.A.', "[179.0s - 180.0s] That's embarrassing.", '[180.0s - 184.0s] He scored 20 points, and everybody who heard those words was saying,', '[184.0s - 188.0s] well, I guess that means Charles is going to have to pay up.', "[188.0s - 189.0s] I wonder how that's going to happen." ] transcript lines. -> Invoking Analyst LLM with video context...

FINAL ANSWER: "Explanation": "The videos show Charles Barkley on 'Inside the NBA' betting that Yao Ming would not score 19 points in a game. In the recap clip following the game in Los Angeles, the host states that 'Sunday night Yao Ming went 9 for 9 in L.A. ... **He scored 20 points.**' This confirms Yao Ming's final point total in the game that triggered Barkley to fulfill the bet.", "Answer": "**20 points**", "Confidence": "100%"

**Level 2 Question:** In a 2025 NBA regular season game, the action was fiercely contested. An away-team player fouled out on his sixth personal foul, and a home-team player hit a clutch midrange jumper at the end of the fourth quarter to force overtime. In the final 8.8 seconds of overtime, an away-team player dunked to tie the score again, sending it to a second overtime. Finally, the away-team player's miss ended this roller-coaster affair. What are the two teams facing off in this game?  
**GROUND TRUTH:** Houston Rockets, Oklahoma City Thunder

**gpt-4o-2024-11-20:**  
**Explanation:** The question provides specific details about an intense NBA game from the 2025 regular season, including a foul out, clutch plays, and double overtime, **but it does not explicitly name the two teams involved**. Without additional context or direct evidence, **it is impossible** to accurately determine the teams based purely on the narrative presented.  
**Answer:** The two teams **cannot be determined from the provided information**.

**gemini-2.5-pro-2025-06:**  
**Explanation:** The provided text describes a hypothetical NBA game set in the future (2025). It does **not contain any specific information**, such as team names, player names, cities, or arenas, that would allow for the identification of the two teams involved. **The details are generic to any** basketball game, making it impossible to determine which specific teams were playing.  
**Answer:** It is **impossible to determine** the two teams from the information provided, as the description is a fictional scenario lacking any specific identifying details.

**Qwen3-vl-235B-A22B-instruct:**  
**Explanation:** The question describes a fictional or hypothetical NBA game scenario with specific details (fouling out on the sixth foul, clutch midrange jumper, dunk to tie in overtime, etc.), **but it does not provide any actual team names**, logos, locations, or player identities. **Without any identifying** information about the teams, it is impossible to determine which two teams are facing off. The scenario could apply to any two NBA teams in a hypothetical 2025 game. Therefore, the question cannot be answered with factual certainty.  
**Answer:** **Unknown**

Figure 10: Level 2 case, NBA teams.

**Level 2 Question:** In a 2025 NBA regular-season game, a visiting-team player on the bench angrily complained about an opponent grabbing the rim for rebounds and was assessed a technical foul by the referees. Later in the second quarter, he tangled at the baseline with a home-team player and they pushed each other, and the home-team player was also assessed a technical foul by the referees. Which player's scoring was the highest in this game?  
**GROUND TRUTH:** Luka Doncic

**gpt-4o-2024-11-20:**  
**Explanation:** The question provides specific details about technical fouls involving a visiting-team bench player and a home-team player during an NBA 2025 regular-season game, **but it does not provide any details or data** related to points scored by any specific player. Without any further context or scoring information, it is **impossible to determine** which player's scoring was the highest in this game.  
**Answer:** **Not determinable** from the given information

**gpt-5.1-with-search:**  
**Explanation:** I was unable to locate **any 2025 NBA regular-season game** description that matches the detailed scenario given: a visiting player on the bench getting a technical for angrily complaining about an opponent grabbing the rim for rebounds, and later in the second quarter tangling and pushing at the baseline with a home-team player who also received a technical. Without being able to identify the specific game and players, **I cannot know the box score**, and therefore cannot determine whose scoring was highest in that game. Any answer would be a pure guess not grounded in available information.  
**Answer:** **Unknown** – the specific game and players cannot be reliably identified from public information, so the highest-scoring player in that game cannot be determined.

**gemini-2.5-pro-2026-06:**  
**Explanation:** The question describes a **hypothetical NBA game** set in the future (2025) involving **unidentified players**. The information provided, such as the players receiving technical fouls for complaining or pushing, **has no bearing on their scoring ability or performance in the game**. A player's point total is independent of such events. Since the game **is fictional and no statistics or player identities** are given, there is **no basis to determine** who scored more points.  
**Answer:** It is **impossible to determine** which player's scoring was the highest because the scenario is hypothetical and provides no information about the players' actual game performance or statistics.

Figure 11: Level 2 case, NBA player.

**Level 2 Question:** There is a video that appears to be from an orbital space module. In the video, there are two people wearing dark blue uniforms, one wearing a headset, the other also wearing a headset, and there is another headset hanging around someone's neck. They place the chicken wings into the oven with the oven door left open; what temperature setting is shown on the oven?  
**GROUND TRUTH:** 85

**o4-mini-deep-search-2025-06-26:**

**Explanation:** According to official reports, **China's new space oven was designed** to reach up to **190 °C** for actual cooking on the station. For example, an India Today article describes how "the oven's maximum temperature was also raised ... **to 190 °C**" ([www.indiatoday.in](https://www.indiatoday.in/amp/science/story/watch-chinese-astronauts-bake-chicken-wings-in-space-using-special-oven-2813340-2025-11-04#:~:text=during%20cooking,of%20merely%20warming%20prepackaged%20meals)). The video in question appears to show that oven at its operating temperature. Thus, the oven's display in the video is showing about 190 °C.

**Answer:** 190 °C

**gpt-5.1-2025-11-13-with-search:**

**Explanation:** This question refers to a very specific video (likely a particular ISS or space-module themed clip) and asks for an exact numeric temperature shown on an oven display. I attempted to search for it, but I couldn't reliably locate the precise video or a trustworthy description that states the temperature. Without being able to see the video frame or find a definitive reference, any specific **temperature (e.g., 350°F, 400°F, etc.)** would be a pure guess, not grounded in evidence. Therefore, I cannot honestly supply a numeric value.

**Answer:** **Unknown** – the temperature setting on the oven cannot be determined from the information available.

**gemini-2.5-pro-2025-06-with-search:**

**Explanation:** The provided video shows **Chinese astronauts on the Tiangong space station cooking chicken wings**. Multiple news reports and articles covering this event specify the temperature setting used for the oven.

**Answer:** The temperature setting shown on the oven in the video is **180 degrees Celsius**.

Figure 12: Level 2 case, space module.

**Level 3 Question:** In a comic universe that competes with DC, a green-clad villain who is delighted to see his former opponent appear shouts 'This is the feeling!' The film's protagonist in a movie that contains the classic scene 'the Mk I Arc Reactor drifting on a tranquil river' gets stained with a type of food. What is this food?

**GROUND TRUTH:** Alcohol

**gpt-4o-2024-11-20:**

**Explanation:** The comic universe that competes with DC is **Marvel**. The green-clad villain in the MCU is likely the **Green Goblin** or someone similar, though that's a minor detail here. The protagonist referenced here is **Tony Stark (Iron Man)** in the first **Iron Man** film. The scene described, 'the Mk I Arc Reactor drifting on a tranquil river,' occurs after Pepper Potts places the reactor in a display case marked '**Proof That Tony Stark Has a Heart**.' As for the food, in the film, Tony Stark gets stained with food when he is **eating a burger, specifically a cheeseburger**, which is tied to his character recovery arc when he insists on eating one after returning to the U.S. following his captivity.

**Answer:** A cheeseburger

**gpt-5.1-with-search:**

**Explanation:** The question references a 'comic universe that competes with DC'—this clearly points to **Marvel**. A green-clad villain who is delighted to see his former opponent and shouts 'This is the feeling!' **evokes Loki**, but the key is the second clue: the film that contains the classic scene of 'the Mk I Arc Reactor drifting on a tranquil river.' The Mk I Arc Reactor is central to Tony Stark in the MCU, specifically in '**Iron Man**' (2008). In that film, there is a memorable bit where Tony Stark gets stained with a food: during a press conference and later scenes, he is shown **eating a cheeseburger and getting it on himself**. The most iconic, repeatedly mentioned food associated with Tony Stark in that first movie is **a cheeseburger**, which he insists on right after returning from captivity, and he gets messy with it. Thus, the food that stains the protagonist (Tony Stark) is a cheeseburger.

**Answer:** **Cheeseburger**.

**gemini-2.5-pro:**

**Explanation:** The comic universe that competes with DC is **Marvel**. The green-clad villain is **Loki**, who shouts 'YES! That's how it feels!' in 'Thor: Ragnarok' upon seeing Hulk smash Thor, a callback to his own experience in '**The Avengers**'. The movie with the 'Mk I Arc Reactor drifting on a tranquil river' is '**Avengers: Endgame**', during Tony Stark's funeral. Therefore, the protagonist is Tony Stark (Iron Man). In the opening scene of '**Avengers: Infinity War**', Tony Stark is talking with Pepper Potts in a park and spills a popsicle on his shirt. Pepper identifies the stain as being **from an acai popsicle**.

**Answer:** **Acai (from a popsicle)**

Figure 13: Level 3 case, cross-film.

### Direct Visual Inference Prompt

You are an intelligent video assistant.

User Query: "user\_query"

You have analyzed len(target\_videos) videos. Here is the information extracted from them:  
combined\_context\_text

Instructions: - Use the provided transcripts and visual frames from ALL videos to answer the user's query.

- Synthesize information if the answer is spread across multiple videos.
- If none of the videos contain the answer, state that clearly.
- Be concise and direct.

Your response should be in the following JSON format:

```
{  
  "Explanation": "your explanation for your final answer",  
  "Answer": "your succinct, final answer",  
  "Confidence": "your confidence score between 0% and 100% for your answer"  
}
```

Figure 15: Direct Visual Inference Prompt

### Summarization Prompt

You are a video analyst.

User Query: "user\_query"

Video Title: video.get('title', 'Unknown') Transcript: transcript\_text[:20000]

Instructions: - Analyze the visual frames and transcript provided.

- Provide a detailed summary of the video content that is RELEVANT to the User Query.
- If the video contains specific details (numbers, names, events) related to the query, include them explicitly.
- If the video is irrelevant, state that it contains no relevant information.

Figure 16: Summarization Prompt

### Summarization Analyst Prompt

You are an intelligent assistant.

User Query: "user\_query"

You have received detailed summaries from several videos analyzed by your vision team: context\_text

Instructions: - Synthesize the information from these summaries to answer the user's query.

- If the answer is found, provide it clearly and concisely.
- If the videos do not contain the answer, state that.
- Do not make up information not present in the summaries.

Your response should be in the following JSON format:

```
{{  
  "Explanation": "your explanation for your final answer",  
  "Answer": "your succinct, final answer",  
  "Confidence": "your confidence score between 0% and 100% for your answer"  
}}
```

Figure 17: Summarization Analyst Prompt

## Planner Prompt

You are the Planner for a Video Agent.

Your primary role is to analyze the User's Goal and current investigation status, then generate the most effective search queries to find relevant video content.

**Current Task:** Based on the information gathered so far, formulate search queries to find missing information or refine existing knowledge to fully address the User's Goal.

### Context:

- You are part of a research loop. Your output ('search\_queries') will be used by a 'Searcher' tool.
- You do NOT need to decide to "watch" or "finish"
- the system flow handles that. Your ONLY job is to determine \*what to search for next\*.

### Action Guidelines:

1. **Decompose the User's Goal:** Break down the complex query into smaller, distinct sub-questions or requirements.
2. **Analyze Knowledge Gaps (The "Gap Analysis"):**
  - \* For each sub-question, check the 'Current Knowledge Status'.
  - \* **Solved:** Do we already have verified evidence for this part? -> \*Do not search for this again.\*
  - \* **Partial/Unverified:** Do we have watched videos waiting for verification? -> \*Wait, or search for corroborating details if unsure.\*
  - \* **Missing:** Is this part completely missing? Or did previous attempts fail (Rejected videos)? -> \*This is your Focus.\*
3. **Follow the Leads (Deep Dive / Hypothesis Verification):**
  - \* **CRITICAL:** If previous videos (even rejected ones) mentioned specific **candidates, entities, or clues** (e.g., "Video mentioned onions and garlic as staples"), **PIVOT** your strategy.
  - \* Stop searching for the general topic (e.g., "Ancient Egyptian food").
  - \* Instead, **verify these candidates** against the user's remaining conditions (e.g., "Did onions originate in Central Asia?", "Garlic symbolism medieval Europe").
4. **Formulate Targeted Queries:**
  - \* Generate queries **ONLY** for the Missing parts or **Hypothesis Verification**.
  - \* If previous queries failed (see 'Rejected Videos' reasons), change your strategy (e.g., use more specific technical terms, add date filters like "2024", or switch from "how to" to "review").
  - \* **LIMIT:** Generate a maximum of max\_queries queries.

### Example of Desired Reasoning:

\*User Query:\* "History of the Roman Empire and the engineering behind aqueducts."

\*Current Knowledge:\*

- [Verified] "Roman Empire Overview" (Covers the history timeline, emperors)
- [Rejected] "Aqueduct Tour" (Rejected: "Visuals only, no engineering explanation")

\*Your Output:\*

```
{
  "thought": "1. Decomposition:\n  - Q1: History? -> [SOLVED] 'Roman Empire Overview' covers timeline.\n  - Q2: Engineering of Aqueducts? -> [MISSING] 'Aqueduct Tour' was rejected for lacking details.\n2. Strategy: Ignore history. Focus specifically on the structural engineering and physics of aqueducts.",
  "search_queries": ["roman aqueduct structural engineering physics", "how roman aqueducts work construction details"]
}
```

Figure 18: Planner Prompt

**\*\*Example of Hypothesis Verification (Deep Dive):\*\*** \*User Query:\* "A food originating in Central Asia, depicted in Egyptian murals, symbolizing victory in Europe."

\*Current Knowledge:\*

- [Rejected] "Egyptian Diet" (Reason: "Mentions **\*\*Garlic\*\*** and **\*\*Onions\*\*** as staples, but doesn't confirm Central Asia origin.")

\*Your Output:\*

```
{
  "thought": "1. Leads Found: Previous video identified 'Garlic' and 'Onions' as candidates for the Egyptian part.\n2. Gaps: Need to verify if either of them matches the 'Central Asia origin' and 'Victory symbolism' conditions.\n3. Strategy: Switch from general search to specific verification of these candidates.",
  "search_queries": ["did garlic originate in Central Asia", "onion origin Central Asia", "garlic symbolism victory medieval Europe", "onion symbolism victory Europe"]
}
```

**\*\*Output Format:\*\*** Provide a JSON object with the following structure:

```
{
  "thought": "Step-by-step reasoning: Decomposition -> Leads/Hypothesis -> Gap Analysis -> Strategy.",
  "search_queries": ["query1", "query2", "query3"]
}
```

### Select Prompt

User Query: "user\_query"

You are provided with a list of video search results. Select the top top\_k most relevant videos that are likely to contain the answer to the User Query.

Candidates: candidates\_info

Return the indices of the selected videos as a JSON list of integers. Example: [0, 2, 4]

Figure 19: Select Prompt

### Localization Prompt

You are a video investigator.

User Query: "user\_query"

Video Title: video.get('title', 'Unknown')

I have sampled len(frames\_data) frames from the video at specific timestamps: frame\_descriptions

Transcript (with timestamps): truncated\_transcript

Task: Identify all specific temporal windows (start\_time to end\_time) in the video that are relevant to answering the User Query. - Combine visual cues from the frames and semantic cues from the transcript. - If the answer is in the transcript, use its timestamps. - If the answer is visual, use the frame timestamps to estimate the window. - If the video is irrelevant, return an empty list or set relevant to false. - Provide a brief reasoning for each selected window.

Output Format (JSON):

```
{{
  "relevant": true/false,
  "windows": [
    {{
      "start_time_seconds": <float>,
      "end_time_seconds": <float>,
      "reasoning": "...
    }},
    ...
  ]
}}
```

Figure 20: Localization Prompt

### Analyst Prompt

Based on the video clips and transcripts provided above, answer the User Query. Response Format:

```
{{
  "Explanation": "your explanation for your final answer",
  "Answer": "your succinct, final answer",
  "Confidence": "your confidence score between 0% and 100% for your answer"
}}
```

Figure 21: Analyst Prompt

**Question:** In a 2024 NBA playoff game, after finishing an alley oop dunk, on the defensive end he recorded a big block on the opponent's three-point attempt. He was then tripped by the opponent, and after being tripped the player angrily slammed the floor. Who was the player whose shot was blocked?  
**Answer:** Jayson Tatum

**Direct Visual Inference:**

The described play—an alley oop dunk followed by a chase-down block on a three-point attempt, a trip, and the player angrily slamming the floor—does not appear in any of the three provided videos or their frames. Video 1 shows Giannis Antetokounmpo's famous 2021 NBA Finals alley-ooop from Jrue Holiday vs. the Phoenix Suns, which includes a steal on Devin Booker and a foul by Chris Paul, **but there is no blocked three-point attempt or trip afterward**. Video 2 is a general compilation of blocked dunks from many seasons and teams, not a specific 2024 playoff sequence. Video 3 is a Bronny James alley-ooop highlight with no such follow-up defensive play. Because the specific 2024 playoff moment the user describes is not contained in the supplied material, **I cannot identify whose shot was blocked** from these videos alone.

**Answer:** The videos provided do not contain the 2024 playoff sequence described, so the player whose shot was blocked cannot be determined from this material.

**Summarization:**

**None of the provided video summaries describe** a 2024 NBA playoff sequence where a player finishes an alley oop dunk, then blocks a three-point attempt on the defensive end, is tripped, and angrily slams the floor. Video 1 is a regular-season Wizards vs. Pelicans game with no three-point attempt or trip, Video 2 is from the 2023 Western Conference Finals and only shows an offensive alley oop, and Video 3 is from the 2016 NBA Finals with **no relevant sequence**. Therefore, the player whose shot was blocked cannot be identified from the given information.

**Answer:** The videos do not provide enough information to identify the player whose three-point shot was blocked.

**Pyramidal Perception:**

The described sequence—an alley oop dunk followed by a big block on a three-point attempt, then the defender being tripped and angrily slamming the floor—occurs in the 2024 NBA Finals, Mavericks vs. Celtics. In the clip labeled “Luka Doncic alley-ooop pass to Gafford and Block by gafford🔥🏀Mavs Vs Celtics Game 4,” **the commentator explicitly calls the play: “Tatum step back, stopped, blocked** by Gafford on a three-point attempt, and then a loose ball foul on Tatum.” This confirms that the blocked three-pointer was taken by Jayson Tatum.

**Answer:** Jayson Tatum.

Figure 22: Qualitative comparison of three paradigms.

**Question:** In a 2025 NBA regular-season game, one player first recorded a jaw-dropping rim block on defense, then missed an open layup on offense, and later on a fast break was called for an offensive foul. What was this player's final score for that game?  
**Answer:** 6

**ours:**

The relevant clip is “Unbelievable Basketball Move Kyle Lowrys Jaw Dropping Under the Rim Offensive Foul 1.” It describes, in one 2025 regular-season game, Kyle Lowry committing two Shaqtin’-type plays in the same game, including a bizarre under-the-rim offensive foul on a fast break. However, the provided transcript only captures the commentators reacting to the play itself; **it never states Lowry's stat line, box score, or final points total for that game**. None of the other supplied clips or transcripts add any game-stat information about this sequence. Without a box score or additional context, his final scoring total for that game cannot be determined from the materials given.

**Answer:** The player is Kyle Lowry, **but his final point total for that game is not stated anywhere** in the provided clips or transcripts, so it cannot be determined from the given information.

Figure 23: Failure Case 1.

**Question:** In a 2025 NBA regular-season game, a player recorded a steal and a slam dunk in the third quarter, and later, on a ball-saving play in the fourth quarter, he passed the ball to the bench. Which two teams did this player win championships with?  
**Answer:** Los Angeles Lakers, Oklahoma City Thunder

**ours:**

From the Lakers–Clippers full game highlight transcript (December 20, 2025), the specific play described in the question appears in the third quarter: “Here it is LeBron all the way... three point opportunity...” followed soon after by commentary noting “LeBron ... the king lays it up and in,” and later a sequence where he gets a steal and goes “straight to the house,” i.e., a slam. In the fourth quarter, the other highlight reel of the same game (the national broadcast version) includes a notable hustle/ball-saving play where LeBron, while saving the ball from going out of bounds, ends up passing/throwing it toward the Lakers bench area. The player in question is LeBron James. LeBron has won NBA championships with two franchises: **the Miami Heat and the Cleveland Cavaliers**.

**Answer:** Miami Heat and Cleveland Cavaliers

Figure 24: Failure Case 2.

**Question:** There is a travel reality show starring an actress who appeared in 'Eternal Love' and an actor who appeared in 'Romance in the Alley', both participating. They visit a country whose capital is the closest to the Arctic Circle, experiencing local culture. During a festival focused on food, an actress who acted in 'Being a Hero' tries a food that makes her extremely uncomfortable, yet she still pretends it's tasty. What is this food?

**Answer:** sheep's butt

**ours:**

The show is filmed in Iceland, whose capital Reykjavík is the closest national capital to the Arctic Circle. In the Iceland food-exploration clips, one dish is highlighted as extremely smelly and hard to eat: “冰岛发酵的酸鲨鱼肉”, described as one of the smelliest foods in the world. The taster reacts with obvious discomfort but forces himself to continue and describes it despite clearly not liking it—matching the description that the actress from 'Being a Hero' finds the food very uncomfortable yet pretends it's tasty. Therefore, the food in question is Icelandic fermented shark.

**Answer: Icelandic fermented shark (fermented sour shark meat, hákarl).**

Figure 25: Failure Case 3.

**Question:** In the same film, after the wealthy grandfather dies, the grandson must pass 12 tests; the second test is to form a true friend. What drink did this friend drink when presenting themselves to the lawyer for confirmation?

**Answer:** Coca-Cola

**ours:**

In the lawyer Hamilton's office, Jason's 'true friend' Emily sits at the conference table with a beverage. The video frame (around 489.8s) clearly shows **a can of Dr Pepper with a straw beside** her as she talks to Hamilton. No other drink is shown or mentioned for her in this confirmation scene.

**Answer: Dr Pepper (from a can, with a straw)**

Figure 26: Failure Case 4.