How Well Do Vision-Language Models Understand Movies?

Emmanouil Zaranis[1,2]*, António Farinhas[18]*,Saul Santos[1,2]*, Beatriz Canaverde[1,2]*, Miguel Moura Ramos[1,2]*, Aditya K Surikuchi[3], André Viveiros[1,2], Baohao Liao[4], Elena Bueno-Benito[5], Nithin Sivakumaran[6], Pavlo Vasylenko[1,2], Shoubin Yu[6], Sonal Sannigrahi[1,2], Wafaa Mohammed[4], Ben Peters[11], Danae Sánchez Villegas[7], Elias Stengel-Eskin[17], Giuseppe Attanasio[2], Jaehong Yoon[16], Stella Frank[7,8], Alessandro Suglia[9], Chrysoula Zerva[2,13,19], Desmond Elliott[7,8], Mariella Dimiccoli[5,15], Mohit Bansal[6], Oswald Lanz[10,14], Raffaella Bernardi[10,14], Raquel Fernández[3,12], Sandro Pezzelle[3,12], Vlad Niculae[4,12], André F. T. Martins[1,2,11,13]

[1]Instituto Superior Técnico, Universidade de Lisboa, [2]Instituto de Telecomunicações, [3]ILLC, University of Amsterdam, [4]Language Technology Lab, University of Amsterdam, [5]Institut de Robòtica i Informàtica Industrial, CSIC-UPC, [6]UNC Chapel Hill, [7]University of Copenhagen, [8]Pioneer Center for AI, [9]University of Edinburgh, [10]Free University of Bozen-Bolzano, [11]TransPerfect, [12]ELLIS Unit Amsterdam, [13]ELLIS Unit Lisbon, [14]ELLIS Unit Trento, [15]ELLIS Unit Barcelona, [16]Nanyang Technological University, [17]The University of Texas at Austin, [18]Sword Health, [19]National Technical University of Athens

Correspondence: emmanouil.zaranis@tecnico.ulisboa.pt, andre.t.martins@tecnico.ulisboa.pt

Abstract

Vision-language models (VLMs) can increasingly handle extended contexts and answer questions about long videos—but can they understand the narrative structure of movies as humans do? To investigate this question, we introduce MF^2 (*Movie Facts and Fibs*), a fully manually curated benchmark designed to evaluate whether models can identify, recall, and comprehend information central to the movie's storyline, using both the visual and textual modalities. MF^2 includes over 50 *open-licensed* movies, each paired with manually constructed sets of claim pairs: one true claim (*fact*) and one plausible but false (*fib*), totaling over 850 pairs. Our analysis across 20 open-source and 2 proprietary models reveals a substantial gap relative to human performance. While subtitles provide a strong signal for narrative understanding, other textual sources (e.g., synopses) are less informative. In contrast, when visual information is distributed across the entire movie and jointly embedded with subtitles, models struggle to identify and effectively use the relevant content. However, when the visual input is restricted to claim-relevant visual content, models improve significantly, indicating substantial room for improvement in processing large multimodal contexts and distinguishing relevant from irrelevant information. Overall, our findings highlight persistent challenges in multimodal narrative understanding for current VLMs. We release our data and code to support further research in this direction.

*Equal contribution.

1 Introduction

The past few years have seen remarkable advances in vision-language models (VLMs), particularly in their capacity to process extended multimodal context and perform question answering across long video streams (Team et al., 2023; Zhang et al., 2024; Bai et al., 2025a; Zhu et al., 2025). Yet their ability to understand movie narratives in a human-like manner remains underexplored. This is an important gap, as narratives provide a natural testbed for studying the long-context capabilities of current models: storylines are shaped by how events, entities, emotional arcs, and causal interactions between characters unfold over time. As a result, successful understanding goes beyond local perception or step-by-step reasoning, and instead depends on integrating information across extended spans, while maintaining coherence over evolving relationships between characters and events.

In textual domains, literal and narrative settings have increasingly emerged as an important aspect of long-context evaluation (Hudson and Al Moubayed, 2022; An et al., 2024; Piper and Bagga, 2024; Karpinska et al., 2024; Hamilton et al., 2026). Similar extensions have begun to arise in multimodal settings, particularly focusing on videos, with recent benchmarks shifting towards longer content (Ataallah et al., 2025; Wang et al., 2025a; Fu et al., 2025a; Wu et al., 2024). However, the majority of these efforts rely on general-purpose video content (e.g., egocentric recordings,

Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing, pages 38762–38801
October 24-29, 2026 ©2026 Association for Computational Linguistics

instructional videos, or vlogs), which, while suitable for assessing perception and reasoning capabilities over extended sequences, are not explicitly designed to capture narrative structure and assess the relationships that shape a storyline. A smaller subset of works instead leverages narrative-rich content, such as movies or TV shows, aligning better with the assessment of storyline development (Huang et al., 2020; Song et al., 2023, 2024; Ataallah et al., 2025). However, these are often constrained by *limited temporal coverage* or *restricted accessibility due to copyright considerations*, often providing only partial spans of the original content. Moreover, some rely on automatically generated questions (Ataallah et al., 2025; Chandrasegaran et al., 2024; Faure et al., 2025), which evaluate content uniformly and may place disproportionate emphasis on non-narrative details rather than on central storyline events.

In this work, we investigate the extent to which VLMs can understand and reason over movie narratives. Building on prior research in narrative understanding across textual and multimodal domains (Papalampidi et al., 2019; Karpinska et al., 2024; Hamilton et al., 2026), we focus on whether models can answer questions related to *core narrative elements*, which shape the narrative trajectory (Papalampidi et al., 2019, 2020). Rather than assessing models uniformly across all parts of the content, as is typically done with automatically generated questions, our analysis targets events central to the storyline that are retained in *human memory*, namely those that viewers can *recall without rewatching the video due to narrative importance*.

To support this analysis, we construct MF^2 (*Movie Facts and Fibs*), a *manually annotated* evaluation set of 868 claim pairs grounded in 53 open-licensed movies. Each pair consists of a true statement (a *fact*) and a plausible but false counterpart (a *fib*) (see Figure 1 and Examples in Appendix F). Following prior work on narrative understanding (Karpinska et al., 2024), we adopt a claim-based evaluation protocol instead of multiple-choice questions. This design is motivated by two considerations: (i) it enables isolated statement-level truth assessment without reliance on contrastive options and (ii) it reduces biases from poorly constructed distractors in multiple-choice setups (Li and Gao, 2024; Loginova et al., 2024; Singh et al., 2025; Molfese et al., 2025).

Our analysis across 20 open-source and 2 proprietary models reveals several key findings: (i) models substantially underperform humans on the task, with the best-performing model lagging behind human performance by 20.3% absolute; (ii) subtitles help models reason better about narratives, while other text forms (e.g., synopses) are not equally informative; and (iii) although claim-relevant visual evidence provides a strong signal when directly provided as input, models struggle to effectively retrieve and utilize it when it is embedded within full-movie context. Together, these findings highlight an important challenge for current VLMs: identifying and reasoning over the relevant information in long-form movies, while ignoring irrelevant details. As part of our contribution to support reproducibility and future research, we release MF^2's data and code under a CC-BY 4.0 license, including raw video content to mitigate link rot and ensure long-term accessibility.[1]

2 Dataset Construction

MF^2 includes 53 full-length, open-licensed movies, each accompanied by subtitles, and 868 human-authored contrastive claim pairs. Each pair tests whether a model can distinguish true from false information based on its understanding of the story. Figure 1 shows some examples (see more in Appendix F). We now describe the dataset construction process in detail, covering movie selection (Section 2.1), the annotation protocol (including claim categorization and reasoning granularities; Section 2.2), and human quality control procedures used to filter ambiguous or low-quality claims (Section 2.3). Figure 2 illustrates these three stages.

2.1 Movie Selection and Subtitles

We started by collecting a pool of movies from the Internet Archive,[2] an online repository of open-licensed media. We specifically selected titles released under the Public Domain 1.0 license to ensure legal reusability and support open-access research. We focused on older films, prioritizing those with limited online visibility, measured by the number of user reviews on IMDb, to reduce the risk of data contamination. We sourced original-language subtitles—the majority of which are in English—from OpenSubtitles.org,[3] a widely used platform that provides subtitles for a large collection of movies, TV shows, and other video content.

[1] https://github.com/deep-spin/MF2-bench-can-vlms-understand-movies

[2] https://archive.org

[3] https://www.opensubtitles.org

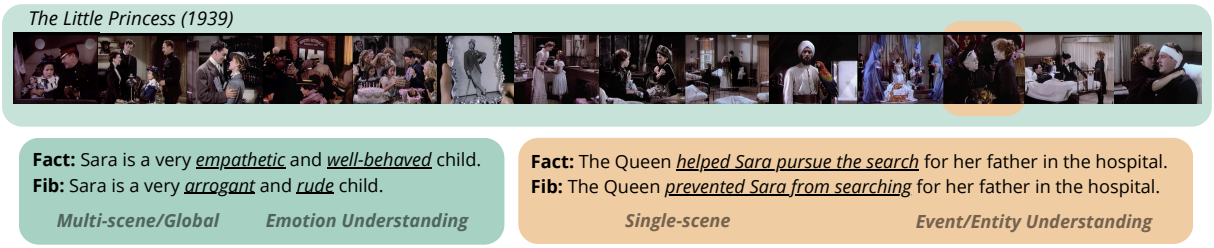


Figure 1: Illustration of two claim pairs (each with a *fact* and a *fib*) from the movie "The Little Princess" (1939). Our claims target memorable events, focusing on key points of the narrative such as emotional arcs and causal relationships between characters, and require reasoning across different granularities (single-scene and multi-scene/global).

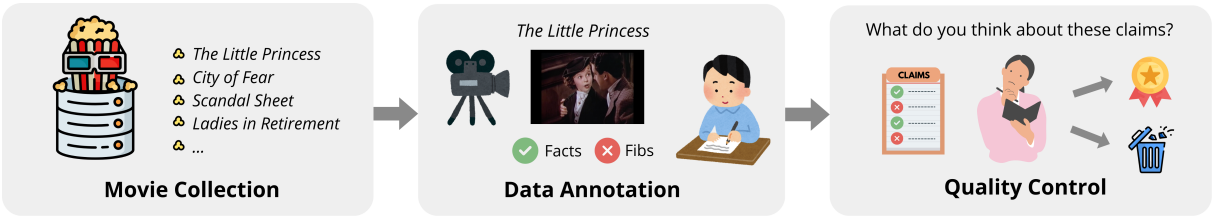


Figure 2: Dataset construction process involving three main stages: movie collection, data annotation, and quality control.

This process yielded a final collection of 53 full-length movies with an average duration of 88.33 minutes, each accompanied by audio and aligned subtitles. The collection spans a wide range of genres as we detail in Appendix A.

2.2 Data Annotation

In contrast to approaches that rely on semi-automatically generated questions, often produced by language models and potentially reflecting model-specific biases, we opted for fully manual annotations, *prioritizing quality over quantity*. The annotation process involved 26 annotators (see Appendix A for details) who watched the full movies, identified key narrative elements, and constructed pairs of contrastive claims: one factually correct statement (*fact*) and one minimally altered, false counterpart (*fib*). Following Karpinska et al. (2024), annotators were instructed to minimize lexical differences between the *fact* and the *fib*, changing only the parts needed to flip the truth value. The data annotation guidelines are presented in Appendix D.

Claim granularity. To capture different levels of reasoning, annotators labeled each *fact* according to the granularity required to verify its truth: *(i) single-scene*: answerable using information from one scene; or *(ii) multi-scene/global*: requiring integration of information across multiple scenes or higher-level understanding that may not be easily attributable to specific scenes, including cases involving accumulated or inferred information across the movie. The dataset exhibits a balanced distribution of single-scene and multi-scene/global facts, comprising 49.4% and 50.6%, respectively.

Comprehension dimensions. In addition to the reasoning granularity, annotators labeled each claim pair with one or more comprehension dimensions, indicating the specific aspects of narrative understanding being tested. These dimensions, informed by prior work (Xiao et al., 2021; Zhang et al., 2023b; Wang et al., 2025a), include event/entity understanding, temporal perception, emotion understanding, and causal reasoning. We provide detailed definitions of each dimension and their distribution in Appendix A.

Grounding Timeframes. For a subset of single-scene claim pairs (12% of the total pairs), annotators were instructed to mark the visual segment (timespan) of the movie required to correctly answer the claim pair. This provides a basis for evaluating models' use of relevant visual content.

2.3 Quality Control

To ensure dataset quality and support the identification of problematic claims, we include a quality-control stage in which annotators answer existing claims and provide constructive feedback. Feedback focused on claims that were ambigu-

ous, poorly phrased, open to interpretation, or too fine-grained to be meaningfully tied to narrative understanding (e.g., needle-in-a-haystack claims). Specifically, annotators first selected a subset of movies they had not previously seen during the data annotation stage. After watching a movie, they classified the corresponding claims as either true or false using a custom annotation interface (see Appendix D for an example and full guidelines). Claims were presented one at a time, and annotators were required to respond based solely on memory, without rewatching the movie or using external tools or notes. Annotators were encouraged to leave comments whenever a claim appeared ambiguous, poorly phrased, open to interpretation, or excessively fine-grained. After this stage, annotators could optionally revisit their earlier responses with access to the movie in order to provide additional validation comments.

Finally, two annotators reviewed all collected comments from the stages described above. Based on the provided feedback, and after discussion to reach consensus, they identified problematic claims and removed them from the dataset. Importantly, no claims were rewritten during this process; they were either accepted or discarded. This filtering stage resulted in the removal of 104 pairs of claims, yielding a cleaner set of 868 high-quality pairs.

3 Experimental Setup

We provide an overview of the models included in our study, as well as important information regarding the input modalities used.

Input Modalities. We evaluate all models in a vision-language setting, where visual input is provided through sampled movie frames. In addition, we investigate the effect of language by providing subtitles either independently or jointly with the visual input.[4] Importantly, models always receive frames sampled *uniformly across entire movie*, rather than from targeted scenes. Consequently, models must process the full provided content to identify the relevant information and determine whether a claim is true or false.

Models. Our evaluation suite covers 2 proprietary and 20 open-weight VLMs. Among the proprietary models, we include GPT-4o (OpenAI et al., 2024), Gemini 2.5 Pro (Team et al., 2023). We cover several open-weight VLM families, including QwenVL (Bai et al., 2025b,a), InternVL (Zhu et al., 2025; Wang et al., 2025b), Ovis2 (Lu et al., 2024), Gemma3 (Team et al., 2025), GLM-4.5V (Team et al., 2026), among others. We additionally include models specialized for long videos, such as LongVILA-R1 (Chen et al., 2025) and MiMo-VL-7B-RL (Xiaomi, 2025). We include all prompt, frame sampling and answer parsing details in Appendix E for reproducibility.

Human baseline. To establish a human baseline for comparison, 24 annotators evaluated the full set of constructed claim pairs, in contrast to prior work that typically relies on smaller subsets. Annotators selected movies they had not previously seen during the data annotation stage. After watching a movie, with optional access to subtitles, they classified the corresponding claims as either true or false using the same interface employed during quality control (see Appendix D for an example and full guidelines). Each movie was evaluated by a single annotator, claims were presented one at a time, and annotators were required to answer based solely on memory, without rewatching the movie, using external tools, or taking notes.[5]

Evaluation protocol. We report two metrics: (i) pairwise accuracy, which measures how often models correctly classify both the true and the false claim in a pair (i.e., credit is given only if both are labeled correctly; no points are awarded for partial correctness); and (ii) standard accuracy, which is computed over individual claims. The random baselines are 25% and 50%, respectively. Following prior work (Karpinska et al., 2024), *models answer each claim independently, without access to the paired structure during prediction*. Pairwise accuracy is computed post-hoc by grouping predictions from the same pair. This particular design (i) allows for assessing isolated statement-level reasoning and truth prediction, and does not involve contrastive reasoning between paired alternatives, (ii) reduces biases caused by poorly constructed distractor options in multiple-choice questions (Loginova et al., 2024; Singh et al., 2025; Molfese et al., 2025), and (iii) avoids the use of automatic judges, which may struggle with rich long-form video (Bavaresco et al., 2025; Liu and Zhang, 2025; Ye et al., 2025b).

[4] Subtitles are appended to the visual input.

[5] Due to the labor-intensive nature of the process, we do not repeat the experiment under alternative input conditions (e.g., video-only), and instead use this setting as a *reference estimate of human performance* for movie narrative understanding.

In Appendix B, we discuss this design choice and the use of pairwise accuracy as our primary metric.

4 Empirical Findings

We now present our empirical findings, starting with the main comparison between models and humans and an analysis across input modalities (Section 4.1). This is followed by a more detailed study of Gemini 2.5 Pro, the best-performing model among those evaluated, in Section 4.2.

4.1 The Role of Language and Vision

In this section, we analyze model performance in comparison to human performance, as reported in Table 1, and examine how different input modalities (subtitles-only, video-only, and video with subtitles) affect performance across open-weight and closed models, identifying several key challenges for VLMs in movie narrative understanding.

Models perform significantly worse than humans. Among closed models, Gemini 2.5 Pro achieves the highest scores, with a pairwise accuracy of 63.8%, followed by InternVL3-72B, which performs 12.5% lower on video and subtitles. Despite their fair performance, *both models rank significantly behind humans*. Among small-scale models ($\leq$9B), Molmo2 achieves the highest performance, outperforming the random baseline by 18% in the video with subtitles setting, while Ovis2 performs best in the medium-scale category (25B–38B). These results underscore the difficulty of the task for current models, while highlighting humans' superior ability to retain, integrate, and reason over critical narrative content.

Subtitles provide a stronger narrative signal than visual information alone. Across nearly all models, performance with subtitles only is remarkably higher than when relying exclusively on visual information sampled across the entire movie. Since subtitles play a central role in constructing the storyline through dialogue, character interactions, and explicit descriptions of events and actions, this result confirms that such textual signals provide a particularly strong source of narrative information for the task.

Coupling subtitles with visual information yields no consistent benefits over subtitles alone for most models. One might expect visual information to provide complementary signals to subtitles, leading to substantial gains when both modalities are available. However, as shown in Table 1, adding visual information yields no consistent improvements across the evaluated models. This suggests that either visual cues are not consistently helpful, or that models do not effectively identify and attend to the relevant visual information when embedded in long multimodal contexts. In the following paragraph, we further examine this hypothesis by investigating whether visual information becomes beneficial when the relevant visual content is explicitly provided to the models in isolation.

Explicit visual cues provide a useful signal. We evaluate whether performance changes when the visual input is restricted to the parts of the movie that are directly relevant to each claim. Specifically, we use a subset of single-scene claims with annotated time spans identifying the relevant visual spans (see Section 2.2). We extract targeted frames from these spans instead of sampling uniformly from the full movie. We evaluate Gemini 2.5 Pro alongside five mid and large scale open-weight models. In Table 2, we report performance on this subset under subtitles-only, video-only, and video-with-subtitles settings, comparing results between full-movie and targeted-frame inputs. Across most models, performance improves when given targeted visual content, with and without subtitles.[6] The consistent improvements observed with targeted visual content demonstrate that *visual cues are indeed useful when explicitly provided to the model*. Therefore, the observed behavior likely stems from difficulties in retrieving and effectively integrating relevant visual information with textual cues in long multimodal contexts, highlighting substantial room for improvement in these aspects.

The key challenge: Overall, our results suggest that models substantially lag behind human performance in movie narrative understanding. Beyond this gap, (i) subtitles provide a strong signal that significantly helps reasoning in narrative understanding, and (ii) the relevant visual cues are nevertheless informative when directly provided as input, but are less effectively utilized in long-context settings (as also illustrated by the qualitative Ex-

[6] Qwen2.5-VL-72B is the only exception, showing a 4.9-point drop in the video-only setting. This behavior appears to be associated with a more conservative prediction tendency, with the model more frequently rejecting both claims; a similar bias toward rejecting claims is also observed across other models (see Appendix B). This tendency disappears when subtitles are provided, suggesting that textual information helps interpret the visual cues.

Method	#Params	Pairwise Accuracy (%)			Accuracy (%)		
		Subs	Video	Video+Subs	Subs	Video	Video+Subs
Baselines							
Random	-	25.0	25.0	25.0	50.0	50.0	50.0
Human	-	–	–	84.1	–	–	90.5
Closed Models							
GPT-4o	-	48.8	19.5	49.8	72.9	55.5	73.2
Gemini 2.5 Pro	-	**58.9**	**37.9**	**63.8**	**77.5**	**64.4**	**79.8**
Open-weight Models							
mPLUG-Owl3*	7B	–	–	–	–	–	–
VideoLLaMA3	7B	36.3	21.1	36.1	64.0	56.8	64.2
Qwen2.5-VL	7B	31.7	24.5	32.8	62.1	56.7	62.0
LLaVA-Video	7B	20.5	6.5	19.1	58.2	51.4	57.9
LongVILA-R1	7B	9.8	11.5	16.9	54.0	50.1	56.6
MiMo-VL-7B-RL	7B	36.4	25.8	38.6	62.1	57.5	65.2
Kangaroo*	8B	–	–	–	–	–	–
InternVL3.5	8B	39.2	18.9	37.2	65.4	55.1	64.5
LLaVAOneVision1.5	8B	37.4	4.9	26.4	65.2	51.0	60.9
Molmo2	8B	41.0	28.3	43.0	67.2	59.0	67.7
Phi4*	14B	26.8	–	–	58.8	–	–
Aria	25B (3.9B)	36.2	2.5	33.4	64.1	50.0	63.2
Gemma3	27B	42.3	31.5	39.7	67.8	61.2	65.3
Qwen3VL-IT	30B	33.5	19.2	43.8	64.7	56.9	69.5
Ovis2	34B	47.2	20.7	48.5	70.3	56.5	71.1
InternVL3.5	38B	47.5	26.6	46.2	70.8	60.0	70.3
Qwen2.5-VL	72B	37.6	29.7	47.5	66.8	58.8	69.0
LLaVA-Video	72B	41.7	15.6	41.8	68.8	54.6	69.1
InternVL3	78B	49.0	22.1	51.3	72.0	58.0	72.7
GLM-4.5V-FP8	108B	38.4	26.6	31.4	64.7	55.0	62.4

Table 1: Performance (%) of open-weight and closed models on MF^2 under different input modalities. Best values overall are **bolded**; best per group are underlined. Asterisk (*) denotes models that failed to follow the instruction prompt in a substantial fraction of samples and their results are therefore not reported. Numbers in parentheses indicate the number of active parameters.

amples 5 and 6, in Appendix F). Taken together, these findings point to a central challenge for current VLMs: effectively identifying, retrieving, and integrating relevant visual evidence with textual information within extended multimodal contexts. Interestingly, our results align well with prior findings from core "vision-centric" tasks, where vision-language models have been shown to rely heavily on language priors and textual inputs (Fu et al., 2025b; Zheng et al., 2025).

4.2 A Closer Look at Gemini 2.5 Pro

We now turn to Gemini 2.5 Pro, the best-performing model, which we use as a representative model to examine its reliance on parametric (world) knowledge and how different forms of textual content (i.e., synopses) affect performance. Lastly, we analyze its behavior across reasoning granularities and comprehension dimensions.

Beyond vision and subtitles: the role of summarized content (synopsis) and world knowledge. Table 3 presents an ablation study examining reliance on textual context beyond subtitles. When provided only with the movie title and release year, the model achieves above-random performance, suggesting an influence of parametric (pretrained) knowledge (e.g., awareness of well-known films and their plots). In contrast, performance drops to near-random levels with only the movie synopsis, indicating that compressed narrative descriptions provide little useful signal for reasoning on the task (a trend that also holds for large-scale open-weight models; see Appendix C). Importantly, both settings remain substantially below human performance, with absolute gaps of 38.7% and 56.2%, respectively, suggesting that *neither memorization nor high-level summarization is sufficient to solve the task*.

Model	Video-only		Video+Subs	
	(full)	(targeted)	(full)	(targeted)
Gemini 2.5 Pro	38.6	49.4 (+10.8)	59.4	74.0 (+14.6)
Gemma3-27B	27.2	36.9 (+9.7)	42.7	49.5 (+6.8)
Qwen3VL-30B	12.6	22.3 (+9.7)	40.8	52.4 (+11.6)
InternVL3.5-38B	29.1	41.7 (+12.6)	51.5	53.4 (+1.9)
Qwen2.5VL-72B	34.0	29.1 (-4.9)	46.6	53.4 (+6.8)
InternVL3-78B	27.2	40.8 (+13.6)	49.5	66.0 (+16.5)

Table 2: Pairwise accuracy when models are provided with "targeted" frames as visual input versus frames sampled uniformly from the entire movie. "Targeted" frames refer to frames extracted from the annotated time intervals that contain the visual content relevant to each claim. Results are reported on a single-scene claim subset under video-only and video with subtitles settings. Values in parentheses indicate the performance difference relative to the full-movie setting. Positive and negative deltas denote performance gains and drops, respectively.

	Input Modality				
Metric	Video	Subs	Video+Subs	Synopsis	Movie Title
Pairwise Acc.	37.9	58.9	63.8	27.9	45.4

Table 3: Performance of Gemini 2.5 Pro across different input modalities. *Video* uses only the visual stream; *Subs* includes only subtitle information; *Synopsis* relies only on the synopsis of the movie obtained from Wikipedia; *Video+Subs* combines both visual and subtitles inputs; and *Movie Title* uses only the claim, along with the movie title and release year, without access to movie content.

Input modality contributions across comprehension dimensions and reasoning granularities. In Figure 3, we present ablation studies examining how different input modalities contribute to performance across comprehension dimensions and reasoning granularities. We observe that the presence of visual information particularly benefits claims involving emotional understanding, with a noticeable gap relative to other comprehension dimensions, especially in the video-with-subtitles setting (a trend that also holds for large-scale open-weight models as well-see Appendix C). In contrast, causal reasoning is notably weaker under video-only conditions, likely due to the need for linguistic disambiguation. This limitation becomes evident when subtitles are introduced: the largest improvement is observed for causal reasoning, highlighting the complementary role of textual context. Beyond comprehension dimensions, reasoning performance remains relatively consistent across reasoning granularities, although single-scene reasoning appears slightly more challenging under the video-only setting.

Success and failure patterns. We analyze representative success and failure cases of Gemini 2.5 Pro. The model performs well on clearly grounded and explicitly depicted information, where visual and textual cues are easy to align (see Examples 3 and 4 in Appendix F). In contrast, failures are more pronounced when the model must rely on implicit cues beyond surface alignment. Specifically, cases involving causal reasoning (e.g., inferring intentions) and implicit event/entity understanding, where relevant event instances must be identified and aggregated over time to form a coherent interpretation (see Examples 1, 2, 5 and 6 in App. F).

5 Related Work

Vision and long context LLMs. The field of VLMs has seen rapid progress, with models becoming increasingly effective at video-language understanding (Deitke et al., 2024; Bai et al., 2025b; Zhu et al., 2025). Early methods focused on short clips and relied on complex spatio-temporal modules, such as Q-formers (Zhang et al., 2023a; Li et al., 2023a), or temporal pooling techniques (Maaz et al., 2024; Luo et al., 2023; Xu et al., 2024a). While not new, projection layers (Li et al., 2023b; Liu et al., 2023; Li et al., 2023a; Liu et al., 2024b) have gained popularity as a simpler and increasingly effective alternative for aligning video and language representations (Bai et al., 2025b; Zhang et al., 2025; Zhu et al., 2025), largely driven by advancements in visual encoders (Radford et al., 2021; Tschannen et al., 2025). In the domain of long video understanding, current approaches primarily focus on compressing tokens (Li et al., 2023b; Zhang et al., 2025), merely extending the context window (Abdin et al., 2024; Liu et al.,

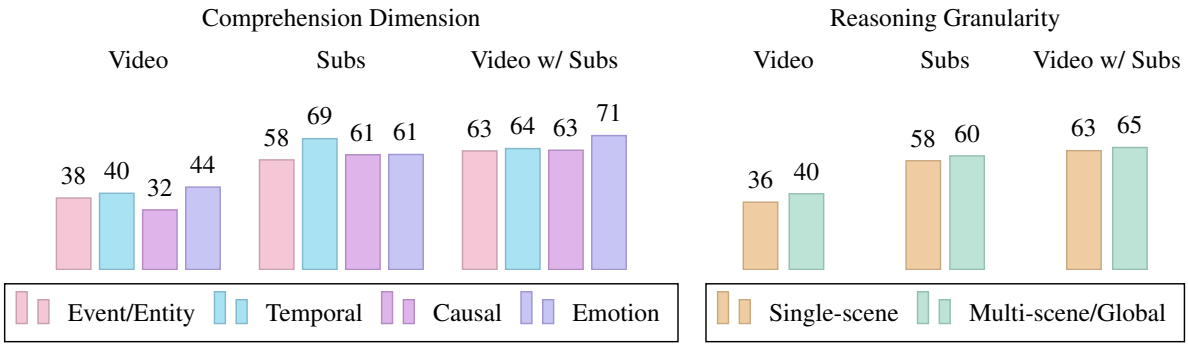


Figure 3: Pairwise accuracy for Gemini 2.5 Pro per comprehension dimension and reasoning granularity when varying the input modalities.

2025) or memory consolidation mechanisms (Balazevic et al., 2024; Song et al., 2023, 2024; Santos et al., 2025). A separate line of work first densely captions videos and then answers questions based on text only (Zhang et al., 2024; Wang et al., 2024a,b). We focus instead on benchmarking VLMs in movie narrative understanding without the use of costly captioning pipelines.

Long video understanding benchmarks. Understanding long videos presents substantial challenges, requiring models to track complex temporal dependencies and retain narrative context over extended durations. While existing benchmarks have driven progress in temporal reasoning over short clips (Xiao et al., 2021; Wu et al., 2021) and in domain-specific settings such as instructional or egocentric videos (Yang et al., 2021; Mangalam et al., 2023; Qiu et al., 2024), most focus on content under three minutes or can be solved with a few keyframes (Yu et al., 2019; Zhang et al., 2023b). More recent benchmarks targeting longer videos (Mangalam et al., 2023; Parmar et al., 2024; Wu et al., 2024; Hu et al., 2025; Shafique et al., 2025), often rely on general-purpose video content making them suitable for assessing perception and reasoning capabilities over extended sequences, but less effective for assessing narrative understanding. Other approaches that also focus on long videos primarily target localized "needle-in-a-haystack" retrieval rather than core narrative elements (Tan et al., 2024; Zohar et al., 2025).

(Movie) Narrative Understanding. Narratives have emerged as an important aspect of long-context evaluation in textual domains, due to their natural temporal and causal structure and rich information-dense content (Xu et al., 2024b; Bonomo et al., 2025; Hamilton et al., 2026; Karpinska et al., 2024). They provide a natural setting for evaluating whether models can integrate information across extended contexts, while maintaining coherence over long event sequences. Similarly, in the video domain, narrative understanding has been studied using movies and TV shows, which naturally contain structured and temporally extended storylines. Papalampidi et al. (2019, 2020) introduce the task of turning point identification in movies. Other works build benchmarks by sourcing movies (Huang et al., 2020; Song et al., 2023, 2024; Rawal et al., 2024), however they are typically limited in temporal coverage. More recent benchmarks (Ataallah et al., 2025) scale to longer content but rely on automatic question generation and treat all movie content uniformly, often producing questions unrelated to the central storyline (e.g., "needle-in-a-haystack" queries). In contrast, our benchmark is designed with a focus on *memorability*, targeting core narrative events that humans can recall without rewatching the movie. This is intended to better reflect what is salient in human memory and narrative recall. Beyond that, all movies are open-licensed, ensuring a fully reproducible setting and supporting future research efforts, unlike prior datasets that rely on copyright-restricted content or recycle the same sources.

6 Conclusions

In this work, we evaluated state-of-the-art models on movie narrative understanding as a long-context multimodal setting. To support this analysis, we introduced MF^2, a fully manually curated benchmark built from open-licensed full-length movies, with emphasis on memorable and storyline-central information that existing long-context evaluations often overlook. Our experiments reveal a substantial gap between human and model performance.

While subtitles provide a strong signal for narrative understanding and visual cues offer complementary information, current VLMs still struggle to effectively leverage and integrate such information in long multimodal contexts. More broadly, our findings suggest that brute-force encoding of visual content remains insufficient for narrative understanding, as models fail to identify, prioritize, and retain the most relevant information across extended videos, unlike humans. We hope MF^2 will encourage future work on models that can selectively consolidate and reason over salient narrative information while discarding irrelevant details.

Limitations

Despite careful design and validation, our dataset is not free from imperfections. Minor issues such as typos may remain, and annotators—although shown one claim at a time—may have recalled earlier claims from the same pair, potentially influencing subsequent judgments. Models do not share this limitation, as they process claims independently. As future work, claims from each pair could be split into disjoint sets and rated by different annotators to better isolate such effects. Furthermore, although our collection spans a broad variety of genres, most films primarily reflect Western cultural contexts, resulting in limited cultural diversity. While capturing diverse cultural perspectives was not the primary goal of this benchmark and lies outside the scope of this work—given that it is likely orthogonal to aspects of narrative understanding—future work could extend the dataset to include films from a wider range of cultural contexts. Finally, the human and model evaluations differ in their available input modalities, as humans have access to audio while models are evaluated using visual inputs and subtitles. Although some VLMs support audio input, current models are not yet capable of efficiently processing long-form video with all three modalities jointly, limiting the feasibility of a fully matched comparison.

Ethical considerations

We adhered to established scientific and ethical standards in constructing and releasing MF^2. All source movies are released under the permissive Public Domain 1.0 license. Claims and annotations were created and validated exclusively by the authors; no external crowdworkers were employed. To encourage a plurality of perspectives in the annotation process, the annotation team consists of individuals from diverse demographic, institutional, and geographic backgrounds. Since MF^2 is derived entirely from fictional movies, it contains no personally identifiable information (PII) of real individuals. Nonetheless, some fictional content may reflect cultural stereotypes or outdated social norms. We caution researchers that models evaluated on MF^2 may inherit such biases, and we recommend appropriate safeguards when interpreting or deploying results. We advise users to employ MF^2 strictly within the scope of this work, namely as a benchmark for evaluating vision–language models on movie narrative understanding, and discourage its use beyond this intended scope.

Acknowledgements

We gratefully acknowledge Miguel Graça, Evan Paces-Wiles, Maya Nachesa, Daniil Larionov, Bryan Sukidi, and José Pombal for their participation in the human evaluation process. This work was supported by the project DECOLLAGE (ERC-2022-CoG 101088763), by a research grant (VIL53122) from VILLUM FONDEN, by the Pioneer Center for AI DNRF grant number P1, by the NNF project 0094281, by FCT/MECI through national funds and when applicable co-funded EU funds under UID/50008: Instituto de Telecomunicações, by DARPA ECOLE Program No. HR00112390060, NSF-AI Engage Institute DRL-2112635, ARO Award W911NF2110220, ONR Grant N00014-23-1-2356, the Microsoft Accelerate Foundation Models Research (AFMR) grant program, by the project PID2023-151351NB-I00 funded by MCIN/ AEI /10.13039/501100011033 and by ERDF, UE. This collaboration resulted from an ELLIS workshop at MFO, the Oberwolfach Research Institute for Mathematics in the German Black Forest. The event was funded by the state of Baden-Württemberg (Germany) and organised in collaboration with the ELLIS Institute Tübingen and the Max Planck Institute for Intelligent Systems.

References

Marah I Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat S. Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, Vishrav Chaudhary, Parul Chopra, and 68 others. 2024. Phi-3

technical report: A highly capable language model locally on your phone. *CoRR*, abs/2404.14219.

Chenxin An, Shansan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. 2024. L-eval: Instituting standardized evaluation for long context language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14388–14411, Bangkok, Thailand. Association for Computational Linguistics.

Xiang An, Yin Xie, Kaicheng Yang, Wenkang Zhang, Xiuwei Zhao, Zheng Cheng, Yirui Wang, Songcen Xu, Changrui Chen, Chunsheng Wu, Huajie Tan, Chunyuan Li, Jing Yang, Jie Yu, Xiyao Wang, Bin Qin, Yumeng Wang, Zizhen Yan, Ziyong Feng, and 3 others. 2025. Llava-onevision-1.5: Fully open framework for democratized multimodal training. In *arXiv*.

Kirolos Ataallah, Eslam Mohamed Bakr, Mahmoud Ahmed, Chenhui Gou, Khushbu Pahwa, Jian Ding, and Mohamed Elhoseiny. 2025. InfiniBench: A benchmark for large multi-modal models in long-form movies and TV shows. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 19485–19512, Suzhou, China. Association for Computational Linguistics.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025a. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Ivana Balazevic, Yuge Shi, Pinelopi Papalampidi, Rahma Chaabouni, Skanda Koppula, and Olivier J Henaff. 2024. Memory consolidation enables long-context video understanding. In *Forty-first International Conference on Machine Learning*.

Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, Alexander Koller, Andre Martins, Philipp Mondorf, Vera Neplenbroek, Sandro Pezzelle, Barbara Plank, David Schlangen, Alessandro Suglia, Aditya K Surikuchi, Ece Takmaz, and Alberto Testoni. 2025. LLMs instead of human judges? a large scale empirical study across 20 NLP evaluation tasks. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 238–255, Vienna, Austria. Association for Computational Linguistics.

Tommaso Bonomo, Luca Gioffré, and Roberto Navigli. 2025. LiteraryQA: Towards effective evaluation of long-document narrative QA. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 34086–34107, Suzhou, China. Association for Computational Linguistics.

Keshigeyan Chandrasegaran, Agrim Gupta, Lea M. Hadzic, Taran Kota, Jimming He, Cristobal Eyzaguirre, Zane Durante, Manling Li, Jiajun Wu, and Fei-Fei Li. 2024. Hourvideo: 1-hour video-language understanding. In *Advances in Neural Information Processing Systems*, volume 37.

Yukang Chen, Wei Huang, Baifeng Shi, Qinghao Hu, Hanrong Ye, Ligeng Zhu, Zhijian Liu, Pavlo Molchanov, Jan Kautz, Xiaojuan Qi, Sifei Liu, Hongxu Yin, Yao Lu, and Song Han. 2025. Scaling rl to long videos. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Christopher Clark, Jieyu Zhang, Zixian Ma, Jae Sung Park, Mohammadreza Salehi, Rohun Tripathi, Sangho Lee, Zhongzheng Ren, Chris Dongjoo Kim, Yinuo Yang, Vincent Shao, Yue Yang, Weikai Huang, Ziqi Gao, Taira Anderson, Jianrui Zhang, Jitesh Jain, George Stoica, Winson Han, and 2 others. 2026. Molmo2: Open weights and data for vision-language models with video understanding and grounding. *Preprint*, arXiv:2601.10611.

Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, YenSung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, and 32 others. 2024. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models. *arXiv preprint arXiv:2409.17146*.

Gueter Josmy Faure, Min-Hung Chen, Jia-Fong Yeh, Ying Cheng, Hung-Ting Su, Yung-Hao Tang, Shang-Hong Lai, and Winston H. Hsu. 2025. MovieCORE: COgnitive REasoning in movies. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 1253–1272, Suzhou, China. Association for Computational Linguistics.

Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, and 1 others. 2025a. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *CVPR*.

Stephanie Fu, tyler bonnen, Devin Guillory, and Trevor Darrell. 2025b. Hidden in plain sight: VLMs overlook their visual representations. In *Second Conference on Language Modeling*.

Sil Hamilton, Matthew Wilkens, and Andrew Piper. 2026. NarraBench: A comprehensive framework for narrative benchmarking. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1:*

Long Papers), pages 3786–3801, Rabat, Morocco. Association for Computational Linguistics.

Kairui Hu, Penghao Wu, Fanyi Pu, Wang Xiao, Yuanhan Zhang, Xiang Yue, Bo Li, and Ziwei Liu. 2025. Video-mmmu: Evaluating knowledge acquisition from multi-discipline professional videos. *arXiv preprint arXiv:2501.13826*.

Qingqiu Huang, Yu Xiong, Anyi Rao, Jiaze Wang, and Dahua Lin. 2020. Movienet: A holistic dataset for movie understanding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 709–727. Springer.

George Hudson and Noura Al Moubayed. 2022. MuLD: The multitask long document benchmark. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3675–3685, Marseille, France. European Language Resources Association.

Marzena Karpinska, Katherine Thai, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2024. One thousand and one pairs: A "novel" challenge for long-context language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17048–17085, Miami, Florida, USA. Association for Computational Linguistics.

Dongxu Li, Yudong Liu, Haoning Wu, Yue Wang, Zhiqi Shen, Bowen Qu, Xinyao Niu, Fan Zhou, Chengen Huang, Yanpeng Li, Chongyan Zhu, Xiaoyi Ren, Chao Li, Yifan Ye, Peng Liu, Lihuan Zhang, Hanshu Yan, Guoyin Wang, Bei Chen, and Junnan Li. 2025. Aria: An open multimodal native mixture-of-experts model. *Preprint*, arXiv:2410.05993.

KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2023a. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*.

Ruizhe Li and Yanjun Gao. 2024. Anchored answers: Unravelling positional bias in gpt-2's multiple-choice questions. *arXiv preprint arXiv:2405.03205*.

Yanwei Li, Chengyao Wang, and Jiaya Jia. 2023b. Llama-vid: An image is worth 2 tokens in large language models. *Preprint*, arXiv:2311.17043.

Hao Liu, Wilson Yan, Matei Zaharia, and Pieter Abbeel. 2025. World model on million-length video and language with blockwise ringattention. In *The Thirteenth International Conference on Learning Representations*.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning.

Jiajun Liu, Yibing Wang, Hanghang Ma, Xiaoping Wu, Xiaoqi Ma, xiaoming Wei, Jianbin Jiao, Enhua Wu, and Jie Hu. 2024a. Kangaroo: A powerful video-language model supporting long-context video input. *arXiv preprint arXiv:2408.15542*.

Ming Liu and Wensheng Zhang. 2025. Is your video language model a reliable judge? In *The Thirteenth International Conference on Learning Representations*.

Zhijian Liu, Ligeng Zhu, Baifeng Shi, Zhuoyang Zhang, Yuming Lou, Shang Yang, Haocheng Xi, Shiyi Cao, Yuxian Gu, Dacheng Li, Xiuyu Li, Yunhao Fang, Yukang Chen, Cheng-Yu Hsieh, De-An Huang, An-Chieh Cheng, Vishwesh Nath, Jinyi Hu, Sifei Liu, and 8 others. 2024b. Nvila: Efficient frontier visual language models. *Preprint*, arXiv:2412.04468.

Olga Loginova, Oleksandr Bezrukov, and Alexey Kravets. 2024. Addressing blind guessing: Calibration of selection bias in multiple-choice question answering by video language models. *arXiv preprint arXiv:2410.14248*.

Shiyin Lu, Yang Li, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and Han-Jia Ye. 2024. Ovis: Structural Embedding Alignment for Multimodal Large Language Model. *arXiv e-prints*, arXiv:2405.20797.

Ruipu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Minghui Qiu, Pengcheng Lu, Tao Wang, and Zhongyu Wei. 2023. Valley: Video assistant with large language model enhanced ability. *Preprint*, arXiv:2306.07207.

Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. 2024. Video-chatgpt: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*.

Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. 2023. Egoschema: A diagnostic benchmark for very long-form video language understanding. *arXiv preprint arXiv:2308.09126*.

Microsoft, :, Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, Dong Chen, Dongdong Chen, Junkun Chen, Weizhu Chen, Yen-Chun Chen, Yi ling Chen, Qi Dai, and 57 others. 2025. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *Preprint*, arXiv:2503.01743.

Francesco Maria Molfese, Luca Moroni, Luca Gioffré, Alessandro Scirè, Simone Conia, and Roberto Navigli. 2025. Right answer, wrong score: Uncovering the inconsistencies of LLM evaluation in multiple-choice question answering. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 18477–18494, Vienna, Austria. Association for Computational Linguistics.

OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb,

Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, and 401 others. 2024. Gpt-4o system card. *Preprint*, arXiv:2410.21276.

Pinelopi Papalampidi, Frank Keller, Lea Frermann, and Mirella Lapata. 2020. Screenplay summarization using latent narrative structure. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1920–1933, Online. Association for Computational Linguistics.

Pinelopi Papalampidi, Frank Keller, and Mirella Lapata. 2019. Movie plot analysis via turning point identification. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1707–1717, Hong Kong, China. Association for Computational Linguistics.

Paritosh Parmar, Eric Peh, Ruirui Chen, Ting En Lam, Yuhan Chen, Elston Tan, and Basura Fernando. 2024. Causalchaos! dataset for comprehensive causal action question answering over longer causal chains grounded in dynamic visual scenes. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Andrew Piper and Sunyam Bagga. 2024. Using large language models for understanding narrative discourse. In *Proceedings of the 6th Workshop on Narrative Understanding*, pages 37–46, Miami, Florida, USA. Association for Computational Linguistics.

Lu Qiu, Yi Chen, Yuying Ge, Yixiao Ge, Ying Shan, and Xihui Liu. 2024. Egoplan-bench2: A benchmark for multimodal large language model planning in real-world scenarios. *arXiv preprint arXiv:2412.04447*.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR.

Ruchit Rawal, Khalid Saifullah, Miquel Farré, Ronen Basri, David Jacobs, Gowthami Somepalli, and Tom Goldstein. 2024. Cinepile: A long video question answering dataset and benchmark. *arXiv preprint arXiv:2405.08813*.

Saul Santos, António Farinhas, Daniel C McNamee, and Andre Martins. 2025. ∞-video: A training-free approach to long video understanding via continuous-time memory consolidation. In *Forty-second International Conference on Machine Learning*.

Bhuiyan Sanjid Shafique, Ashmal Vayani, Muhammad Maaz, Hanoona Abdul Rasheed, Dinura Dissanayake, Mohammed Irfan Kurpath, Yahya Hmaiti, Go Inoue, Jean Lahoud, Md. Safirur Rashid, Shadid Intisar Quasem, Maheen Fatima, Franco Vidal, Mykola Maslych, Ketan Pravin More, Sanoojan Baliah, Hasindri Watawana, Yuhao Li, Fabian Farestam, and 10 others. 2025. A culturally-diverse multilingual multimodal video benchmark & model. *Preprint*, arXiv:2506.07032.

Shrutika Singh, Anton Alyakin, Daniel Alexander Alber, Jaden Stryker, Ai Phuong S Tong, Karl Sangwon, Nicolas Goff, Mathew de la Paz, Miguel Hernandez-Rovira, Ki Yun Park, and 1 others. 2025. It is too many options: Pitfalls of multiple-choice questions in generative ai and medical education. *arXiv preprint arXiv:2503.13508*.

Enxin Song, Wenhao Chai, Guanhong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Xun Guo, Tian Ye, Yan Lu, Jenq-Neng Hwang, and 1 others. 2023. Moviechat: From dense token to sparse memory for long video understanding. *arXiv preprint arXiv:2307.16449*.

Enxin Song, Wenhao Chai, Tian Ye, Jenq-Neng Hwang, Xi Li, and Gaoang Wang. 2024. Moviechat+: Question-aware sparse memory for long video question answering. *arXiv preprint arXiv:2404.17176*.

Chaolei Tan, Zihang Lin, Junfu Pu, Zhongang Qi, Wei-Yi Pei, Zhi Qu, Yexin Wang, Ying Shan, Wei-Shi Zheng, and Jian-Fang Hu. 2024. Synopground: A large-scale dataset for multi-paragraph video grounding from tv dramas and synopses. *Preprint*, arXiv:2408.01669.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. Gemma 3 technical report. *Preprint*, arXiv:2503.19786.

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 74 others. 2026. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *Preprint*, arXiv:2507.01006.

Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin,

Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier Hénaff, Jeremiah Harmsen, Andreas Steiner, and Xiaohua Zhai. 2025. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*.

Weihan Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Ming Ding, Xiaotao Gu, Shiyu Huang, Bin Xu, and 1 others. 2025a. Lvbench: An extreme long video understanding benchmark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22958–22967.

Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, and 56 others. 2025b. Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *Preprint*, arXiv:2508.18265.

Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. 2024a. Videoagent: Long-form video understanding with large language model as agent. In *European Conference on Computer Vision*, pages 58–76. Springer.

Ziyang Wang, Shoubin Yu, Elias Stengel-Eskin, Jaehong Yoon, Feng Cheng, Gedas Bertasius, and Mohit Bansal. 2024b. Videotree: Adaptive tree-based video representation for llm reasoning on long videos. *arXiv preprint arXiv:2405.19209*.

Bo Wu, Shoubin Yu, Zhenfang Chen, Joshua B. Tenenbaum, and Chuang Gan. 2021. STAR: A benchmark for situated reasoning in real-world videos. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.

Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. 2024. Longvideobench: A benchmark for long-context interleaved video-language understanding. In *Advances in Neural Information Processing Systems*, volume 37, pages 28828–28857. Curran Associates, Inc.

Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. 2021. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9777–9786.

LLM-Core-Team Xiaomi. 2025. Mimo-vl technical report. *Preprint*, arXiv:2506.03569.

Lin Xu, Yilin Zhao, Daquan Zhou, Zhijie Lin, See Kiong Ng, and Jiashi Feng. 2024a. Pllava : Parameter-free llava extension from images to videos for video dense captioning. *Preprint*, arXiv:2404.16994.

Zhe Xu, Jiasheng Ye, Xiaoran Liu, Xiangyang Liu, Tianxiang Sun, Zhigeng Liu, Qipeng Guo, Linlin Li, Qun Liu, Xuanjing Huang, and 1 others. 2024b. Detectiveqa: Evaluating long-context reasoning on detective novels. *arXiv preprint arXiv:2409.02465*.

Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. 2021. Just ask: Learning to answer questions from millions of narrated videos. In *ICCV*.

Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. 2025a. mPLUG-owl3: Towards long image-sequence understanding in multi-modal large language models. In *The Thirteenth International Conference on Learning Representations*.

Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh V Chawla, and Xiangliang Zhang. 2025b. Justice or prejudice? quantifying biases in LLM-as-a-judge. In *The Thirteenth International Conference on Learning Representations*.

Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. 2019. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *AAAI*, pages 9127–9134.

Boqiang Zhang, Kehan Li, Zesen Cheng, Zhiqiang Hu, Yuqian Yuan, Guanzheng Chen, Sicong Leng, Yuming Jiang, Hang Zhang, Xin Li, Peng Jin, Wenqi Zhang, Fan Wang, Lidong Bing, and Deli Zhao. 2025. Videollama 3: Frontier multimodal foundation models for image and video understanding. *arXiv preprint arXiv:2501.13106*.

Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, and Gedas Bertasius. 2024. A simple llm framework for long-range video question-answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21715–21737.

Hang Zhang, Xin Li, and Lidong Bing. 2023a. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*.

Hongjie Zhang, Yi Liu, Lu Dong, Yifei Huang, Zhen-Hua Ling, Yali Wang, Limin Wang, and Yu Qiao. 2023b. Movqa: A benchmark of versatile question-answering for long-form movie understanding. *Preprint*, arXiv:2312.04817.

Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2024. Video Instruction Tuning With Synthetic Data. *arXiv e-prints*, arXiv:2410.02713.

Xu Zheng, Chenfei Liao, Yuqian Fu, Kaiyu Lei, Yuanhuiyi Lyu, Lutao Jiang, Bin Ren, Jialei Chen, Jiawen Wang, Chengxin Li, Linfeng Zhang, Danda Pani Paudel, Xuanjing Huang, Yu-Gang Jiang, Nicu Sebe, Dacheng Tao, Luc Van Gool, and Xuming Hu. 2025. Mllms are deeply affected by modality bias. *Preprint*, arXiv:2505.18657.

Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, and 32 others. 2025. InternVL3: Exploring Advanced Training and Test-Time Recipes for Open-Source Multimodal Models. *arXiv e-prints*, arXiv:2504.10479.

Orr Zohar, Rui Li, Andres Marafioti, Xiaohan Wang, Stanford AI Team, and Hugging Face. 2025. Timescope: How long can your video large multimodal model go? https://huggingface.co/blog/timescope-video-lmm-benchmark.

A Data Annotation Details

Movie Collection Metadata We collected 53 full-length movies from the Internet Archive[7], focusing on titles released under the Public Domain 1.0 license to ensure legal reusability and support open-access research. Original-language subtitles, mostly in English, were obtained from OpenSubtitles.org[8]. For one movie without available subtitles, we used whisper-1 (Radford et al., 2023)[9] to generate a transcript and manually post-edited to ensure high quality. Each movie includes audio and aligned subtitles, with an average duration of 88 minutes. In Tables 8 and 9, we provide detailed information on the 53 released movies, including their genre, original language, and duration. Collectively, the movies span a wide temporal range of approximately five decades (1920–1970) and cover a diverse set of genres.

Annotator Demographics All annotations were performed manually by 26 co-authors from 12 institutions across 7 countries, with a roughly balanced distribution of career stages (PhD students, postdocs, and faculty members) and gender. We did not rely on crowdsourcing or semi-automatic approaches, as constructing claims about central narrative events—often requiring reasoning across multiple scenes—demands careful human judgment. Having co-authors perform the annotations also provided strong intrinsic motivation and accountability, ensuring high-quality, consistent work that might not be achievable with anonymous or paid crowdworkers. Detailed guidelines, multiple rounds of discussion, and collective revisions were used to ensure consistency and high quality, while annotators were not informed of any specific hypotheses or expected outcomes, reducing the risk of bias.

Annotated Claim Pairs Figure 4 shows the number of annotated claim pairs per movie.

Comprehension Dimensions Each claim pair was labeled according to one or more comprehension dimensions, which capture specific aspects of narrative understanding. These dimensions are informed by prior work (Xiao et al., 2021; Zhang et al., 2023b; Wang et al., 2025a) and are defined in Table 4.

[7] https://archive.org

[8] https://www.opensubtitles.org

[9] https://platform.openai.com/docs/models/whisper-1

B Pairwise Accuracy vs Simple Accuracy

Our evaluation protocol is motivated as follows: multiple-choice settings introduce distractor-driven biases, where models exploit superficial cues or distractor order (Molfese et al., 2025). Following prior work on narrative understanding (Karpinska et al., 2024), we evaluate each minimally differing claim independently, removing distractors and reducing such shortcut behaviors. Under this setup, pairwise accuracy is stricter and more informative; a model is counted as correct only if it predicts both labels accurately, minimizing the chance that it appears "correct for the wrong reason" (Karpinska et al., 2024). In contrast, in simple accuracy each prediction contributes independently to the score making it less suitable for our case and more suitable for multiple-choice formats, which we intentionally avoid for the aforementioned reasons. In our evaluation, the relationship between the two predictions is essential. **The model must implicitly distinguish what truly happened from a minimally edited alternative that is also plausible, even though it never sees the two statements jointly during inference.** That said, the evaluation depends on the relationship between the two predictions, not the predictions alone.

As a further illustration of why interpreting results based only on simple accuracy can be misleading, we report in Table 5 several metrics (pairwise accuracy, simple accuracy, accuracy on facts and accuracy on fibs) for some of the open-weight models and for Gemini 2.5 Pro, across both the video-only and video with subtitles settings. *Simple accuracy can appear high simply because the model performs well on the fibs, while failing on the facts*, highlighting a bias towards predicting False. Taken together **these considerations motivate our choice of pairwise accuracy as the primary evaluation metric**. Nevertheless, we also report accuracy for completeness.

C Mid and Large scale Model Analysis

Beyond vision and subtitles: the role of summarized content (synopsis) and pretrained (world) knowledge. Table 6 reports an ablation study examining the extent to which models rely on textual information beyond subtitles. When conditioned only on the movie title and release year, performance remains at or only slightly above chance level, indicating that parametric knowledge acquired during pretraining contributes only

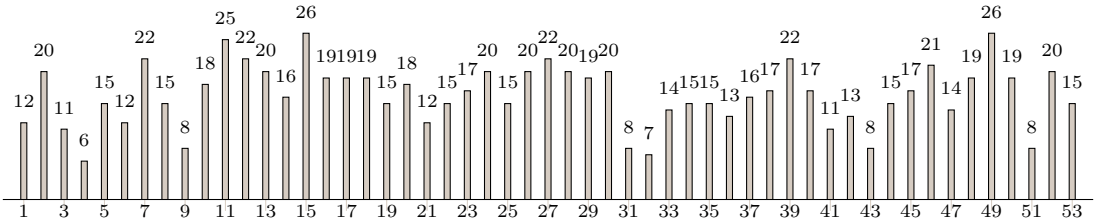


Figure 4: Number of claim pairs per movie.

Comprehension Dimension	Definition
Event/Entity Understanding	Involves identifying key entities (e.g., people, places, or objects) and understanding the events they participate in. This includes tracking entities across scenes, interpreting their roles, and recognizing their interactions and relationships throughout the narrative.
Temporal Perception	Requires reasoning about the timeline of events—determining whether actions occur before, after, or simultaneously—and may also include counting or sequencing events. The focus is on broader temporal relationships within the narrative.
Emotion Understanding	Involves recognizing the emotional states of characters and interpreting how these emotions evolve throughout the story.
Causal Reasoning	Focuses on identifying cause-and-effect relationships between events or actions, including both explicit and implicit dependencies that may span multiple scenes.

Table 4: Definitions of comprehension dimensions.

marginally to the task, unlike the stronger effect observed for Gemini 2.5 Pro as reported in the main text. Likewise, providing only the movie synopsis results to near-random performance for most models, suggesting that high-level textual summaries alone are insufficient for reliable prediction. In both settings, performance remains far below human level, indicating that *neither memorized knowledge nor condensed narrative summaries are sufficient to successfully solve the task*.

A fine-grained view of large-scale model performance across comprehension dimensions and reasoning granularities. Fig. 5 shows that, among the large-scale models, Gemini 2.5 Pro still demonstrates inferior performance, ranking second to humans in various categories. Other models like LLaVA-Video and InternVL3 generally show lower scores, suggesting areas for improvement. The results also highlight varying degrees of difficulty across the tasks, with emotion comprehension appearing to be a strong point for humans and models, while temporal perception is the weakest for both. Interestingly, the analysis on reasoning granularity reveals an interesting pattern between humans and models: as reasoning shifts from single-scene to multi-scene/global, model performance tends to decrease for most models, while human performance remains stable.

D Detailed Guidelines for Data Annotation and Human Evaluation

D.1 Data Annotation Guidelines

Figures 6 and 7 present the detailed guidelines provided to annotators during the data annotation process. These include instructions for constructing contrastive claim pairs, and labeling each pair with the appropriate reasoning granularity and comprehension dimensions. Furthermore, in Figures 8 and 9, we include a subset of illustrative examples shown to annotators to guide their annotations of reasoning granularity and comprehension dimensions, respectively.

We note that among the comprehension dimensions annotators could assign to each claim pair, an "Other" category was included to account for cases that did not clearly align with any of the prede-

Model	Pairwise Acc	Average Acc	Acc on Facts	Acc on Fibs
		Video		
Gemini 2.5 Pro	37.9	64.4	47.4	81.5
Gemma3-27B	31.5	61.2	57.4	65.1
Qwen3VL-IT-30B	19.2	56.9	24.3	89.5
InternVL3.5-38B	26.6	60.0	38.8	81.1
Qwen2.5VL-72B	29.7	58.8	45.4	72.1
InternVL3-78B	22.1	58.0	37.7	78.3
		Video+Subtitles		
Gemini 2.5 Pro	63.8	79.8	72.2	87.3
Gemma3-27B	39.7	65.3	57.6	73.0
Qwen3VL-IT-30B	43.8	69.5	53.0	86.0
InternVL3.5-38B	46.2	70.3	61.6	79.0
Qwen2.5VL-72B	47.5	69.0	59.3	78.7
InternVL3-78B	51.3	72.7	67.5	77.9

Table 5: Reporting pairwise accuracy, average accuracy, accuracy on facts and accuracy on fibs for some open-weight models and Gemini2.5. Results highlight that simple accuracy on its own can be misleading.

Model	Movie Title	Synopsis	Subs	Video	Video+Subs
Gemini 2.5 Pro	45.4	27.9	58.9	37.9	63.8
Gemma3-27B	28.1	29.6	42.3	31.5	39.7
Qwen3VL-IT-30B	12.9	26.6	33.5	19.2	43.8
InternVL3.5-38B	27.0	31.8	47.5	26.6	46.2
Qwen2.5-VL-72B	15.2	25.7	37.6	29.7	47.5
LLaVA-Video-72B	17.1	38.7	41.7	15.6	41.8
InternVL3-78B	30.4	30.4	49.0	22.1	51.3

Table 6: Extended version of Table 3 reporting pairwise performance for Gemini 2.5 Pro and mid- to large-scale open-weight models across different input modalities. *Video* uses only the visual stream; *Subs* includes only subtitle information; *Synopsis* relies only on the synopsis of the movie obtained from Wikipedia; *Video+Subs* combines both visual and subtitles inputs; and *Movie Title* uses only the claim, along with the movie title and release year, without access to movie content.

fined dimensions. As this label was selected rarely (0.49% of the data), it is excluded from the figures presented in the main text. Additionally, although annotators were instructed to distinguish between single-scene, multi-scene, and global reasoning granularities, annotator feedback indicated that the distinction between multi-scene and global reasoning was not always clear-cut in practice, with substantial overlap between the two categories. Based on this feedback, we therefore merge them into a single category (multi-scene/global) in the analysis reported in the main text.

Consent and Data Use. We confirm that all contributors whose data were used in this project provided informed consent prior to participation. The data consist solely of contributions created for benchmark construction; no personal data were collected. Contributors were informed that their contributions would be released under a CC-BY 4.0 license.

D.2 Human Evaluation Guidelines

In Figures 10 and 11, we provide the full set of guidelines shared to participants during the human evaluation process, which consists of two stages: an initial stage in which evaluators respond without revisiting the movie, and an optional second stage that allows revisiting. While we only analyze the results from Stage 1—as our goal is to assess movie understanding based on memorable events without allowing participants to rewatch parts of the film—

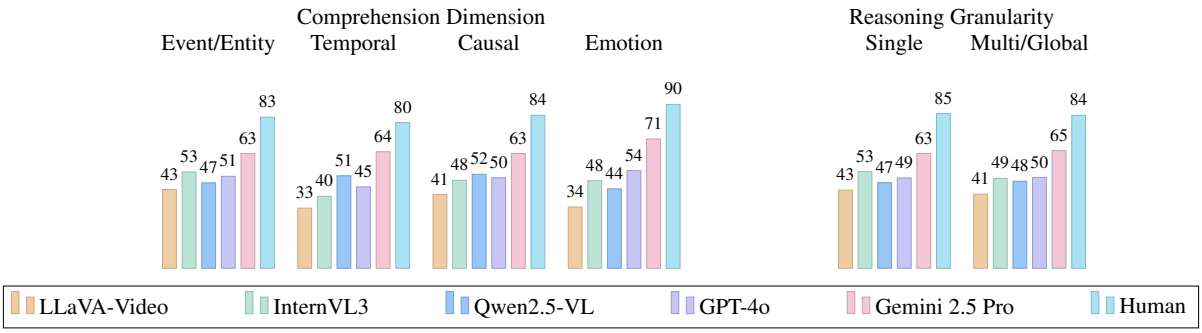


Figure 5: Pairwise accuracy for large-scale models with video and subtitles, and human baseline per comprehension dimension and reasoning granularity.

we include the complete instructions for both stages to offer full context. Additionally, we provide an illustration of the evaluation interface to clarify the evaluation setup.

E Details on Experimental Setup

E.1 Frame selection

For all models except GPT-4o, we first downsample videos to 1 frame per second. From these frames, we then uniformly sample a subset, adjusting the number of frames based on each model's maximum context length and original training settings. For GPT-4o, frames are uniformly sampled directly from the original videos without prior downsampling. The exact number of frames sampled for each model is reported in Table 7.

E.2 Prompt Templates & Answer Extraction

In Figures 12 and 13 we present the direct and explanation prompt templates used for open-weight and closed models, respectively. The former requests only a True/False response, while the latter additionally asks for a brief justification before the final answer. We found that the direct prompt yielded better performance for open-weight models, while the explanation prompt proved more effective for closed models. When experimenting with different input modalities—such as adding the synopsis, subtitles, or movie title—we adapt the prompts accordingly. To extract predictions, we use regular expressions to identify True/False answers in the model outputs, selecting the last valid match.

E.3 Resources

Our infrastructure consists of a single machine equipped with 2 NVIDIA H200 GPUs (140GB each) and 12 Intel Xeon Gold 6348 CPUs (2.60GHz, 1TB RAM). Most experiments were conducted on a single GPU, except for evaluations involving mid and large scale open-weight models, where all 2 GPUs were used to accelerate inference.

F Illustrative Examples

In this part, we include a small set of examples (Tables 10, 11, 12, 13, 15, 16, 17, 18, 19, 20) showing when visual or text cues are needed to resolve the claims and how Gemini 2.5 Pro (best performing model) succeeds or fails on them, along with its corresponding explanations.

AI Assistance

We would like to note that large language models (ChatGPT) were used to assist in drafting and polishing the writing of this work.

Guidelines for Data Annotation (Part 1)

We are conducting a research study on long movie understanding as part of a broader effort to explore how well viewers comprehend and recall complex narratives. Your task is to create claims that test a viewer's comprehension of a movie after watching it. These claims will be used in a human evaluation study to assess how well participants understand and recall key events from the movie. We appreciate your participation in this data collection process.

General Task Instructions Select a movie from the current "Pool" of movies (the "Pool" can be found in <LINK>). Make sure this movie is not selected by another annotator.

- Watch the entire movie carefully.
- We highly recommend reading the example claims provided to gain a better understanding of the task you need to fulfil.
- Start writing down your claims following the template available in <LINK> (you will find two tabs available: the "Examples" tab contains claim examples, and the "Annotations Template" tab is the template you should follow). Please create another sheet with your claims–do not directly use the current template–and send it to us once it is completed.

Annotation Process

1. Writing Claims You are asked to create pairs of contrastive claims, where one claim is true (fact) and the counterfactual version is false (fib). The two claims should differ by minimal edits, meaning they should be as similar as possible while maintaining contrast. Each claim should differ in a subtle but meaningful way, challenging comprehension without being overly obvious.
Example:
Fact: The first bomb exploded in the bus.
Fib: The first bomb exploded in the aquarium.
Why this works: The counterfactual claim is created with minimal edits, maintaining contrast while testing the understanding of a key event.

2. Select Claim Granularity For each pair of claims you constructed, indicate whether answering them correctly requires reasoning based on a single scene, multiple scenes, or globally within the movie.
Definition of scene:
A scene in film refers to a complete unit of storytelling, usually consisting of a sequence of events and dialogue taking place in a specific location and time. It often involves one or more characters and is usually shot in one continuous take or consisting of a sequence of shots.
Reasoning Granularity Labels:

- **Single-scene:** Claims that are answerable using information from a single scene.
- **Multi-scene:** Claims falling into this granularity require information/evidence from multiple distinct scenes, but not from the whole film. In this case, details are usually spread out between the multiple scenes. The supporting information/evidence is distributed, but explicit and locatable (timestamps/scenes can be clearly identified and referenced)
- **Global:** Claims falling into this granularity require a holistic understanding of the movie narrative. They cannot be easily tied to specific scenes or timestamps, and need to infer or accumulate information/evidence that emerges across the entire narrative (timestamps/scenes can not be clearly identified and referenced).

Note: Reasoning granularity labels should be selected based on the fact (true claim). Check the examples provided in the "Examples for Reasoning Granularity" part.

Figure 6: Guidelines provided for the data annotation procedure (Part 1).

Guidelines for Data Annotation (Part 2)

3. Claim Categorization Identify the comprehension dimensions the constructed pair of claims examines. Sometimes more than one dimension is examined, so we allow for multiple labels.

Comprehension Dimension Labels:

- **Event/Entity Understanding:** it refers to claims that require the identification of key entities (such as people, places, or objects) and understanding of actions or events involving those entities throughout the narrative. Understanding these claims involves tracking the presence and role of entities across scenes, extracting relationships among them, observing and interpreting their actions, and linking them to relevant events in the narrative.

- **Temporal Perception:** temporal perception refers to claims that require understanding of the timeline of events. It involves reasoning about the order in which events or actions occur—e.g., determining whether an event/action takes place before, after or at the same time as another—and may also require counting the number of specific actions or events. Unlike tasks focused on localizing a specific action in time, temporal perception emphasizes comprehension of broader temporal relationships within the evolving storyline.

- **Emotion Understanding:** emotional understanding refers to claims that involve recognizing and interpreting the emotional development of characters throughout the narrative.

- **Causal Reasoning:** causal reasoning refers to claims that require identifying cause-and-effect relationships between events or actions, where the relationship may be either direct or implicit.

- **Other:** If none of the above fit, select "Other" and suggest a new category.

Note: The categorization is based on both claims (fact and fib). Check the examples provided in the "Examples for Comprehension Dimensions" part.

Important Points To Consider

- **Ensure claims assess the viewer's understanding of the movie**. To put it simply, claims should refer to **significant moments** in the movie, **avoiding trivial details or Needle in a Haystack (NIAH)-style claims**, such as: "The detective wears a red T-shirt" (if this detail is not important in the movie).

- **Claims must be clear and unambiguous in isolation**, meaning they should be understandable without requiring additional context but should still require reasoning based on the movie. **Each claim should be self-contained and make sense independently**, without referencing its counterfactual version. Also, **avoid highly subjective or interpretive claims**. Each claim should still have a definitive answer based on the movie's content.

- **Avoid providing unnecessary contextual details**. For example, do not use phrases like "in the beginning of the movie, . . . ", "in the final scene, . . . " unless such information is essential to understanding the claim.

- Ensure that claims focus on memorable and salient narrative content that can be recalled without rewatching the movie.

- Once you finish the annotation process, please **go through your claims and confirm that they are in line with the points raised above** (these points are important to be covered to ensure good quality of annotations).

Figure 7: Guidelines provided for the data annotation procedure (Part 2).

Guidelines for Data Annotation (Part 3)

Examples for Reasoning Granularity In this part, we provide examples to illustrate how to assign reasoning granularity labels.

Example 1:
Fact: According to the Hattley, the individual shown in the photograph (Marakelli) worked with Constain.
Fib: According to Hattley, the individual shown in the photograph (Marakelli) had no connection or working relationship with Constain.
Reasoning Granularity: Single-scene.
Justification: This event is categorized as single-scene because it takes place within one specific scene: Hattley shows the photograph to Conley, they are having a discussion and it is implied that Marakelli worked with Constain in the mafia.

Example 2:
Fact: Hattley appeared visibly bothered with the discussion he had in his office with Constain's attorney.
Fib: Hattley appeared pleased with the discussion he had in his office with Constain's attorney.
Reasoning Granularity: Single-scene.
Justification: That is again a single scene event. Constain's attorney enters the office and they are having a discussion. After a while, Hattley kicks him out.

Example 3:
Fact: Miss Conley received a dress as a personal gift from the policeman.
Fib: Miss Conley received a dress as a gift from the government, delivered by the policeman.
Reasoning Granularity: Multi-scene.
Justification: That is a multi-scene event, that we need to ground on 2 independent scenes to answer the question correctly. In the first scene Miss Conley receives a gift from the policeman, who says that the gift is from the government. After a while (some scenes are interleaved), she understands that the policeman bought the gift for her and not the government. So to answer correctly, we need to ground on these 2 specific scenes.

Example 4:
Fact: Conley's statement about her occupation, describing herself as a "gang buster," implicitly refers to Constain.
Fib: Conley's statement about her occupation, describing herself as a "gang buster," implicitly refers to Pete Tinelli.
Reasoning Granularity: Global
Justification: There is a single scene in the end of a movie during which Conley characterises herself as a "gang buster". Although it is a single scene, it is impossible to understand solely by this scene why she said it and to whom she is referring to. We need to watch a big part of the movie (if not all of it) to understand that refers to Constain.

Figure 8: Guidelines provided for the data annotation procedure (Part 3). This part of the guidelines provides examples given to annotators to illustrate how to assign reasoning granularity labels. While more examples were shared during the annotation process, we include a selection here for illustrative purposes.

Guidelines for Data Annotation (Part 4)

Examples for Comprehension Dimensions In this part we provide examples to illustrate how to assign comprehension dimension labels.

Example 1:
Fact: At Jim's bar, the Connel keeps drinking as he talks to the fake John Doe, expressing his frustration and concern.
Fib: At Jim's bar, the Connel keeps drinking as he talks to the fake John Doe, expressing hope and happiness.
Comprehension Dimension: emotion understanding
Justification: We need to understand what emotion Connel expressed, to answer the pair of claims correctly.

Example 2:
Fact: Conley's statement about her occupation describing herself as a "gang buster", implicitly refers to Constain.
Fib: Conley's statement about her occupation describing herself as a "gang buster", implicitly refers to Pete Tinelli.
Comprehension Dimension: entity/event understanding
Justification: We need to understand to whom the expression "gang buster" refers to. So, the comprehension dimension is entity understanding.

Example 3:
Fact: Hallet brought Conley's sister to the hotel with the intent to make Conley testify in the trial.
Fib: Hallet brought Conley's sister to the hotel with the intent to make her feel safe.
Comprehension Dimension: causal reasoning
Justification: Here we need to understand why Hallet brought Conley's sister to the hotel. So it examines a causal-and-effect relationship.

Example 4:
Fact: Conley decided to testify only after Wiloughby's death.
Fib: Conley had already decided to testify before Wiloughby's death.
Comprehension Dimension: temporal perception
Justification: that pair examines the temporal dimension (if the decision was taken before or after Wiloughby's death).

Figure 9: Guidelines provided for the data annotation procedure (Part 4). This part of the guidelines provides examples given to annotators to illustrate how to assign comprehension dimension labels. While more examples were shared during the annotation process, we include a selection here for illustrative purposes.

Guidelines for Human Evaluation (Part 1)

This evaluation study aims to assess how well people comprehend and recall key events from a movie. You will watch a movie and then evaluate a series of claims about its content. Your goal is to determine whether each claim is True or False, based solely on what was shown in the movie. We appreciate your participation in this study.

Task Instructions

- Assign to yourself the movies you want to watch and do the test (we expect 2 movies per person). Please add your name to the Human-Eval column, on this LINK.
- Visit the platform for evaluation LINK.
- Provide your email to receive access to the movie (it will be used as your unique identifier).
- Once you submit your email, you should carefully select from the drop-down list the corresponding movie you assigned yourself and proceed with the evaluation. You will be shown with the movie link. Please open it in a new tab.

The test is divided in **2 stages**: The **first stage** is **mandatory** and should be completed by everyone (*during this stage you are not allowed to go back to the movie while answering the questions*). The **second stage** is **optional** (*during this stage you are allowed to go back to the movie while answering the questions*).

Stage 1:

1. **Watch the entire movie carefully before proceeding to the evaluation**. Pay attention to details and context in the movie, as some claims may be subtle or require careful reasoning.
2. After watching, it's time to proceed to Stage 1. **Please do not go back to the movie until Stage 1 of the test is completed.** Press the "Start Classifying Claims" button, and you will be shown with **one claim at a time**. For each claim shown, you need to do the following:
    - **Classify the claim as True/False** (you should always answer truthfully, without aiming to maximise you score).
    - Mark your **confidence** about your answer. This is helpful for stage 2, where you will have the opportunity to revise your claims (by looking back at the movie).
    - Leave a comment if any of the following applies: If a claim is **ambiguous, unclear, open to interpretation, has a bad phrasing or typos, you may leave an optional comment explaining your concerns**. You can also comment on the claim in case it is **needle-in-a-haystack style** and you think it is too detailed and doesn't test the understanding of the movie.
    - Once you answered, click "Save" to submit your response and move on to the next claim.

Important details: Once you submit an answer, you cannot go back and change it. At this stage, you are **strictly prohibited from searching back in the movie, rewinding, or rewatching scenes while answering the claims**. Your responses **should be based on your memory** and **understanding**. You must **NOT use any AI tools or external sources to verify or generate answers**. The goal of this study is to assess human understanding of long movies, not automated retrieval or AI-assisted responses. Also you are not allowed to take any paper notes, while watching the movie.

Figure 10: Guidelines provided for human evaluation (Part 1).

Guidelines for Human Evaluation (Part 2)

Stage 2:

Once you complete Stage 1, you will see a message asking you if you want to proceed to Stage 2 (Stage 2 is optional).

During Stage 2, you will be shown again with the choices you selected during Stage 1, but now **you can revise your answers by looking back to the movie** (you can reuse the movie link we provided you). You will be shown for each claim with the choices you did in Stage 1. You are free to change them and proceed to the next claims. Don't worry your answers will not be overwritten. Once you finish with Stage 2, you will be shown with a confirmation message.

If you have any questions or encounter any technical issues, please report them to our team! Thank you for your time and effort!

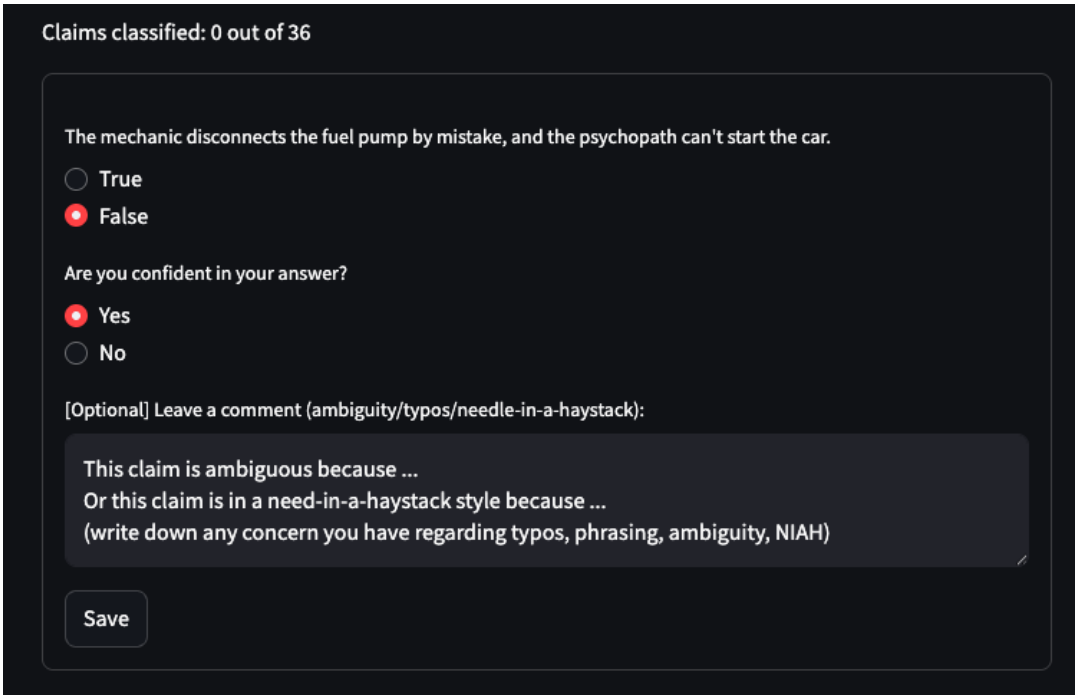


Illustration of the human evaluation interface.

Figure 11: Guidelines provided for human evaluation (Part 2).

Model	#Params	#Frames	HF Checkpoint
		Closed Models	
GPT-4o (OpenAI et al., 2024)	-	50	–
Gemini 2.5 Pro (Team et al., 2023)	-	120	–
		Open-weight Models	
mPLUG-Owl3* (Ye et al., 2025a)	7B	–	mPLUG/mPLUG-Owl3-7B-240728
VideoLLaMA3 (Zhang et al., 2025)	7B	180	DAMO-NLP-SG/VideoLLaMA3-7B
Qwen2.5-VL (Bai et al., 2025b)	7B	1024	Qwen/Qwen2.5-VL-7B-Instruct
LLaVA-Video (Zhang et al., 2024)	7B	64	lmms-lab/LLaVA-Video-7B-Qwen2
LongVILA-R1 (Chen et al., 2025)	7B	180	Efficient-Large-Model/LongVILA-R1-7B
MiMo-VL-7B-RL (Xiaomi, 2025)	7B	256	XiaomiMiMo/MiMo-VL-7B-RL
Kangaroo* (Liu et al., 2024a)	8B	24	KangarooGroup/kangaroo
InternVL3.5 (Zhu et al., 2025)	8B	64	OpenGVLab/InternVL3_5-8B
LLaVAOneVision1.5 (An et al., 2025)	8B	32	lmms-lab/LLaVA-One-Vision-1.5-8B-Instruct
Molmo2 (Clark et al., 2026)	8B	256	allenai/Molmo2-8B
Phi4* (Microsoft et al., 2025)	14B	256	microsoft/Phi-4-multimodal-instruct
Aria (Li et al., 2025)	25B(3.9B)	64	rhymes-ai/Aria
Gemma3 (Team et al., 2025)	27B	180	google/gemma-3-27b-it
Qwen3VL-IT (Bai et al., 2025a)	30B	1024	Qwen/Qwen3-VL-30B-A3B-Instruct
Ovis2 (Lu et al., 2024)	34B	10	AIDC-AI/Ovis2-34B
InternVL3.5 (Wang et al., 2025b)	38B	64	OpenGVLab/InternVL3_5-38B-Instruct
Qwen2.5-VL (Bai et al., 2025b)	72B	1024	Qwen/Qwen2.5-VL-72B-Instruct
LLaVA-Video(Zhang et al., 2024)	72B	64	lmms-lab/LLaVA-Video-72B-Qwen2
InternVL3 (Zhu et al., 2025)	78B	64	OpenGVLab/InternVL3-78B
GLM-4.5V-FP8 (Team et al., 2026)	108B	512	zai-org/GLM-4.5V-FP8

Table 7: Overview of the evaluated vision-language models, including parameter scale, number of input frames, and hf checkpoint.

Direct Prompt Template

System: You are a helpful AI assistant. Your task is to carefully analyze the provided content and determine whether statements made about it are true or false based on the available information.

User: You are provided with a movie and a statement. Your task is to carefully watch the movie and then determine whether the statement is true or false.
Answer TRUE if the statement is true in its entirety based on the movie.
Answer FALSE if any part of the statement is false based on the movie.

Statement: {claim}
Based on the movie, is the above statement TRUE or FALSE?
Provide only your final answer.

Figure 12: Direct prompt template used for open-weight models.

Explanation Prompt Template

System: You are a helpful AI assistant. Your task is to carefully analyze the provided content and determine whether statements made about it are true or false based on the available information.

User: You are provided with a movie and a statement. Your task is to carefully watch the movie and then determine whether the statement is true or false.
Answer TRUE if the statement is true in its entirety based on the movie.
Answer FALSE if any part of the statement is false based on the movie.

Statement: {claim}
Based on the movie, is the above statement TRUE or FALSE?
First provide an explanation of your decision-making process in at most one paragraph, and then provide your final answer.

Figure 13: Explanation prompt template used for closed models.

Movie (Year)	Genre (IMDB)	Lang.	Duration
The Last Chance (1945)	Drama, War	en, it	93.84
They Made Me a Criminal (1939)	Boxing, Film Noir, Crime, Drama, Sport	en	91.21
Tokyo After Dark (1959)	Drama	en	81.23
The Sadist (1963)	Horror, Thriller	en	91.63
Suddenly (1954)	Film Noir, Psychological Thriller, Crime, Drama, Thriller	en	76.71
Sabotage (Hitchcock) (1936)	Psychological Thriller, Spy, Crime, Thriller	en	75.92
Murder By Contract (1958)	Film Noir, Crime, Drama, Thriller	en	80.45
Pushover (1954)	Film Noir, Crime, Drama, Thriller	en	87.77
Go for Broke (1951)	Drama, History, War	en	90.85
Meet John Doe (1941)	Political Drama, Satire, Comedy, Drama, Romance	en	122.87
Scarlet Street (1945)	Film Noir, Tragedy, Crime, Drama, Thriller	en	102.39
Little Lord Fauntleroy (1936)	Period Drama, Drama, Family	en	100.72
Deadline - U.S.A. (1952)	Film Noir, Crime, Drama	en	87.06
My Favorite Brunette (1947)	Hard-boiled Detective, Comedy, Crime, Mystery, Romance, Thriller	en	87.34
Woman in the Moon (1929)	Adventure, Comedy, Drama, Romance, Sci-Fi	de	168.73
Lonely Wives (1931)	Comedy, Romance	en	85.35
Nothing Sacred (1937)	Satire, Screwball Comedy, Comedy, Drama, Fantasy, Romance	en	73.57
Fingerman (1955)	Film Noir, Crime, Drama, Thriller	en	82.06
Borderline (1950)	Film Noir, Crime, Drama, Thriller	en	88.16
Babes in Toyland (1934)	Screwball Comedy, Slapstick, Comedy, Family, Fantasy, Musical	en	77.26
The Man From Utah (1934)	Drama, Western	en	51.49
The Man With The Golden Arm (1955)	Drug Crime, Psychological Drama, Crime, Drama, Romance	en	119.07
A Star Is Born (1937)	Tragic Romance, Drama, Romance	en	110.98
Africa Screams (1949)	Farce, Action, Adventure, Comedy	en	79.13
Dementia 13 (1963)	Slasher Horror, Horror, Thriller	en	74.94
Fear and Desire (1952)	Drama, Thriller, War	en	70.19
The Little Princess (1939)	Costume Drama, Comedy, Drama, Family, Musical	en	92.77
Father's Little Dividend (1951)	Comedy, Drama, Romance	en	81.74
Kansas City Confidential (1952)	Conspiracy Thriller, Film Noir, Heist, Crime, Drama, Thriller	en	99.27
Of Human Bondage (1934)	Dark Romance, Film Noir, Medical Drama, Tragedy, Tragic Romance, Drama, Romance	en	82.77
Half Shot at Sunrise (1930)	Comedy, Musical	en, fr	78.04
Bowery at Midnight (1942)	B-Horror, Crime, Horror, Thriller	en	62.05

Table 8: Metadata of collected movies (Part 1).

Movie (Year)	Genre (IMDB)	Lang.	Duration
The Emperor Jones (1933)	Drama, Music	en	76.29
The Deadly Companions (1961)	Adventure, Drama, Western	en	93.62
The Red House (1947)	Film Noir, Drama, Mystery, Thriller	en	100.39
Trapped (1949)	Film Noir, Crime, Drama, Thriller	en	79.4
City of Fear (1959)	Crime, Drama, Thriller	en	75.18
Kid Monk Baroni (1952)	Action, Drama, Sport	en	79.56
Tight Spot (1955)	Film Noir, Crime, Drama, Thriller	en	95.99
Captain Kidd (1945)	Costume Drama, Swashbuckler, Adventure, Biography, Drama, History	en	87.53
The Front Page (1931)	Dark Comedy, Satire, Screwball Comedy, Comedy, Crime, Drama, Mystery, Romance	en	101.14
The Hitch-Hiker (1953)	Film Noir, Crime, Drama, Thriller	en	70.8
Obsession (1949)	Film Noir, Psychological Thriller, Crime, Thriller	en	92.39
Thunderbolt (1929)	Film Noir, Crime, Drama, Music, Romance	en	91.27
Cyrano de Bergerac (1950)	Swashbuckler, Adventure, Drama, Romance	en	112.87
Scandal Sheet (1952)	Film Noir, Crime, Drama, Romance, Thriller	en	81.75
Ladies in Retirement (1941)	Film Noir, Crime, Drama	en	92.31
Detour (1945)	Film Noir, Crime, Drama	en	69.09
The Crooked Way (1949)	Film Noir, Crime, Drama, Thriller	en	85.95
A Bucket of Blood (1959)	Comedy, Crime, Horror	en	65.84
Love Affair (1939)	Holiday Romance, Comedy, Drama, Romance	en	89.62
The Jackie Robinson Story (1950)	Biography, Drama, Sport	en	76.82
The Last Time I Saw Paris (1954)	Tragedy, Tragic Romance, Drama, Romance	en	116.02

Table 9: Metadata of collected movies (Part 2).

Example 1 (Part 1) - Movie: *"Suddenly"*



Pidge brings the water



Setting up the circuit



Pop spills the water



The electrocution

Fact (✓): The grandfather, Peter Benson, intentionally spilled the water on the floor.
Fib (✗): The grandfather, Peter Benson, accidentally spilled the water on the floor.

Granularity: Multi-scene/Global

Context (Spoiler alert !): Early in the film, the TV breaks. Later, when the criminals have already broken into the house, pop (the grandfather) orchestrates the setup to save everyone (without explicitly stating it). He feigns a heart attack and asks for water, Pidge brings it. The TV technician pretends to repair the broken TV as part of a sabotage plan, while the circuit. When the moment comes, Pop *intentionally* spills the water on the floor (placing it on the edge of the table). Shortly after, the electrocution occurs when one of the criminals steps on the spilled water. *Both visual and textual information are needed to understand that Pop indeed spilled the water and what was his intention. Gemini fails to predict both claims correctly in every modality setting.*

Gemini's Predictions:

Subs only setting

Prediction on Fact: The statement is "The grandfather, Peter Benson, intentionally spilled the water on the floor." To determine if this is true, we need explicit evidence in the transcript: 1. That Peter Benson spilled water; 2. That he did so intentionally. Transcript evidence: Peter Benson (Pop) feigns a heart attack. He asks Pidge: "Pidge, put down that gun and get some water for your grandfather." Pidge presumably brings the water. Later, just before the climax, the line "Oh, butterfingers!" appears. This implies something was dropped or spilled, likely the water. However, the transcript does not explicitly state who exclaimed "Oh, butterfingers!" or who spilled the water. Since the transcript does not explicitly confirm that Peter Benson spilled the water, the statement cannot be verified. **FALSE** *(Ground-truth: TRUE)*

Prediction on Fib: The transcript mentions Pidge being asked to get water for his grandfather, Peter Benson (Pop), when Pop feigns a heart attack [..] However, the transcript does not state that Peter Benson subsequently spilled this water, or any water, on the floor. [...] it is not explicitly stated who Jud is referring to or that the fumbling involved spilling water. The context is the electrocution and shooting, not an earlier incident of spilling water. There is no direct statement or clear implication in the transcript that Peter Benson accidentally spilled water on the floor. **FALSE** *(Ground-truth: FALSE)*

Table 10: **Illustrative Example 1 from the movie "Suddenly"**. Gemini fails to predict both claims correctly in every modality setting.

Example 1 (Part 2) - Movie: *"Suddenly"*

Video only setting

Prediction on Fact: The grandfather, Peter Benson, is seen at 2:17 reaching for a glass of water. He then appears to deliberately drop the glass, causing the water to spill on the floor. This creates a distraction, allowing the woman to crawl towards a gun. The context of the scene, where the family is being held hostage and looking for an opportunity to fight back, combined with the grandfather's actions and the subsequent events, strongly suggests the spill was intentional. **TRUE** *(Ground-truth: TRUE)*

Prediction on Fib: The video introduces an older character who is later identified as Peter Benson (around 0:46-0:47). Much later in the video, during a tense hostage situation, Peter Benson is seen sitting in an armchair holding a glass of water. At timestamp 2:18, he is startled and gets up abruptly, causing the water in his glass to spill onto the floor. The context of the scene, where he reacts to the unfolding events, suggests the spill was accidental due to his sudden movement or reaction. **TRUE** *(Ground-truth: FALSE)*

Video w/ subs setting

Prediction on Fact: The grandfather, Peter Benson, is seen getting up with a glass of water around timestamp 55:20. As he walks, he stumbles and drops the glass, spilling the water on the floor near the table leg and electrical wiring. Immediately after, he clutches his chest and feigns a heart attack. The transcript reflects this with dialogue about his pain. There is no visual or textual evidence to suggest the spill was intentional; rather, it appears accidental, possibly due to his feigned or real physical distress, and he then uses the moment to further distract the captors by exaggerating his heart condition.**TRUE** *(Ground-truth: TRUE)*

Prediction on Fib: The grandfather, Peter Benson, takes a glass of water and walks towards the table where the assassin has set up the rifle. He places the glass on the edge of the metal table (around 55:23). As he moves to talk to Jud, who is working on the TV, the glass tips over and spills the water on the floor (around 55:24). The spilling appears to be an unintentional consequence of his actions and the precarious placement of the glass. Therefore, it was an accident.**TRUE** *(Ground-truth: FALSE)*

Table 11: **Illustrative Example 1 from the movie "Suddenly"**. Gemini fails to predict both claims correctly in every modality setting.

Example 2 (Part 1) - Movie: ***"Sabotage (Hitchcock)"***

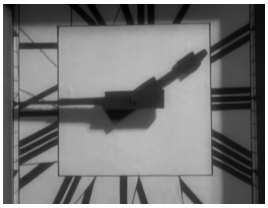


Aquarium


Planning the sabotage


Meet in cinema


Kid enters the bus


The time arrives


Bus explosion


People in the cinema


Explosion in cinema

Fact (✓): The second bomb exploded in the cinema.
Fib (✗): There was a second bomb about to explode, but it didn't.

Granularity: Global
Context (Spoiler alert !): Mr. Verloc and the terrorist meet at the aquarium, where the terrorist first mentions the plan to explode a bomb (he mentions "a parcel of fireworks" and asks Mr. Verloc to meet his friend the bombmaker). The terrorists meet again in the cinema, to discuss about the sabotage. Later in the movie the explosion happens inside the bus. While the second explosion happens inside the cinema. *Gemini fails to predict both claims correctly in every modality setting.*

Gemini's Predictions:

Subs only setting

Prediction on Fact: To determine if the statement is true or false, we need to identify the sequence of bomb-related incidents in the transcript. 1. **First Incident (Sabotage at the Cinema):** The transcript begins with a power outage at the cinema, described as "Sand. Sabotage. Wrecking. Deliberate." This is the first act of sabotage or a device causing disruption that Verloc is connected with. [...] This incident occurred at the cinema. 2. **Second Incident (Stevie's Bomb on the Bus):** Mr. Verloc is then tasked with a more serious act: to leave a "small parcel... a parcel of fireworks" (a bomb) in Piccadilly Circus. His young brother-in-law, Stevie, is sent to deliver this parcel. This bomb explodes prematurely. [...] This is the second distinct sabotage/bomb event in the plot, and this bomb exploded on a bus, not in the cinema. 3. **Third Incident (Bomb Explosion at the Cinema):** Towards the end of the transcript, the old man who makes the bombs is cornered in the cinema. [...] This bomb then explodes: [...] This explosion occurs at the cinema. Based on this chronological sequence of events: * The first destructive "bomb-like" event was the sabotage at the cinema (power outage). * The second bomb (Stevie's parcel) exploded on a bus in the West End. * The third bomb (detonated by the old man) exploded in the cinema. [...] Therefore, the statement is false. **FALSE** *(Ground-truth: TRUE)*

Table 12: **Illustrative Example 2 from "Sabotage (Hitchcock)"**. Gemini predicts the pair of claims incorrectly in every modality setting.

Example 2 (Part 2) - Movie: *"Sabotage"*

Prediction on Fib: The transcript describes two main bomb-related incidents. The first is the explosion that kills Stevie. The second occurs at the cinema where the old man (the bomb maker) threatens to "blow the place sky high." The transcript later confirms an explosion there: "The roof's gone. The whole of the back of the place - clean blown off." This indicates the second bomb *did* explode. Therefore, the statement that a second bomb was "about to explode, but it didn't" is false because it did explode. **FALSE** *(Ground-truth: FALSE)*

Video-only setting

Prediction on Fact: The video shows one clear explosion involving a bus on a street around the 1:46 mark. Following this, there are various scenes, including one where people are watching a cartoon in a cinema. However, there is no depiction or indication of a second bomb exploding, nor is there any explosion shown occurring inside the cinema. **FALSE** *(Ground-truth: TRUE)*

Prediction on Fib: The video montage shows scenes from Alfred Hitchcock's "Sabotage." The primary bomb plot shown results in an explosion on a bus around 1:46. Later in the montage, around 2:17-2:19, the Professor (the bomb maker) is seen suspiciously entering the cinema building. This implies he might be carrying a second bomb and intending to cause another explosion. However, the video montage ends shortly after this, at 2:23, without showing what happens to the Professor or if there was indeed a second bomb that was "about to explode." Crucially, the video does *not* show this second bomb failing to explode or being prevented from exploding. Since the video does not provide any information about the outcome of this potential second bomb (neither showing it explode nor showing it *not* explode), the assertion "but it didn't" is not supported by the available video content. Therefore, the statement cannot be confirmed as true based on the provided video. **FALSE** *(Ground-truth: FALSE)*

Table 13: **Illustrative Example 2 from "Sabotage (Hitchcock)"**. Gemini predicts the pair of claims incorrectly in every modality setting.

Example 2 (Part 3) - Movie: *"Sabotage"*

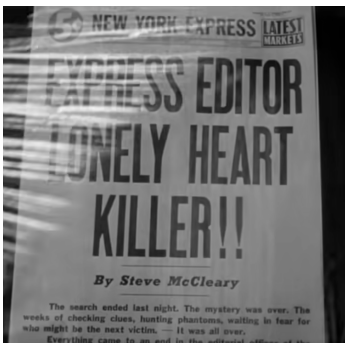
Video w/ subs setting

Prediction on Fact: The first act of sabotage was the power station blackout. The second planned act of sabotage, referred to as "a parcel of fireworks", was to be left in the cloakroom at Piccadilly Circus underground station [...] This parcel, carried by Stevie on a bus, explodes on the bus at 1:45. A newspaper headline later confirms "WEST END BOMB SENSATION. Fully-laden bus was blown up in Piccadilly this afternoon". [...] An explosion then occurs at the cinema, as described by an officer: "The roof's gone. The whole of the back of the place - clean blown off". This is a separate, third bomb event. Therefore, the second bomb (the one Stevie was carrying) exploded on the bus, not in the cinema. The explosion in the cinema was a distinct event caused by a different bomb wielded by the bird shop owner. **FALSE** *(Ground-truth: TRUE)*

Prediction on Fib: The statement claims: "There was a second bomb about to explode, but it didn't." Let's analyze the bombs in the film: 1. The first act of sabotage is the power station blackout. While not explicitly shown as a bomb explosion, it is an act of "wilful destruction." 2. The second major plot device is the bomb intended for Piccadilly Circus on Lord Mayor's Show day. Verloc gives this bomb to Stevie. This bomb explodes prematurely on the bus, killing Stevie and many others [...]. So this bomb *did* explode. 3. The third instance involving a bomb is at the cinema, where the Professor (the bomb maker) is cornered by the police. [..] This bomb *does* explode [...]. In all instances where a bomb is primed or ready for its purpose within the narrative, it either explodes (bus bomb, Professor's bomb in the cinema) or fulfills its sabotage (power station). There is no scene in the film where a bomb is "about to explode" and then "didn't." The bombs that are set up to go off, do go off. Therefore, the statement is false. **FALSE** *(Ground-truth: FALSE)*

Table 14: **Illustrative Example 2 from "Sabotage (Hitchcock)"**. Gemini predicts the pair of claims incorrectly in every modality setting.

Example 3 (Part 1) - Movie: *"Scandal Sheet"*



News about the killer



Sales increasing..



Sales increasing..

Fact (✓): The sales of the newspaper increased after it was revealed that the editor was the killer.
Fib (✗): The sales of the newspaper decreased after it was revealed that the editor was the killer.

Granularity: Single-scene

Context (**Spoiler alert !**): In the film's climax, Steve McCleary discovers that his editor-in-chief, Mark Chapman, is the "Lonely Hearts Killer." After Chapman is exposed and confronted, McCleary—now acting as editor—publishes the full story, revealing Chapman's true identity and crimes. Following the revelation, the newspaper's sales surge dramatically. *This pair of claims relies primarily on visual information. Gemini predicts them incorrectly when given only a single modality, but when provided with both the transcript and the video, it correctly identifies the pair of claims.*

Gemini's Predictions:

Subs only setting

Prediction on Fact: The transcript ends with Mark Chapman, the editor, being apprehended as the killer. He tells the reporter McCleary to "Write it up big, kid. It'll sell a lot of papers." This indicates an expectation that the story will boost sales, but the transcript concludes before any information about the actual sales figures *after* this revelation is provided. Therefore, the statement that sales *increased* after the revelation cannot be verified from the given text. **FALSE** *(Ground-truth: TRUE)*

Prediction on Fib: The transcript ends with Mark Chapman, the editor, being apprehended as the killer. His last words are, Ẅrite it up big, kid. It'll sell a lot of papers.T̈his is a prediction about future sales, not a statement of what actually happened to the sales. The transcript provides no information about the newspaper's sales figures *after* it was revealed that the editor was the killer. Therefore, the statement cannot be confirmed as true based on the provided text. **FALSE** *(Ground-truth: FALSE)*

Table 15: **Illustrative Example 3 from "Scandal Sheet"**. The pair of claims relies primarily on visual information. Gemini predicts them incorrectly when given only a single modality, but when provided with both the transcript and the video, it correctly identifies the pair of claims.

Example 3 (Part 2) - Movie: *"Scandal Sheet"*

Video only setting

Prediction on Fact: The video shows a chart of the "NET PAID CIRCULATION NEW YORK EXPRESS" at various points. Initially, at 0:10, the circulation is shown with an upward trend. After the editor is revealed to be the killer and is subsequently shot (around 2:26-2:28), the final shot of the movie (2:29) shows a newspaper headline "KILLER!!" with the byline "By Steve McCleary" superimposed over the same circulation chart. At this point, the line on the chart indicating sales has reached its highest point, implying a significant increase in sales coinciding with or immediately following the revelation and capture of the killer editor. **TRUE** *(Ground-truth: TRUE)*

Prediction on Fib: The movie shows a chart of the "New York Express" net paid circulation at 0:10, indicating an upward trend in sales from October to September, reaching 650,000. Later, after the editor (the man standing by the chart at 0:10) is revealed as the killer and is shot (around 2:27), the final shot at 2:29 shows a newspaper headline "KILLER!!" superimposed over the *same* circulation chart. The numbers and the upward trend on this chart are identical to the one shown at 0:10. The movie does not provide any information or visual evidence suggesting that the newspaper's sales decreased after the editor was revealed as the killer. The chart displayed at the end is the same one that showed increasing sales. **TRUE** *(Ground-truth: FALSE)*

Video w/ subs setting

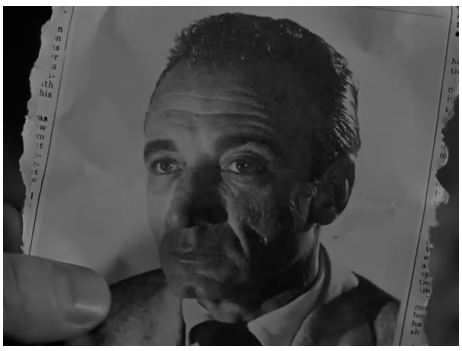
Prediction on Fact: The statement's truthfulness hinges on whether newspaper sales increased after the revelation that the editor, Mark Chapman, was the killer. 1. At 3:40, Mark Chapman discusses the newspaper's circulation, showing a graph and mentioning that when circulation reaches 750,000, a big bonus will be paid. The graph shows circulation rising, with the last point around June/July showing approximately 620,000-650,000. 2. The revelation that Mark Chapman is the killer unfolds towards the end of the film. 3. At 59:39, the newspaper headline "EXPRESS EDITOR LONELY HEART KILLER!!" by Steve McCleary is shown. 4. Immediately following this (59:45 - 59:49), a large circulation meter for the "N.Y. EXPRESS" is shown. A hand physically moves the needle on the meter, stopping at 750,000. "The End" is then superimposed over this image. This sequence strongly implies that the story revealing the editor as the killer led to the circulation reaching 750,000. Earlier, Chapman had also told McCleary to write the story big because "It'll sell a lot of papers" (59:06). The visual of the circulation hitting 750,000 immediately after the headline about the editor being the killer is the final image before "The End," suggesting this was the peak sales driven by this sensational story. **TRUE** *(Ground-truth: TRUE)*

Prediction on Fib: At the end of the movie, after the editor, Mark Chapman, is revealed as the killer and the headline "EXPRESS EDITOR LONELY HEART KILLER!!" is shown (59:40-59:43), the circulation meter for the New York Express is shown. A hand moves the needle to 750,000 (59:45). Earlier in the film (around 4:24), it was mentioned that reaching 750,000 in circulation would trigger a big bonus. The film ends with the circulation at this high number, indicating that sales reached this target *after* the revelation about the editor. There is no indication that sales decreased; in fact, the implication is that this sensational story would boost sales, as McCleary says, "Write it up big, kid. It'll sell a lot of papers." (59:14-59:16). **FALSE** *(Ground-truth: FALSE)*

Table 16: **Illustrative Example 3 from "Scandal Sheet"**. The pair of claims relies primarily on visual information. Gemini predicts them incorrectly when given only a single modality, but when provided with both the transcript and the video, it correctly identifies the pair of claims.

Example 4 (Part 1) - Movie: *"Tight Spot"*





Fact (✓): When Hallet shows Conley a photograph of an individual who was on the yacht, she immediately recognizes the person.
Fib (✗): When Hallet shows Conley a photograph of an individual who was on the yacht, she claims not to recognize the person.

Granularity: Single-scene

Context (Spoiler alert !): This pair of claims to be answered correctly requires *both* the *visual* and *text* cues. Hallet shows Conley a photograph (visual information) and asks whether she recognizes the person (Senior Maracelli). Conley recognizes the person. *When the model is given only the subtitles, it misclassifies both the fact and the fib (as visual cues are absent). In contrast, when provided with both the video and the subtitles, it predicts each one correctly.*

Gemini's Predictions:

Subs only setting

Prediction on Fact: The statement claims that "When Hallet shows Conley a photograph of an individual who was on the yacht, she immediately recognizes the person." We need to check two things in the transcript: 1. Does Hallett show Conley a photograph of an individual from the yacht? 2. If so, does Conley immediately recognize the person? In the transcript, the relevant exchange is: HALLETT: "And at one of these islands, this man came aboard the yacht." CONLEY: "Signor Maracelli. The Italian garden expert. He taught me to say his name himself." [...] *The transcript does not use the word "photograph" or describe the action of showing one in this specific interaction.* Since the statement's premise that Hallett shows Conley a photograph is not explicitly confirmed by the transcript, a part of the statement is not verifiably true based on the provided text. **FALSE** *(Ground-truth: TRUE)*

Prediction on Fib: When Mr. Hallett mentions Signor Maracelli coming aboard the yacht, Conley immediately identifies him as "The Italian garden expert" and adds, "He taught me to say his name himself." This indicates she clearly recognizes the person Hallett is referring to (and presumably showing a photograph of, though the act of showing the photo isn't explicitly stated, Hallett's line "this man came aboard the yacht" implies he's identifying someone specific). She does not claim not to recognize him. **FALSE** *(Ground-truth: FALSE)*

Table 17: **Illustrative Example 4 from "Tight Spot"**. Gemini fails to predict both claims correctly when given only the subtitles or only the visual cues. However, when provided with both modalities together, it correctly predicts both claims.

Example 4 (Part 2) - Movie: *"Tight Spot"*

Video only setting

Prediction on Fact: The statement asks whether Conley immediately recognizes the person in the photograph shown to her by Hallet. In the video, at 00:30, Hallet takes out a photograph. At 00:31, he shows it to Conley. Conley looks down at the photo, then at 00:32, she looks up at Hallet with an expression that seems more contemplative or questioning rather than one of instant, clear recognition. She then glances back at the photo before looking at Hallet again and starting to speak at 00:33. Her reaction involves a pause and a sequence of looks, which does not convey *immediate* recognition. "Immediate" suggests an instantaneous, reflexive understanding, which is not evident in her more measured response. **FALSE** *(Ground-truth: TRUE)*

Prediction on Fib: The statement claims that when Hallet shows Conley a photograph of an individual from a yacht, Conley claims not to recognize the person. Observing the video, there is a scene from approximately 0:29 to 0:36 where a man (presumably Hallet, played by Edward G. Robinson) shows a photograph to a woman (presumably Conley, played by Ginger Rogers). The woman looks at the photograph (0:31-0:32). However, the video segment does not include any audible dialogue or clear visual cue (like shaking her head definitively or mouthing "I don't know") from the woman that would confirm she "claims not to recognize the person." Her facial expression is somewhat pensive or surprised but not an explicit denial of recognition. She then proceeds to make a phone call (0:34). The context that the photograph is of "an individual who was on the yacht" is also not explicitly provided by the visual information in this clip. Since the crucial part of the statement - her claiming not to recognize the person - is not depicted or audible in the provided video evidence, the statement cannot be verified as true based on the available information. **FALSE** *(Ground-truth: FALSE)*

Video w/ subs setting

Prediction on Fact: When Hallett shows Conley the photograph (around 00:40), he identifies the man as "Signor Maracelli." Conley immediately confirms recognition by saying, "The Italian garden expert. He taught me to say his name himself." This demonstrates she knew who the person was. **TRUE** *(Ground-truth: TRUE)*

Prediction on Fib: When Mr. Hallett mentions Signor Maracelli coming aboard the yacht, Conley immediately identifies him as "The Italian garden expert" and adds, "He taught me to say his name himself." This indicates she clearly recognizes the person Hallett is referring to (and presumably showing a photograph of, though the act of showing the photo isn't explicitly stated, Hallett's line "this man came aboard the yacht" implies he's identifying someone specific). She does not claim not to recognize him. **FALSE** *(Ground-truth: FALSE)*

Table 18: **Illustrative Example 4 from "Tight Spot"**. Gemini fails to predict both claims correctly when given only the subtitles or only the visual cues. However, when provided with both modalities together, it correctly predicts both claims.

Example 5 (Part 1) - Movie: *"Tight Spot"*



Pete in taxi



Walking the stairs of the church



Pete being shot

Fact (✓): Pete Tinelli was shot.
Fib (✗): Pete Tinelli was poisoned.

Granularity: Multi-scene

Context (Spoiler alert !): Early in the film, we are introduced to Pete during a scene in which he rides in a taxi with two other men. Later, the group arrives in front of a church, and as they walk up the steps, Pete is suddenly shot while the other two men look around, trying to identify the shooter. *This pair of claims requires both visual and text cues to be resolved correctly, as the model needs to understand who is Pete (in the first scene) and connect it with his murder (being shot). Gemini fails to predict correctly the pair of claims in every modality setting*.

Gemini's Predictions:

Subs only setting

Prediction on Fact: The transcript states that "Pete's dead" and that "He was murdered by the Costain mob." [..] However, the transcript does not specify the method of murder. It does not explicitly say that Pete Tonelli was shot. Therefore, because the specific detail of "shot" is not confirmed in the transcript, the statement is false. **FALSE** *(Ground-truth: TRUE)*

Prediction on Fib: The transcript states, "Pete's dead... He was murdered by the Costain mob." It does not specify the method of murder. Since the transcript doesn't mention poisoning as the cause of Pete Tonelli's death, and only confirms he was murdered, the statement that he was poisoned cannot be confirmed as true based on the provided text. **FALSE** *(Ground-truth: FALSE)*

Table 19: **Illustrative Example 5 from "Tight Spot"**. Gemini fails to predict the pair of claims correctly, in every modality setting.

Example 5 (Part 2) - Movie: *"Tight Spot"*

Video only setting

Prediction on Fact: A man is visibly shot around the 1:57-1:58 mark in the video. In this scene, a man is held by two others, another man (resembling Edward G. Robinson) nods, a gunshot sound is heard, and the man being held slumps. However, throughout the provided video clip, there is no dialogue, on-screen text, or any other information that identifies the man who was shot as "Pete Tinelli." [...] **FALSE** *(Ground-truth: TRUE)*

Prediction on Fib: The video contains several scenes that lead to the conclusion that Pete Tinelli was poisoned. At 1:58, a group of men are shown in what appears to be a garage, and one man is forcibly made to drink something. He subsequently collapses at 2:01. Later, at 2:05, a woman asks, "What happened to Pete Tinelli?" and is told, "He won't bother you anymore." Shortly after, at 2:08, one man informs another, "Pete Tinelli's been taken care of." When asked "How?", the response is "Poison." This sequence of events and dialogue directly confirms that Pete Tinelli was poisoned.**TRUE** *(Ground-truth: FALSE)*

Video w/ subs setting

Prediction on Fact: The provided transcript states at 00:54-00:55 that "Pete's dead" and "He was murdered by the Costain mob." However, neither the transcript nor the accompanying video footage specifies the method of his murder. The statement "Pete Tinelli was shot" asserts a specific cause of death (being shot) which is not explicitly mentioned or shown in the provided materials. Since the method of murder is not confirmed, we cannot ascertain that he was shot based solely on the given information. **FALSE** *(Ground-truth: TRUE)*

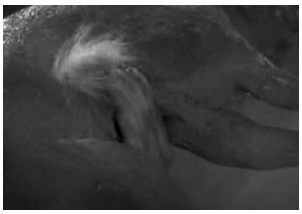
Prediction on Fib: The transcript states at 00:42, "He was murdered by the Costain mob." While this confirms Pete Tonelli was murdered, it does not specify the method of murder. The statement claims he was poisoned. Since the available information does not mention poisoning as the cause of death, the specific claim that he was poisoned cannot be verified as true from the provided movie and transcript. **FALSE** *(Ground-truth: FALSE)*

Table 20: **Illustrative Example 5 from "Tight Spot"**. Gemini fails to predict the pair of claims correctly, in every modality setting.

Example 6 (Part 1) - Movie: *"A Bucket of Blood"*


Holding the cat sculpture


Drops it


The damage


Moment of realization

Fact (✓): The gallery owner dropped the cat sculpture, causing minor damage.

Fib (✗): The gallery owner dropped the cat sculpture, but it remained intact.

Granularity: Single-scene

Context (Spoiler alert !): When the gallery owner accidentally drops the cat sculpture in the cafe and sees real flesh beneath the clay, he realizes that the sculpture is not a realistic imitation but an actual dead cat covered in clay. *Gemini does not correctly predict the claim pair, as it fails to recognize that the sculpture was damaged.*

Gemini's Predictions:

Subs-only setting

Prediction on Fact: The transcript details the creation of the "Dead Cat" sculpture by Walter and its reception at the Yellow Door, owned by Leonard De Santis. Leonard agrees to display the cat sculpture. Later, a customer takes the sculpture home to show his wife and then makes an offer to buy it for $500. Leonard tells this customer, "Someone has the cat just now, but I'll have him back in a few days. Then you can have it for the $500." There is no mention in the transcript of Leonard, the gallery/club owner, dropping the cat sculpture or of it sustaining any damage. **FALSE**

Prediction on Fib: The transcript mentions that Mr. De Santis, the owner of the Yellow Door (where the art is displayed, akin to a gallery owner in this context), arranged for a potential buyer to take the cat sculpture home to show his wife. It states, "A man wanted to pay me $100 for the cat. In fact, he's taking it home to show to his wife." Later, the same man calls back to buy it, and Mr. De Santis says, "Someone has the cat just now, but I'll have him back in a few days." There is no mention in the transcript of Mr. De Santis or anyone else dropping the cat sculpture. Therefore, the premise that the gallery owner dropped it is false. **FALSE**

Table 21: **Illustrative Example 6 from the movie: "A Bucket of Blood"**. Gemini fails to predict the pair of claims correctly, in every modality setting.

Example 6 (Part 2) - Movie: *"A Bucket of Blood"*

Video-only setting

Prediction on Fact: The statement is: "The gallery owner dropped the cat sculpture, causing minor damage." To verify this, we need to identify a "gallery owner," a "cat sculpture," an act of "dropping" by the gallery owner, and resulting "minor damage." 1. **Gallery Owner:** * Leonard de Santis is the owner of "The Yellow Dog" cafe, which also functions as an art gallery where poetry is read and art is displayed [...] Walter Paisley creates a sculpture of a cat by encasing his landlady's dead cat in clay (0:16-0:20). This is the most prominent cat sculpture. * At 0:29, Leonard and Carla are seen holding other, more stylized white cat sculptures, indicating these might be common art pieces in their circle. 3. **Dropped the cat sculpture, causing minor damage:** * At 0:29, Leonard (the cafe/gallery owner) is shown holding one of the stylized white cat sculptures. He does not drop it. [...] No scene in the provided video shows a character identifiable as "the gallery owner" dropping a cat sculpture and causing minor damage. The closest event involves a police officer tapping Walter's cat sculpture, causing a tiny chip. Therefore, the statement is false. **FALSE**

Prediction on Fib: The statement says: "The gallery owner dropped the cat sculpture, but it remained intact." Let's break this down: 1. **Identity of the "gallery owner"**: In the movie clip, the gallery owner is Maxwell Brock, often seen wearing a beret or a suit, and he runs "The Yellow Door" gallery. 2. **Did the gallery owner drop a cat sculpture?**: * At 0:27, the main character, Walter Paisley (a busboy who becomes an artist), is carrying a large abstract sculpture (not a cat) when he is bumped by Carla. Walter drops this sculpture, and it shatters. Walter is not the gallery owner. * Walter Paisley creates a sculpture of a cat (by encasing his landlady's dead cat in clay) around 0:22. This sculpture is later handled but not dropped by the gallery owner. * At 0:50, Maxwell Brock (the gallery owner) is seen handling a *different* white, stylized cat sculpture. He picks it up, examines it carefully, and places it back down. He does *not* drop this sculpture. Since there is no instance in the provided video where the gallery owner (Maxwell Brock) drops any cat sculpture, the first part of the statement ("The gallery owner dropped the cat sculpture") is false. Therefore, the entire statement is false. **FALSE**

Video w/ subs setting

Prediction on Fact: The video and transcript show the gallery owner, Leonard, handling the cat sculpture on a couple of occasions. First, when Walter initially brings it to "The Yellow Door" (around 13:36-13:55), Leonard is skeptical but eventually agrees to display it. He takes the sculpture from Walter and places it on a surface without dropping it. Later, after the sculpture gains attention (around 25:05), Leonard is seen picking up the cat sculpture, looking at it, and then placing it back down on a table, again without dropping or damaging it. The transcript contains no mention of the gallery owner dropping the sculpture. **FALSE**

Prediction on Fib: The statement claims that the gallery owner dropped the cat sculpture. Reviewing the video and transcript, at 13:47, Leonard De Santis, the owner of the Yellow Door gallery/cafe, picks up the cat sculpture made by Walter Paisley. Shortly after, at 13:48, Walter Paisley himself is holding the sculpture. At 13:49, Walter Paisley (the artist, not the gallery owner) drops the cat sculpture onto the floor. The sculpture appears to remain intact. Since it was Walter, the artist, and not Leonard, the gallery owner, who dropped the sculpture, the first part of the statement is false. **FALSE**

Table 22: **Illustrative Example 6 from the movie: "A Bucket of Blood"**. Gemini fails to predict the pair of claims correctly, in every modality setting.