GOOSE: *Which Video Does Not Belong?* MULTI-VIDEO REASONING FOR OUT-OF-CONTEXT DETECTION IN EVENT STREAMS

Anonymous authors
Paper under double-blind review

ABSTRACT

Recent multimodal large language models (MLLMs) show strong multi-video reasoning, yet social-media event streams pose a distinct challenge: authentic footage can still be *contextually false* when it originates from a different event. Existing benchmarks typically ask predefined questions about known targets, rather than requiring models to reconstruct an event and discover conflicting posts. To address this gap, we introduce **GOOSE** (**G**roup reasoning **O**ver **O**ut-of-context videos in **S**ocial-media **E**vents), a benchmark for discovering genuine footage falsely presented as evidence of another event. GOOSE contains 281 real-world cases, 351 OOC videos, and 4,085 videos across 6 domains, covering six evidence mechanisms involving scene, entity, temporal–causal, physical, multi-view, and provenance conflicts. It emphasizes **group-level evidence aggregation**, **event-level reasoning**, and **open-ended investigation**: models must discover unspecified OOC posts, while truthfully captioned off-event posts prevent relevance-based shortcuts. We evaluate two settings: Direct IO jointly processes uniformly sampled frames, whereas a VideoAgent-style ReAct framework iteratively selects and inspects videos. Discovery is measured by case-macro OOC@K, and a rubric-guided native-video MLLM Judge evaluates grounded cross-video explanations; the Evidence-Grounded Score (EGS) jointly captures both aspects. Across 15 leading MLLMs, Gemini 3.8 Flash retrieves **86.73%** of gold OOCs but achieves only **58.56% EGS**, showing that correct discovery does not guarantee a valid explanation. VideoAgent also often trails Direct IO because fragmented inspection weakens global event reconstruction and candidate–anchor comparison. These findings identify grounded cross-video verification, rather than visual plausibility or provenance recall alone, as the central bottleneck in real-world OOC detection.

1 INTRODUCTION

MLLMs increasingly process and reason across multiple videos, yet seeing many clips is not equivalent to reconstructing an event. This distinction becomes critical when an AI agent monitors or manages a social-media account, where no moderator marks which post deserves verification. An out-of-context (OOC) post reuses authentic footage from one event under the caption of another and may therefore appear fully plausible in isolation. False context accounts for **35%** of false or misleading visual content in an international fact-checking corpus, compared with only **10%** involving visual manipulation (Brennen et al., 2021); misinformation exposure has also produced measurable behavioral effects (Loomba et al., 2021). We ask: *can an agent discover what needs verification before it can verify it?*

Existing benchmarks remain misaligned with this setting in three respects (Table 2). **(1) Sample-level verification.** Multimodal misinformation benchmarks typically present an image–text pair, video, or claim that has already been selected for verification (Luo et al., 2021; Papadopoulos et al., 2024; Qi et al., 2023a; 2024; Zhang et al., 2026a). Retrieval-grounded and agentic methods may reconstruct provenance, search the Web, or invoke external tools, but they still begin from a supplied target (Tonglet et al., 2024; Yu et al., 2026; Xu et al., 2026c). **(2) Question-conditioned video groups.** Multi-video benchmarks evaluate comparison, association, temporal alignment, and

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

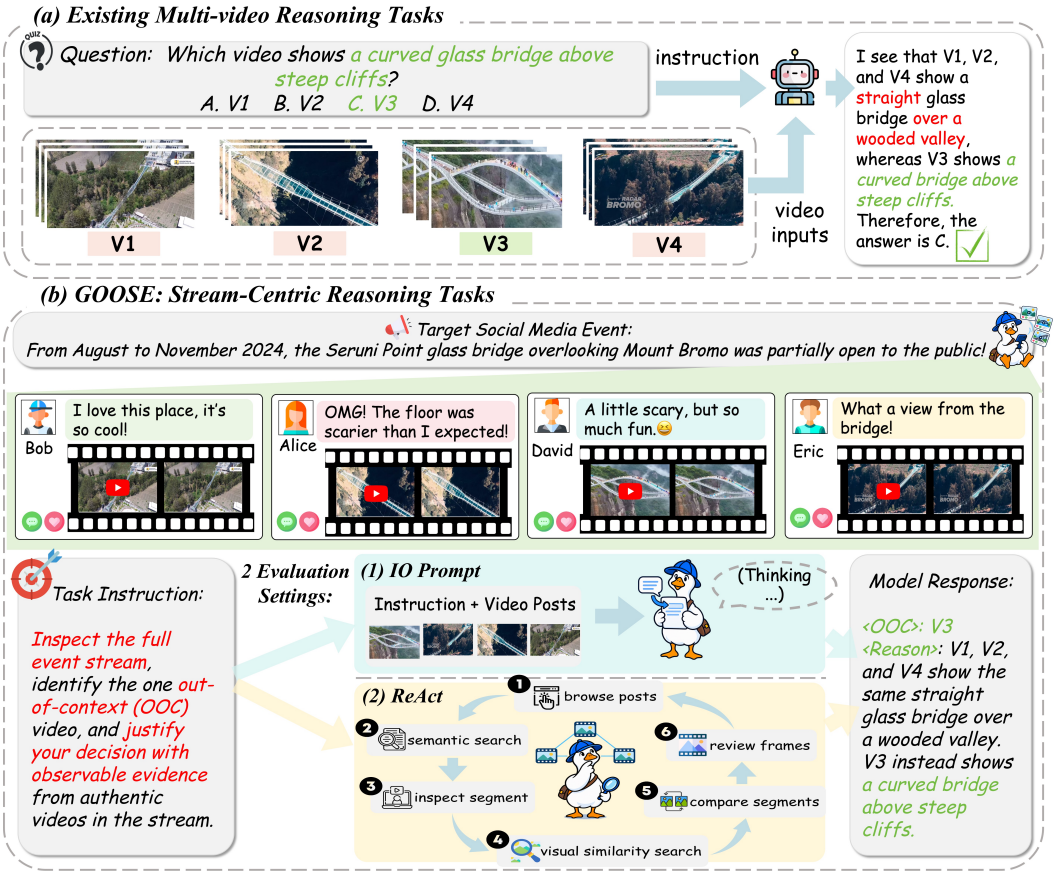

Figure 1: Examples showing the GOOSE tasks and the difference from existing multi-video tasks.

structured reasoning, but provide a question, reference video, or curated video set before inference (Li et al., 2026b; Zhu et al., 2025; Zhang et al., 2026c; Sun et al., 2026). NEED uses same-event neighbors to improve fake-video detection, yet still classifies an already designated candidate (Qi et al., 2023b). **(3) Unaudited use of group evidence.** Existing tasks primarily evaluate a prediction or generated answer; even when multiple videos are available, they rarely require the model to identify which posts establish the conflict. Consequently, a correct prediction need not show that the model has formed an event-level account or used the relevant cross-video evidence. Across these paradigms, the benchmark designer provides the suspicion before the model begins reasoning.

We call the missing capability *group reasoning*: constructing a bounded account of an event from multiple independent posts, discovering locally plausible observations that conflict with this account, and identifying the observations that make each conflict verifiable. Group reasoning is not visual outlier detection or majority voting. Different viewpoints, event phases, recording conditions, and reposts can create legitimate visual differences, while a truthfully captioned video from another event may appear more unusual than the misinformation itself. The mismatch must instead be established through positive, event-specific relations distributed across videos.

To address this gap, we introduce GOOSE, **G**roup reasoning **O**ver **O**ut-of-context videos in **S**ocial-media **E**vents. GOOSE contains **281** real-world event streams, **351** OOC videos, and **4,085** videos across 6 domains. Each stream contains 9–20 posts and mixes three semantic roles: REAL videos that truthfully depict the target event, authentic OOC videos falsely attributed to that event, and UNRELATED videos whose off-event captions remain truthful. The last category provides matched hard negatives, preventing models from treating every off-event video as misinformation.

GOOSE differs from prior benchmarks in three key aspects. **First, target-free investigation:** models receive the complete event stream without a marked suspect and must discover the OOC videos themselves. **Second, group-resolvable construction:** every core OOC remains plausible in isolation, contains no decisive isolated shortcut, and is contradicted by at least two independently sourced

Real anchors. **Third, evidence-grounded output:** models must both identify the OOC videos and provide the in-stream evidence supporting each decision. We organize the annotated evidence into a six-way taxonomy of OOC evidence mechanisms: *scene fingerprint mismatch, entity configuration mismatch, temporal–causal conflict, physical-mechanism conflict, multi-view inconsistency, and provenance-integrity failure*. This mechanism-centric taxonomy serves as the primary axis for diagnosing which forms of observable cross-video evidence models can actually use.

We evaluate 15 proprietary and open-weight MLLMs under two settings: direct input–output inference (IO) and a ReAct-style VideoAgent setting (Wang et al., 2024; Yao et al., 2022). The strongest result is obtained by Gemini 3.8 Flash under IO, with 86.73% Accuracy and 58.56% Evidence-Grounded Score (EGS), revealing a substantial gap between discovering OOC videos and supporting them with valid cross-video evidence. Performance also varies markedly across the six evidence mechanisms, with temporal–causal conflicts posing the greatest challenge. Moreover, the VideoAgent setting does not consistently improve performance: averaged across models, it reduces Accuracy by 6.24 points and EGS by 9.44 points. These results show that an agentic workflow alone does not improve event-stream reasoning; the main challenge remains grounding OOC decisions in valid cross-video evidence.

In summary, our contributions are threefold. **(1)** We formulate target-free investigative OOC reasoning, shifting evaluation from verifying a supplied candidate to discovering suspicious videos within an event stream. **(2)** We construct GOOSE with group-resolvable OOC cases, independent Real anchors, truthfully captioned off-event distractors, and a six-way evidence-mechanism taxonomy for fine-grained diagnosis of cross-video reasoning. **(3)** Through a comprehensive evaluation of 14 MLLMs under both direct inference and VideoAgent-style agentic reasoning, we reveal a persistent gap between OOC discovery and evidence-grounded verification, and show that active evidence acquisition alone is insufficient to overcome cross-video evidence-binding failures.

2 Task Overview

This section formalizes stream-level OOC discovery in GOOSE, introduces the evidence mechanisms used to characterize contextual conflicts, and describes the benchmark's construction process and dataset statistics. We then compare its evaluation setting with existing multi-video reasoning and multimodal misinformation benchmarks.

2.1 Task Definition

Let E denote a concrete target event and $\mathcal{S}_E = \{p_1, \ldots, p_N\}$ its social-media stream, where $p_i = (c_i, v_i, t_i, m_i)$ contains a caption, video, publication time, and benchmark-visible metadata. A post is out of context (OOC) if v_i is authentic footage of an identifiable source event $E_i^{\text{src}} \neq E$ but c_i falsely attributes it to E (Luo et al., 2021; Papadopoulos et al., 2024). A REAL post truthfully depicts E, whereas an UNRELATED post truthfully describes another event and serves as a hard distractor. Synthetic, substantially manipulated, and unauditable media are excluded.

Given the complete stream without marked candidates, a model ranks all posts and explains each predicted OOC using in-stream evidence. We evaluate the fraction of gold OOCs appearing within the top-K predictions, and use a native-video Judge to score the quality of the supporting explanation. Each retained OOC is plausible in isolation and must be contradicted by at least two independently sourced Real anchors, making group evidence necessary for both discovery and explanation.

2.2 Evaluation Metrics

We report three complementary scores. For a single event stream, let $\mathcal{O}$ be the gold OOC set, $K = |\mathcal{O}| \geq 1$, π the predicted ranking, and $\mathcal{H} = \text{Top}_K(\pi) \cap \mathcal{O}$. **Mismatch Detection Accuracy (Acc.)**, or OOC@K, is $|\mathcal{H}|/K$ and reduces to Top-1 accuracy for single-OOC streams. **Explanation Quality (EQ)** is the Judge score $s_j \in [0, 5]$ for a retrieved gold OOC, defined as the minimum of evidence grounding, relational sufficiency, and faithfulness. The **Evidence-Grounded Score (EGS)** jointly measures discovery and explanation as $\frac{1}{K}\sum_{j\in\mathcal{H}} s_j/5$: missed OOCs and invalid predictions receive zero. reported results average these per-case Acc. and EGS values over event streams

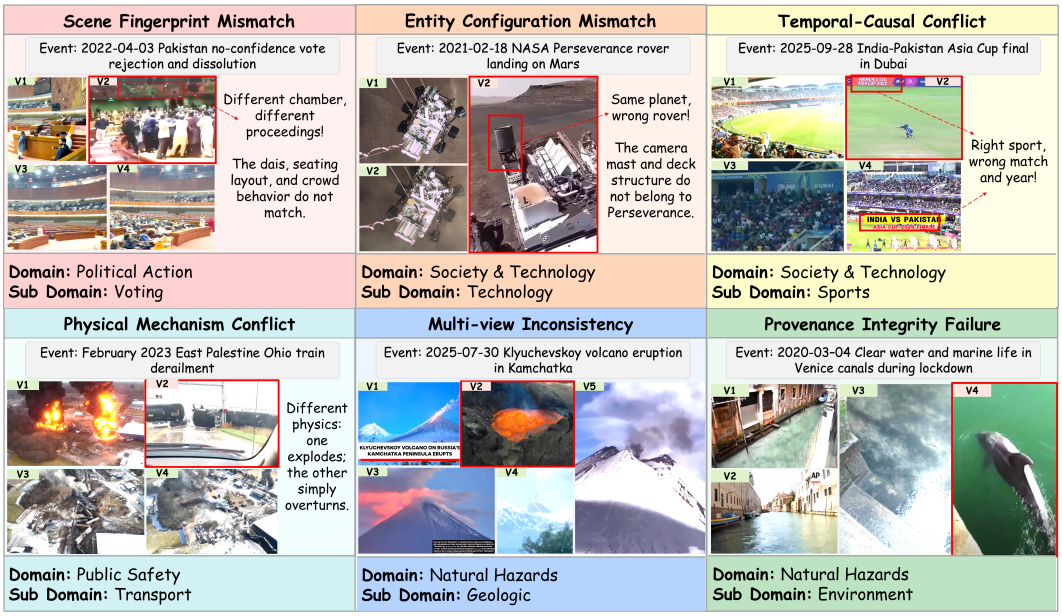

Figure 2: **Overview of the six evidence mechanisms in GOOSE.** Each mechanism specifies an observable cross-video relation through which independently sourced posts of the target event expose an otherwise plausible OOC attribution.

and express them as percentages. If a retrieved OOC lacks a Judge score, its EGS is reported as unavailable rather than imputed. Full Judge prompts and scoring details are provided in Appendix G.

2.3 Evidence Mechanisms

To characterize how an out-of-context video deviates from the target event, we organize the cases into six evidence mechanisms: **(i) Scene Fingerprint Mismatch (SFM)**, conflicts in stable scene cues such as terrain, architecture, weather, lighting, or landmarks; **(ii) Entity Configuration Mismatch (ECM)**, differences in the identities, numbers, attributes, or arrangements of visible entities; **(iii) Temporal–Causal Conflict (TCC)**, inconsistencies in time, event phase, or causal sequence; **(iv) Physical Mechanism Conflict (PMC)**, contradictions in motion, physical interaction, damage progression, or observable outcomes; **(v) Multi-view Inconsistency (MVI)**, failure to align with the spatial geometry, landmarks, or instance details established by authentic views; and **(vi) Provenance Integrity Failure (PIF)**, source-level conflicts arising from recycled, mislabeled, or mixed-event footage.

3 The GOOSE Benchmark

3.1 Data Construction

The construction of GOOSE follows three principles. **P1. Event-stream discovery:** Models must identify OOC videos within a complete event stream without being given suspicious candidates. **P2. Cross-video reasoning:** Each OOC video is plausible in isolation, and its false attribution becomes evident only through comparison with multiple independent videos of the target event. **P3. Visual grounding:** Every decision must be supported by directly observable cross-video evidence rather than source-event memorization or external knowledge. It consists of three stages: OOC-seeded event collection, group-resolvable evidence annotation, and independent human quality assurance. Figure 3 further motivates our emphasis on cross-video visual evidence rather than isolated source recognition.

OOC-seeded event and video collection. We begin with documented OOC claims from professional fact-checking organizations and collect the corresponding videos primarily from publicly accessible posts on YouTube, X, Facebook, and Instagram. Each candidate case is screened using

the following checklist: (1) **Concrete event**: the falsely claimed target must correspond to a specific occurrence with an identifiable time and place; (2) **Sufficient event coverage**: multiple independently sourced videos must capture different viewpoints or phases of the target event; (3) **Visually verifiable conflict**: the OOC attribution must be refutable through observable relations with at least two independent Real videos, without requiring knowledge of its authentic source event; and (4) **Realistic distractors**: Unrelated videos must be truthfully captioned and visually or semantically similar to the target event. For each admitted case, annotators initially collect 25 candidate Real videos and retain 5–25 after verification and deduplication, together with 1–5 verified OOC videos and 3–17 Unrelated videos.

Group-resolvable evidence annotation. For each stream containing at least ten posts, human annotators compare every OOC candidate with the collected Real videos, without using Web search. A candidate is retained only when at least two independently sourced Real anchors reveal a positive, observable conflict that cannot be resolved from the candidate alone or explained by legitimate within-event variation. Human annotators then cite the decisive anchors, assign one of the six evidence mechanisms, and write the group-reasoning evidence by describing the concrete visual discrepancy between the Real and OOC videos. The annotated anchors ensure that each case is resolvable from evidence within the stream, but models are not required to use the same anchors; any valid visual evidence supporting the OOC decision is accepted. Because several valid conflicts may exist, the annotation provides a verified reference rather than the only acceptable explanation. Separately, DeepSeek-V4.1-Flash (Xu et al., 2026b) summarizes the verified true-source-event evidence for each OOC video from the corresponding fact-checking materials; it does not produce the group-reasoning annotation. Both annotations are subsequently checked during human quality assurance.

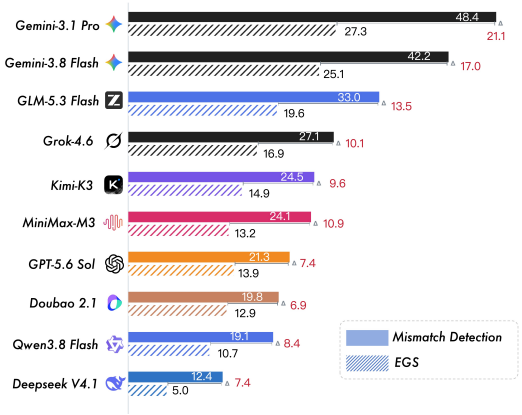

Figure 3: **Isolated OOC performance.** Detection consistently exceeds EGS, exposing a gap between identifying mismatches and grounding them in visual evidence.

Human quality assurance. Each completed stream is independently reviewed by two annotators using the three-part checklist illustrated in Figure 6: (1) **Playable and auditable**: all videos remain accessible, captions clearly make the annotated event claims, and no benchmark metadata reveals the labels; (2) **Independent and sufficient anchors**: each OOC is supported by at least two non-duplicate Real sources whose observable evidence rules out benign differences in viewpoint, time, or event phase; and (3) **Verified provenance**: the target event, authentic source event, and misleading attribution are traceable to reliable public sources. Disagreements are resolved by a third adjudicator, while cases that fail any criterion are corrected, returned for additional collection, or excluded. The complete construction workflow, protocol, and checklist are detailed in Appendix A (Figure 6).

3.2 Dataset Statistics

GOOSE comprises 281 event streams and 4,085 videos spanning 6 real-world domains and 20 subdomains. Across the 281 evaluation streams, each event stream contains an average of 14.54 posts: 1.25 OOC, 8.11 Real, and 5.18 Unrelated. The benchmark covers six evidence mechanisms for cross-video OOC reasoning, as illustrated in Figure 2. Video durations range from 4.93 to 21,485.26 seconds, with a mean of 203.00 seconds and a median of 125.99 seconds; the complete distribution is shown in Figure 8. Detailed dataset statistics are reported in Table 1, while Table 2 compares GOOSE with related multi-video reasoning and misinformation benchmarks.

Table 1: Statistics of GOOSE.

Dataset	Statistic
Event stream	
Case count	281
OOC Mechanism Types	6
Domain count	6
Videos/Case (min/max/avg)	9/20/14.54
Video	
Video count	4,085
OOC video count	351
Video Duration (min) (min/max/avg)	0.08/358.09/3.38

Table 2: **Comparison with related multi-video reasoning and video-misinformation benchmarks.** Avg. #Vids/Ex. denotes the average number of input videos per instance; Dur. denotes the average duration of one video in seconds. O/C denotes open-/closed-ended outputs, and A/M denotes automatic/manual benchmark annotation. Evidence eval. indicates separate evaluation of retrieved source videos or explicit cross-video differences. "–" indicates that a comparable statistic could not be established.

Benchmark	#Videos	Avg. #Vids/Ex.	Dur.	Input level	Form	Construction taxonomy	Target-free	Evidence eval.	Anno.
Multi-video reasoning benchmarks									
CrossVid (Li et al., 2026b)	5,331	2.87	215.0	Video set	O/C	4 dimensions / 10 tasks / 7 content types	✗	✗	A&M
CVBench (Zhu et al., 2025)	1,315	2.79	106.6	Video set	C	3 reasoning levels / 15 subtasks / 5 domains	✗	✗	A&M
MVU-Eval (Peng et al., 2026)	4,959	4.71	–	Video set	C	Perception and reasoning / 8 abilities	✗	✗	A&M
MVX-Bench (Zhang et al., 2026c)	4,255	–	–	Video set	C	3 cognitive levels / 11 tasks	✗	✗	A&M
VideoSetDiff (Qiu et al., 2025)	3,132	4.00	55.4	Four-video set	C	9 comparison aspects / 27 categories	✗	✗	A&M
VidDiffBench (Burgess et al., 2025)	1,098	2.00	8.8	Video pair	O/C	5 domains / 18 activities / 148 differences	✗	✓	M
EgoExoBench (He et al., 2026)	–	–	–	Ego–exo set	C	3 dimensions / 11 tasks	✗	✗	A&M
X-Stream (Sun et al., 2026)	932	2.22	948.0	Sync. streams	O/C	8 domains / 3 levels / 11 tasks	✗	✗	A&M
Video-misinformation benchmarks									
FakeSV / NEED (Qi et al., 2023a;b)	5,538	10.29	–	Candidate + neighbors	C	2 authenticity labels / 3 video roles	✗	✗	M
FakeVV (Zhang et al., 2026a)	∼100K	1.00	–	Video + title	O/C	4 manipulated-entity types	✗	✗	A
EVID-Bench (Yu et al., 2026)	222	4	–	Candidate + retrieval	O	3 manipulation categories / 9 types / 6 domains	✗	✓	A&M
GOOSE (Ours)	**4,085**	**14.54**	**203.00**	**Event stream**	**O**	**6 evidence mechanisms / 6 domains / 20 subdomains**	✓	✓	**A&M**

3.3 Comparison with Existing Benchmarks

Table 2 compares GOOSE with closely related multi-video reasoning and multimodal misinformation benchmarks along evaluation unit, target visibility, evidence scope, required cross-video reasoning, and grounding output. Unlike sample-level or designated-candidate benchmarks, GOOSE embeds a latent verification target in an event stream and requires the model to both discover and justify it using endogenous multi-video evidence.

4 Experiments

4.1 Experimental Settings

We evaluate 15 MLLMs, comprising seven proprietary models from the GPT-5.6, Gemini 3.x, Grok 4.6, Qwen3.8, and Doubao-Seed families (OpenAI, 2026; Google DeepMind, 2026a;c; xAI, 2026; Qiu et al., 2026; Volcengine, 2026), and eight open-weight models: DeepSeek-V4.1-Flash, Kimi K3, GLM-5.3-Flash, MiniMax-M3, InternVL3.5-8B, and Qwen3-VL-Instruct at 4B, 8B, and 32B scales (Xu et al., 2026a; Team et al., 2026; Zeng et al., 2026a; Lai et al., 2026; Wang et al., 2025). Exact endpoints, checkpoints, access dates, and decoding configurations are provided in Appendix E.

All models are evaluated zero-shot on the same anonymously ordered event streams, without post roles or the number of OOCs. The main comparison uses a benchmark-controlled frame interface with identical resolution, temporal order, captions, and output schema. Direct inference uses 16 uniformly sampled frames per post, capped at 250 images per stream, whereas our adaptation of VideoAgent (Wang et al., 2024) uses a ReAct-style loop (Yao et al., 2022) and customized tools to survey the stream, inspect selected clips, compare videos, and cite evidence under a maximum budget of 512 delivered frames. For the human reference, **7** participants independently solve a stratified subset of **40** streams using the same interface and without Web access. Complete prompts, sampling settings, tool definitions details are provided in Appendices E, G.

4.2 Overall Results

Table 3 reports direct inference (IO), our primary setting, and the VideoAgent-style baseline (Agent), where available. Results are grouped by six mechanisms that make an OOC video falsifiable: Scene Fingerprint Mismatch (SFM), Entity Configuration Mismatch (ECM), Temporal–Causal Conflict (TCC), Physical Mechanism Conflict (PMC), Multi-view Inconsistency (MVI), and Provenance Integrity Failure (PIF). Acc. measures the fraction of gold OOC posts retrieved within the hidden top-K, while EGS measures whether those detections are supported by grounded evidence.

Group reasoning improves both OOC discovery and evidence-grounded verification. Across seven models and 2,320 paired OOC predictions, detection rises from 32.07% in the isolated consistency experiment to 75.69% in the event-stream group reasoning experiment, a gain of 43.62 points. The model-average EGS also rises from 19.3% to 46.9%, with all seven models improving. Authentic posts therefore serve as event-specific anchors rather than merely acting as distractors. To

Table 3: Main results (%) by evidence mechanism under direct inference (IO) and ReAct baseline (Agent). Green and red cells indicate the best and second-best scores within each model group, respectively; best scores are bold.

Model	Setting	Avg.		SFM		ECM		TCC		PMC		MVI		PIF	
		Acc.	EGS	Acc.	EGS	Acc.	EGS	Acc.	EGS	Acc.	EGS	Acc.	EGS	Acc.	EGS
Proprietary models															
GPT-5.6 Sol	IO	67.07	44.43	71.74	44.71	59.57	45.96	67.02	35.64	58.87	38.23	73.91	52.83	71.46	49.28
	Agent	70.35	35.26	66.30	34.49	68.09	34.54	63.65	28.05	78.01	39.43	80.43	43.91	65.83	31.40
GPT-5.6 Luna	IO	64.26	39.34	63.59	40.36	67.02	42.62	47.52	24.26	69.86	42.70	70.65	45.43	67.01	40.80
	Agent	41.97	23.32	37.50	20.54	44.68	28.30	31.74	16.06	55.67	30.43	33.70	18.91	48.12	25.50
Gemini 3.1 Pro Preview	IO	75.49	44.82	80.43	42.72	77.66	46.67	68.09	35.85	82.27	47.73	78.26	53.26	66.60	42.89
	Agent	69.83	30.29	69.02	27.39	63.83	32.13	65.43	26.99	75.53	31.91	75.00	34.57	70.24	28.79
Gemini 3.8 Flash	IO	**86.73**	**58.56**	85.33	58.51	**79.79**	**60.85**	**87.23**	**47.80**	**90.78**	**56.81**	**89.13**	**65.87**	**88.12**	**61.60**
	Agent	81.07	38.38	**87.50**	42.14	67.02	33.90	83.87	34.75	85.82	40.64	80.43	42.83	81.88	36.22
Grok 4.6	IO	75.28	52.71	84.78	**60.07**	75.53	57.45	65.25	38.51	74.82	50.43	77.17	58.70	74.41	51.44
	Agent	70.39	40.65	78.80	43.51	68.09	40.14	66.31	34.33	71.63	39.86	67.39	43.70	70.24	42.46
Qwen3.8 Flash	IO	51.63	30.04	53.80	31.52	50.35	30.57	35.82	17.98	53.55	28.16	54.35	34.78	61.77	37.23
	Agent	51.66	28.66	48.91	25.43	47.16	26.31	35.99	18.48	65.60	36.10	59.78	36.52	52.64	29.18
Doubao-Seed-2.1-Turbo	IO	64.92	40.80	71.74	45.14	59.57	39.36	55.14	32.34	62.77	35.67	75.00	49.78	65.62	42.75
	Agent	56.10	31.71	50.00	27.17	45.74	29.93	53.90	25.74	60.28	32.20	66.30	39.35	60.35	35.87
Open-weight models															
DeepSeek-V4.1-Flash	IO	65.74	42.08	63.95	42.39	60.64	41.99	65.25	32.38	66.31	40.99	67.39	48.48	70.76	46.32
	Agent	42.35	21.10	37.32	16.74	34.04	17.66	39.01	18.58	47.52	20.71	56.52	32.17	39.97	20.88
Kimi K3	IO	**75.46**	**52.98**	**77.72**	**56.27**	69.15	**54.33**	**75.89**	**42.98**	**79.08**	**54.11**	78.26	**58.48**	72.85	**51.92**
	Agent	70.42	39.47	71.20	43.08	59.57	33.69	63.30	31.45	76.24	40.07	**82.61**	48.70	69.90	40.08
GLM-5.3-Flash	IO	69.15	48.95	66.85	51.30	65.96	48.94	68.44	38.87	76.95	52.62	73.91	56.09	62.95	46.15
	Agent	66.96	35.77	56.34	29.57	**70.57**	42.98	56.21	26.67	74.11	36.60	66.30	36.96	**77.74**	41.60
MiniMax-M3	IO	47.95	25.97	45.65	21.74	45.74	26.31	40.43	18.94	51.42	25.96	57.61	37.17	47.05	25.87
	Agent	51.50	25.86	50.91	23.22	53.19	28.37	48.94	21.70	46.10	23.97	57.61	30.87	52.36	27.03
InternVL3.5-8B	IO	5.82	1.46	1.45	0.00	0.00	0.00	6.38	2.13	14.89	4.26	7.61	1.96	4.58	0.42
	Agent	1.71	0.07	2.17	0.00	2.13	0.00	2.66	0.00	1.06	0.43	0.00	0.00	2.19	0.00
Qwen3-VL-4B-Instruct	IO	6.05	2.92	6.52	2.61	2.13	1.28	2.13	0.85	8.51	3.83	8.70	5.22	8.33	3.75
	Agent	3.56	0.50	2.17	0.00	0.00	0.00	0.00	0.00	6.38	0.85	8.70	2.17	4.17	0.00
Qwen3-VL-8B-Instruct	IO	14.59	5.76	11.96	4.13	12.77	3.40	8.51	3.40	13.83	4.68	21.74	10.87	18.75	8.13
	Agent	2.92	0.50	2.17	0.00	4.26	1.70	0.00	0.00	6.38	1.28	2.17	0.00	2.50	0.00
Qwen3-VL-32B-Instruct	IO	18.09	7.72	14.13	7.61	10.64	2.98	12.77	5.96	21.99	7.23	21.74	9.13	27.08	13.33
	Agent	13.88	5.35	10.87	3.26	10.64	4.26	8.51	2.55	19.15	7.66	17.39	7.65	16.67	6.67
Baselines															
Best Performance of Models		86.73	58.56	87.50	60.07	79.79	60.85	87.23	47.80	90.78	56.81	89.13	65.87	88.12	61.60
Human		93.62	78.36	95.43	79.62	92.34	79.17	90.21	73.86	94.18	78.37	96.07	79.74	93.56	79.42

separate explanation quality from the change in detection coverage, we further compare 427 single-OOC model–case pairs that are correctly identified in both experiments. Conditional EQS rises from 3.01 to 3.70 out of 5, with all seven models improving. The largest rubric gain occurs in relational sufficiency (+0.77), followed by faithfulness (+0.76) and evidence grounding (+0.60). The additional OOCs recovered only through group reasoning remain harder to explain, averaging 3.20 EQS versus 3.70 for shared detections. **Group reasoning therefore improves not only discovery but also the quality of cross-video explanations, particularly the relation between the OOC post and authentic event anchors. However, the newly recovered difficult cases remain less well grounded and continue to constrain the overall EGS.**

Evidence mechanisms expose distinct capability gaps. Gemini 3.8 Flash achieves the strongest IO result, with 86.75% Acc., yet its EGS is only 58.57%; Gemini 3.1 Pro Preview similarly obtains 75.55% Acc. but 44.85% EGS, while DeepSeek-V4-Flash obtains 65.72% Acc. but 42.09% EGS. The mechanism breakdown localizes this gap. Averaged over the seven models in the paired analysis, MVI is the easiest mechanism by Acc. (74.53%), followed by SFM (72.80%) and PMC (71.94%), whereas ECM (66.87%) and TCC (67.02%) are the hardest. EGS produces an even sharper distinction: MVI reaches 56.58%, while TCC falls to 38.63%. Gemini 3.8 Flash, for example, retrieves 87.23% of TCC cases but obtains only 47.80% EGS. Figures 4b and 4c further show that these capability gaps are structured rather than isolated. Although performance transfers across most mechanisms, the relatively weak correlation between PMC and PIF ($r = 0.49$) indicates that physical plausibility and provenance verification require partly distinct capabilities; the radar profiles likewise reveal uneven evidence- and domain-level strengths even among the strongest models. **Cross-view disagreement and visible scene fingerprints are therefore comparatively accessible, whereas temporal and causal conflicts remain difficult to turn into grounded OOC–Real explanations.**

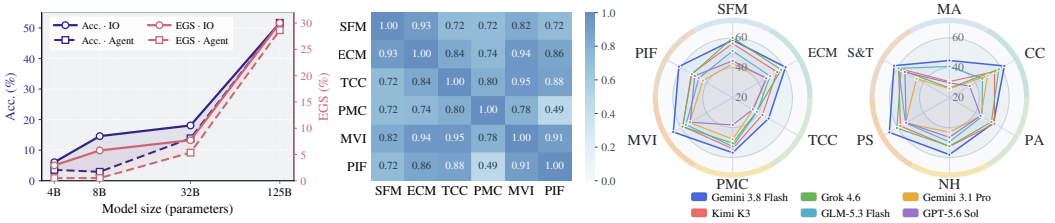

(a) Qwen-family scaling. (b) Evidence correlations. (c) Evidence- & domain-wise performance.

Figure 4: (a) Acc. and EGS scaling across Qwen-family model sizes under IO and Agent settings. (b) Correlations among the six evidence mechanisms across the strongest Direct-IO models. (c) Direct-IO EGS profiles across evidence mechanisms and event domains.

The VideoAgent-style baseline does not consistently improve IO. Across the four proprietary models with both settings, average Acc. decreases from 76.18% under IO to 72.93% under Agent (-3.25 points). For the three models with complete EGS in both settings, average EGS falls from 52.06% to 36.46% (-15.60 points). The effect is not universal: GPT-5.6 Sol and MiniMax-M3 gain 3.29 and 3.54 Acc. points, respectively. The Qwen-family scaling results in Figure 4a reinforce this finding: increasing model scale substantially improves both Acc. and EGS, but the Agent setting generally remains below IO, suggesting that scale alone does not resolve evidence selection and retention failures. **These results show that access to tools is not itself a reasoning capability; the benefit depends on whether the model can acquire, retain, and relate the right evidence.**

5 Analysis

5.1 Agent Is Not All You Need

Sequential exploration fragments cross-video evidence. The substantially larger Agent drop in EGS than in Acc. indicates that sequential exploration often finds a suspicious post but loses the explicit OOC–Real relation needed to justify it: candidates and authentic anchors are inspected in different rounds and compressed into a fragmented trajectory, whereas Direct IO preserves a synchronized view of the stream. A global-to-local refinement, which surveys every post before targeted inspection, partly addresses this failure on Qwen3.8 Flash, raising Acc. from 51.67% to 56.67% and EGS from 28.66% to 33.22%. Thus, global context is necessary but not sufficient: active inspection helps only when the controller can reliably preserve and bind candidate evidence to authentic anchors.

5.2 When Do External Tools Help?

Table 4: Task 3 OOC1 (%) under tool and Web time-window conditions.

Setting	Gemini 3.8 Flash	Qwen3.8 Flash	Kimi K3
Tool conditions			
Baseline(agent)	67.5	**53.0**	59.6
Web	75.8	46.7	65.7
ASR	80.3	39.5	65.5
Web+ASR	**80.8**	30.6	**73.1**
Web time-window conditions			
+2 days	71.7	39.1	**60.0**
+7 days	73.9	**47.8**	52.0
+30 days	**78.3**	37.0	54.0
Unrestricted	**78.3**	41.3	58.0

Tools help only when models bind their evidence correctly. Table 4 shows that Web+ASR gives the best result for Gemini 3.8 Flash (80.8% OOC@1) and Kimi K3 (73.1%), but reduces Qwen3.8 Flash from 46.7% with Web to 30.6%. Trajectories explain this reversal: Qwen averages 2.13 ASR calls per scored stream, yet these runs contain only one wrong-to-correct flip and six correct-to-wrong flips. In five regressions, it replaces the gold OOC with an honestly captioned Unrelated post, treating off-event content as suspicious instead of testing whether the video contradicts its own caption. Web evidence also helps only when used: Qwen searches in only 17 of 122 scored runs. The publication-window results reveal a further model-dependent effect: Gemini improves monotonically from 71.7% at +2 days to 73.9%, 78.3%, and 78.3% as the window expands, whereas Qwen fluctuates between 39.1% and 47.8%, and Kimi similarly varies from 60.0% to 52.0%, 54.0%, and 58.0%. Thus, external evidence cannot replace relational reasoning; without correct post-role and caption–video binding, broader retrieval can reinforce an event-relevance shortcut rather than resolve contextual falsity. The Web time-window results further show that access to more mature reporting benefits Gemini but not all models; detailed analysis is provided in Appendix F.3.

5.3 Internal Knowledge Is a Lead, Not Evidence

Remembering the source is not the same as verifying it. Across 351 paired OOCs, Gemini 3.1 Pro Preview achieves the highest isolated OOC rejection rate (48.43%), but its detection gain from isolated verification to group reasoning is 27.92 points, substantially below 45.58 for Gemini 3.8 Flash, 50.00 for Grok 4.6, and 56.93 for Kimi K3. This smaller gain is consistent with a model that can often form a plausible provenance hypothesis from internal knowledge before seeing authentic event anchors. Such recognition is useful but insufficient: Gemini 3.1 reaches 75.55% IO Acc. yet only 44.85% EGS, showing that many correct detections are not supported by an equally strong grounded relation. A representative case concerns footage of *Keir Starmer*: **in isolation**, the model identifies it as *the October 2023 Labour Party Conference*, although that date is not visible in the cited frames; **with the event stream**, it instead contrasts the suspect video's *"BRITAIN'S FUTURE"* stage with *"CHANGE BEGINS"* on two Real anchors. Internal knowledge therefore provides a valuable lead, but verification requires converting that lead into observable, cross-video evidence.

5.4 Error Analysis

Figure 5 reveals complementary failure modes. Gemini 3.1 Pro Preview often relies on internal knowledge without converting it into observable cross-video evidence, explaining its gap between detection and EGS. Grok 4.6 more often misses the decisive scene, entity, or temporal relation despite inspecting relevant posts. Both models also exhibit reasoning, knowledge-application, and hallucination errors. Thus, recognizing a suspicious video remains easier than producing a faithful cross-video justification. Representative failures are examined in Appendices F.6 and F.7.

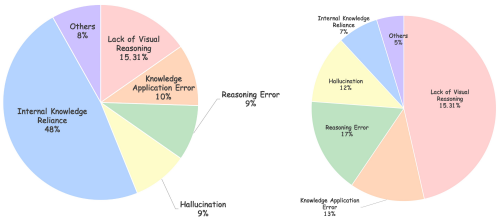

(a) Gemini 3.1 Pro. (b) Grok 4.6.

Figure 5: Representative error distributions.

6 Related Work

Cross-video and multi-stream understanding. Cross-video benchmarks evaluate implicit, comparative, and long-range reasoning (Zhu et al., 2026; Li et al., 2026a;b; Zhu et al., 2025; Peng et al., 2026; Bai et al., 2026; Zhang et al., 2026c), including video differences, synchronized streams, and distributed context (Qiu et al., 2025; Burgess et al., 2025; He et al., 2026; Sun et al., 2026; Yuan et al., 2026). These settings specify the target or question before inference; GOOSE requires discovering both the suspicious post and its supporting videos in a noisy event stream.

Multimodal misinformation verification and agents. Misinformation benchmarks classify supplied media–text pairs or retrieve and explain evidence for designated targets (Luo et al., 2021; Papadopoulos et al., 2024; Qi et al., 2023a; 2024; Tonglet et al., 2024; Xu et al., 2026d; Yu et al., 2026; Zhang et al., 2026a; Zeng et al., 2026b); NEED also assumes a given target (Qi et al., 2023b). Agentic video systems acquire evidence through trajectories, visual tools, and platform context (Yao et al., 2022; Wang et al., 2024; Fan et al., 2024; Zhang et al., 2026b; Zhou et al., 2026; Xu et al., 2026c; Xue et al., 2026; Xu et al., 2026f;e). GOOSE instead evaluates candidate discovery and separates misleading OOC reuse from truthfully captioned off-event posts.

7 Conclusion

We introduce GOOSE for target-free OOC verification in social-media event streams. Results show that detecting suspicious posts is easier than grounding their mismatch, while sequential agents do not consistently outperform direct inference because evidence becomes fragmented. By separating discovery from explanation across six falsifiable mechanisms, GOOSE provides a controlled testbed for evidence-grounded verification in noisy event streams.

ETHICS STATEMENT

The benchmark studies misinformation about real social events and therefore requires careful handling of potentially distressing footage, identifiable social-media accounts, and copyrighted media. We will minimize collection of unnecessary personal information, preserve only metadata needed for verification, document platform and copyright constraints, and avoid redistributing media when redistribution rights are unclear. Annotators will be warned about potentially disturbing content and may skip items. The benchmark is intended to support fact-checking and integrity research; it is not designed to rank individuals, infer private traits, or automate punitive moderation decisions without human oversight. We will document known failure modes and discourage deployment claims that exceed the benchmark's measured capabilities.

REPRODUCIBILITY STATEMENT

We will release the benchmark annotations and event-level splits together with the evaluation code, prompts, agent configurations, and exact model versions used in our experiments. The supplementary material documents the dataset construction, case-selection criteria, human-review procedure, and evaluation metrics. For videos that cannot be redistributed, we provide the original source references where permitted and available. All systems are evaluated using the same data partitions and resource budgets.

AI USE STATEMENT

Generative AI tools were used during early-stage research ideation, literature organization, and drafting assistance for this proposal. The authors manually reviewed the resulting research claims, checked cited prior work against primary paper or proceedings sources, and remain responsible for the final scientific framing, benchmark design, experiments, code, annotations, and manuscript text. Generative AI tools will not be treated as authoritative sources for dataset labels or provenance; benchmark annotations and final evidence relationships will be independently audited by human reviewers. This statement will be updated before submission to accurately reflect all uses covered by the ICLR 2027 AI policy.

REFERENCES

Purui Bai, Tao Wu, Jiayang Sun, Xinyue Liu, Huaibo Huang, and Ran He. Mvpbench: A multi-video perception evaluation benchmark for multi-modal video understanding. *arXiv preprint arXiv:2603.22756*, 2026.

J Scott Brennen, Felix M Simon, and Rasmus Kleis Nielsen. Beyond (mis) representation: Visuals in covid-19 misinformation. *The international journal of press/politics*, 26(1):277–299, 2021.

James Burgess, Xiaohan Wang, Yuhui Zhang, Anita Rau, Alejandro Lozano, Lisa Dunlap, Serena Yeung, et al. Video action differencing. In *International Conference on Learning Representations*, volume 2025, pp. 15049–15081, 2025.

Dongping Chen, Ruoxi Chen, Shilin Zhang, Yinuo Liu, Yaochen Wang, Huichi Zhou, Qihui Zhang, Yao Wan, Pan Zhou, and Lichao Sun. Mllm-as-a-judge: Assessing multimodal llm-as-a-judge with vision-language benchmark. *arXiv preprint arXiv:2402.04788*, 2024.

Yue Fan, Xiaojian Ma, Rujie Wu, Yuntao Du, Jiaqi Li, Zhi Gao, and Qing Li. Videoagent: A memory-augmented multimodal agent for video understanding. In *European Conference on Computer Vision*, pp. 75–92. Springer, 2024.

Google DeepMind. Gemini 3.1 Pro model card. `https://deepmind.google/models/model-cards/gemini-3-1-pro/`, February 2026a. Accessed: 2026-09-23.

Google DeepMind. Gemini 3.7 Flash model card. `https://deepmind.google/models/model-cards/gemini-3-7-flash/`, August 2026b. Accessed: 2026-09-23.

Google DeepMind. Gemini 3.8 Flash model card. `https://deepmind.google/models/model-cards/gemini-3-8-flash/`, September 2026c. Accessed: 2026-09-23.

Yuping He, Yifei Huang, Guo Chen, Baoqi Pei, Jilan Xu, Tong Lu, and Jiangmiao Pang. Egoexobench: A benchmark for first-and third-person view video understanding in mllms. *Advances in Neural Information Processing Systems*, 38, 2026.

Wonkyun Kim, Changin Choi, Wonseok Lee, and Wonjong Rhee. An image grid can be worth a video: Zero-shot video question answering using a vlm. *IEEE Access*, 12:193057–193075, 2024.

Xunhao Lai, Weiqi Xu, Yufeng Yang, Qiaorui Chen, Yang Xu, Lunbin Zeng, Xiaolong Li, Haohai Sun, Haichao Zhu, Vito Zhang, et al. Minimax sparse attention. *arXiv preprint arXiv:2606.13392*, 2026.

Seongyun Lee, Seungone Kim, Sue Park, Geewook Kim, and Minjoon Seo. Prometheus-vision: Vision-language model as a judge for fine-grained evaluation. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 11286–11315, 2024.

Jiachun Li, Shaoping Huang, Zhuoran Jin, Chenlong Zhang, Pengfei Cao, Yubo Chen, Kang Liu, and Jun Zhao. Mmr-life: Piecing together real-life scenes for multimodal multi-image reasoning. *arXiv preprint arXiv:2603.02024*, 2026a.

Jingyao Li, Jingyun Wang, Molin Tan, Haochen Wang, Cilin Yan, Likun Shi, Jiayin Cai, Xiaolong Jiang, and Yao Hu. Crossvid: A comprehensive benchmark for evaluating cross-video reasoning in multimodal large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 6244–6252, 2026b.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: Nlg evaluation using gpt-4 with better human alignment. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, pp. 2511–2522, 2023.

Sahil Loomba, Alexandre De Figueiredo, Simon J Piatek, Kristen De Graaf, and Heidi J Larson. Measuring the impact of covid-19 vaccine misinformation on vaccination intent in the uk and usa. *Nature human behaviour*, 5(3):337–348, 2021.

Grace Luo, Trevor Darrell, and Anna Rohrbach. Newsclippings: Automatic generation of out-of-context multimodal media. In *Proceedings Of The 2021 conference on empirical methods in natural language processing*, pp. 6801–6817, 2021.

Jinxiang Meng, Shaoping Huang, Fangyu Lei, Jingyu Guo, Haoxiang Liu, Jiahao Su, Sihan Wang, Yao Wang, Enrui Wang, Ye Yang, et al. Dv-world: Benchmarking data visualization agents in real-world scenarios. *arXiv preprint arXiv:2604.25914*, 2026.

OpenAI. GPT-5.6: Frontier intelligence that scales with your ambition. `https://openai.com/index/gpt-5-6/`, 2026. Accessed: 2026-09-23.

Stefanos-Iordanis Papadopoulos, Christos Koutlis, Symeon Papadopoulos, and Panagiotis C Petrantonakis. Verite: a robust benchmark for multimodal misinformation detection accounting for unimodal bias. *International Journal of Multimedia Information Retrieval*, 13(1):4, 2024.

Tianhao Peng, Haochen Wang, Yuanxing Zhang, Noah Wang, Zili Wang, Ge Zhang, Jian Yang, Shihao Li, Yanghai Wang, Xintao Wang, et al. Mvu-eval: Towards multi-video understanding evaluation for multimodal llms. *Advances in Neural Information Processing Systems*, 38, 2026.

Peng Qi, Yuyan Bu, Juan Cao, Wei Ji, Ruihao Shui, Junbin Xiao, Danding Wang, and Tat-Seng Chua. Fakesv: A multimodal benchmark with rich social context for fake news detection on short video platforms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 14444–14452, 2023a.

Peng Qi, Yuyang Zhao, Yufeng Shen, Wei Ji, Juan Cao, and Tat-Seng Chua. Two heads are better than one: Improving fake news video detection by correlating with neighbors. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 11947–11959, 2023b.

Peng Qi, Zehong Yan, Wynne Hsu, and Mong Li Lee. Sniffer: Multimodal large language model for explainable out-of-context misinformation detection. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13052–13062. IEEE, 2024.

Yue Qiu, Yanjun Sun, Takuma Yagi, Shusaku Egami, Natsuki Miyata, Ken Fukuda, Kensho Hara, and Ryusuke Sagawa. Videosetdiff: Identifying and reasoning similarities and differences in similar videos. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12242–12252. IEEE, 2025.

Zihan Qiu, Zekun Wang, Xiao Li, Yanpeng Li, Yang Xu, Yixuan Wang, Huaqing Zhang, Rui Men, Bochao Mao, Chengruidong Zhang, et al. On the design of qwen3. 8-next architecture: Evaluation, efficiency, and training stability. *arXiv preprint arXiv:2608.30320*, 2026.

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

Peiwen Sun, Xudong Lu, Huadai Liu, Yang Bo, Dongming Wu, Huankang Guan, Minghong Cai, Jinpeng Chen, Xintong Guo, Shuhan Li, et al. X-stream: Exploring mllms as multiplexers for multi-stream understanding. *arXiv preprint arXiv:2606.02482*, 2026.

Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, Jianfeng Cai, Xinyuan Cai, Peizhou Cao, Yuxuan Cao, Ziwei Chai, Y Charles, et al. Kimi k3: Open frontier intelligence. *arXiv preprint arXiv:2607.24653*, 2026.

Jonathan Tonglet, Marie Francine Moens, and Iryna Gurevych. "image, tell me your story!" predicting the original meta-context of visual misinformation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 7845–7864, 2024.

Volcengine. Doubao-Seed-2.1-Pro model documentation. `https://www.volcengine.com/docs/82379/2549861?lang=en`, 2026. Accessed: 2026-09-23.

Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, et al. Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025.

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

xAI. Introducing Grok 4.6. `https://x.ai/news/grok-4-6`, August 2026. Accessed: 2026-09-23.

Anyi Xu, B Li, Bangcai Lin, Bing Xue, BingCheng Xian, Bingzheng Xu, Bochao Wu, Bowei Zhang, Boyi Deng, CC Yu, et al. Deepseek-v4. 1-flash: Pushing the limits of kv cache compression. *arXiv preprint arXiv:2609.19969*, 2026a.

Anyi Xu, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, Chenchen Ling, et al. Deepseek-v4: Towards highly efficient million-token context intelligence. *arXiv preprint arXiv:2606.19348*, 2026b.

Cheng Xu, Changhong Jin, Yingjie Niu, Nan Yan, Yuke Mei, Shuhao Guan, Liming Chen, and Tahar Kechadi. Livefact: A dynamic, time-aware benchmark for llm-driven fake news detection. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11881–11910, 2026c.

Danni Xu, Shaojing Fan, Harry Cheng, and Mohan Kankanhalli. Rw-post: Auditable evidence-grounded multimodal fact-checking in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10964–10973, 2026d.

Shuwen Xu, Zhitao He, and Yi R Fung. Redact: Redacting agent capability traces for procedural skill protection. *arXiv preprint arXiv:2606.10813*, 2026e.

Shuwen Xu, Yao Xu, Jiaxiang Liu, Chenhao Yuan, Wenshuo Peng, Jun Zhao, and Kang Liu. Graph-walker: Agentic knowledge graph question answering via synthetic trajectory curriculum. *arXiv preprint arXiv:2603.28533*, 2026f.

Dizhan Xue, Jing Cui, Shengsheng Qian, Chuanrui Hu, and Changsheng Xu. Some: A realistic benchmark for llm-based social media agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 1391–1399, 2026.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.

Tao Yu, Yujia Yang, Shenghua Chai, Zhang Jinshuai, Haopeng Jin, Hao Wang, Minghui Zhang, Zhongtian Luo, Yuchen Long, Xinlong Chen, et al. When seeing is not believing–a benchmark for search-grounded video misinformation detection. *arXiv preprint arXiv:2606.04098*, 2026.

Chenhao Yuan, Yinhao Xu, Shuwen Xu, Xizhi Yang, Jiaxiang Liu, Chenxi Zhou, Shaoping Huang, Haolin Ren, Pengfei Cao, Jun Zhao, et al. Longcrafter: Towards diverse long-context understanding via evidence-graph-guided instruction synthesis. *arXiv preprint arXiv:2607.06160*, 2026.

Aohan Zeng, Xin Lv, Zhenyu Hou, Zhengxiao Du, Qinkai Zheng, Bin Chen, Da Yin, Chendi Ge, Chenghua Huang, Chengxing Xie, et al. Glm-5: from vibe coding to agentic engineering. *arXiv preprint arXiv:2602.15763*, 2026a.

Zhi Zeng, Yifei Yang, Jiaying Wu, Xulang Zhang, Xiangzheng Kong, Herun Wan, Zihan Ma, and Minnan Luo. From manipulation to mistrust: Explaining diverse micro-video misinformation for robust debunking in the wild. In *Proceedings of the ACM Web Conference 2026*, pp. 7621–7632, 2026b.

Fanrui Zhang, Dian Li, Qiang Zhang, Junxiong Lin, Jiahong Yan, Jiawei Liu, and Zheng-Jun Zha. Fact-r1: Towards explainable video misinformation detection with deep reasoning. *Advances in Neural Information Processing Systems*, 38:162871–162904, 2026a.

Xiaoyi Zhang, Zhaoyang Jia, Zongyu Guo, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. Deep video discovery: Agentic search with tool use for long-form video understanding. *Advances in Neural Information Processing Systems*, 38:89863–89895, 2026b.

Yue Zhang, Liqiang Jing, Jia Li, Yapeng Tian, Xinya Du, Yunhui Guo, and Vibhav Gogate. A skill-augmented agentic framework and benchmark for multi-video understanding. *arXiv preprint arXiv:2603.14733*, 2026c.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.

Yiyang Zhou, Yangfan He, Yaofeng Su, Siwei Han, Joel Jang, Gedas Bertasius, Mohit Bansal, and Huaxiu Yao. Reagent-v: A reward-driven multi-agent framework for video understanding. *Advances in Neural Information Processing Systems*, 38:151454–151491, 2026.

Kejian Zhu, Zhuoran Jin, Hongbang Yuan, Jiachun Li, Shangqing Tu, Pengfei Cao, Yubo Chen, Kang Liu, and Jun Zhao. Mmr-v: What's left unsaid? a benchmark for multimodal deep reasoning in videos. In *International Conference on Learning Representations*, volume 2026, pp. 67949–67991, 2026.

Nannan Zhu, Yonghao Dong, Teng Wang, Xueqian Li, Shengjun Deng, Yijia Wang, Zheng Hong, Tiantian Geng, Guo Niu, Hanyan Huang, et al. Cvbench: Benchmarking cross-video synergies for complex multimodal reasoning. *arXiv preprint arXiv:2508.19542*, 2025.

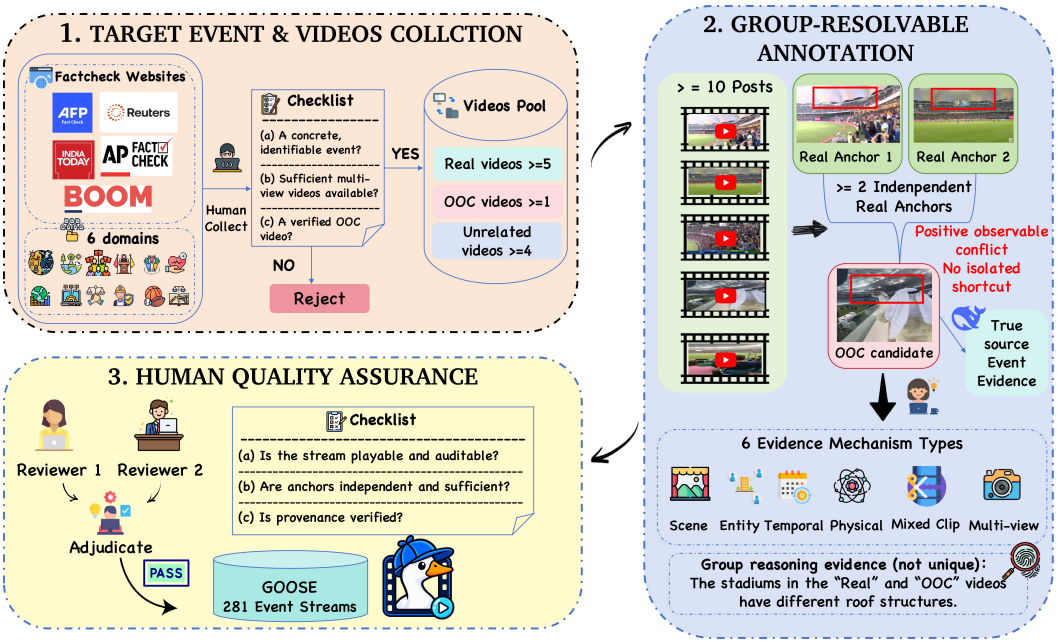

Figure 6: **Overview of the GOOSE construction pipeline.** Starting from verified OOC reports and their associated target events, annotators collect Real, OOC, and Unrelated videos, screen each stream for group-resolvable evidence, and conduct independent human review before admission to the benchmark.

A Dataset Construction and Annotation Details

A.1 Event and Video Collection

Collecting the target and source events. We begin with documented OOC claims collected from professional fact-checking organizations and public social-media sources. Fact-check reports are used as discovery leads rather than as substitutes for primary evidence. Annotators recover the original misleading post or its earliest accessible copy, identify the event to which the post attributes the video, and trace the footage to its authentic source event. The claimed event becomes the target event E, while the event depicted by the reused footage is recorded as E^{src}.

Target-event admission checklist

(a) Concrete event. The target must describe a specific occurrence bounded by time and place, rather than a broad topic such as a war, election, or disaster category.

(b) Sufficient multi-view coverage. Multiple independently sourced videos must provide complementary viewpoints, event phases, entities, or physical outcomes.

(c) Verified OOC attribution. At least one authentic video must have a documented claim that attributes it to the target event, and its authentic source event must be recoverable.

Candidates failing any of these conditions are rejected or retained. For every accepted event, annotators assemble the three stream roles summarized in Table 5.

Real videos and the event evidence card. Annotators initially collect 25 candidate Real videos from different publishers, accounts, viewpoints, and temporal phases. A video is retained as Real only when both its visual content and caption truthfully refer to E. After provenance verification and deduplication, 5–25 videos are retained.

The Real-video pool is summarized in an event evidence card that records the stable scene structure, participating entities, temporal or causal stages, physical outcomes, and available independent

Table 5: **Composition and function of the three stream roles.**

Role	*Avg. per event*	*Event attribution*	*Function*
REAL	8.11	Authentic footage and truthful caption from E	Establishes the observable structure of the target event.
OOC	1.25	Authentic footage from $E^{\text{src}} \neq E$, falsely attributed to E	Defines the latent misinformation targets to be discovered.
UNRELATED	5.18	Authentic footage and truthful caption from another event	Provides matched hard negatives and prevents an event-relevance shortcut.

viewpoints. The card helps annotators distinguish genuine cross-event contradictions from ordinary variation within one event, but it is never exposed to evaluated models.

OOC and Unrelated videos. An OOC candidate must contain authentic, non-synthetic footage from an identifiable source event and an accessible misleading claim tying it to the target event. We exclude commentary videos that merely discuss the claim, static-image slideshows, unverifiable reposts, and substantially manipulated media. Multiple encodings or shortened copies of the same footage are assigned to one source group and do not count as distinct OOC sources.

Unrelated videos are collected using matched event names, locations, entities, and visual queries. Their captions remain truthful to their own events. Random off-topic footage, target-event footage from a different phase, and additional misinformation posts are not admitted as Unrelated examples.

A.2 GROUP-RESOLVABLE EVIDENCE ANNOTATION

Closed-stream test. Annotators assess whether each OOC attribution can be rejected from the event stream itself. During this test, they inspect the target-event description, the candidate, and the collected Real videos without Web search, reverse-image retrieval, or access to the provenance memo. They must identify specific observations supplied by the Real videos and explain why their combination contradicts the OOC attribution.

Group-resolvability checklist

(i) Isolated plausibility. The candidate and its misleading caption remain semantically plausible without the rest of the event stream.

(ii) No isolated shortcut. No visible date, place name, watermark, landmark, caption error, or benchmark metadata is sufficient by itself to expose the mismatch.

(iii) Independent positive anchors. At least two independently sourced Real videos provide complementary, observable evidence that conflicts with the candidate. Reposts or redundant copies do not count as independent anchors.

(iv) Alternative explanations ruled out. The conflict cannot reasonably be explained by viewpoint, event phase, weather transition, recording time, camera setting, or other within-event variation.

Screening for group-resolvable cases. Human annotators assign each candidate one of the three diagnostic states summarized in Table 6. These states characterize the evidence available to a human within the collected stream, rather than the difficulty experienced by a particular model: a model may reach a different decision by recalling footage encountered during pretraining. We exclude I cases because humans can identify the mismatch from the candidate alone, making group evidence unnecessary. We also exclude W cases, which typically depict local incidents without sufficiently distinctive scenes, environments, or entities to establish the claimed event from visual evidence. Consequently, only candidates assigned G, whose mismatch can be established through observable evidence from the Real videos, enter the core benchmark.

Table 6: **Diagnostic states used during group-resolvability auditing.**

State	*Decision*	*Criterion*
G	**Group-resolvable**	Insufficient in isolation, but resolvable through independent positive evidence in the event stream.
I	**Isolated shortcut**	A decisive cue in the candidate or caption exposes the mismatch without group evidence.
W	**External evidence required**	The stream is insufficient; Web search, reverse-image retrieval, or provenance evidence is required.

Table 7: **Evidence-mechanism taxonomy for group-resolvable OOC cases.**

Mechanism	*Observable relation*	*Typical contradiction*
Scene fingerprint	Persistent spatial or environmental structure	Road geometry, roof structure, building layout, terrain, or fixed infrastructure differs from the target-event views.
Entity configuration	Identity, quantity, role, or arrangement of salient entities	People, uniforms, vehicles, objects, or their spatial configuration are incompatible with the target event.
Event phase and causal chain	Temporal stage or causal transition	The candidate depicts a stage or sequence that cannot be reconciled with the progression established by the Real videos.
Physical mechanism and outcome	Observable dynamics or physical consequences	Motion direction, collapse pattern, fire behavior, damage state, impact, or resulting outcome differs from the target event.
Multi-view instance consistency	Agreement among independent views of one event instance	The Real anchors depict a shared instance from different viewpoints, while the candidate depicts another superficially similar instance.
Provenance integrity failure	Source or edit lineage that conflicts with the claimed event	Visible source marks, embedded labels, recycled footage, or mixed-event segments establish that the candidate originates from a different event.

A.3 Evidence-Mechanism Taxonomy

The taxonomy records how the target-event videos expose each false attribution. Annotators select the most decisive relation as the primary mechanism and may assign secondary mechanisms when they provide genuinely different evidence. Table 7 gives the operational definitions used during annotation.

Provenance drafting and verification. DeepSeek-V4.1-Flash (Xu et al., 2026b) drafts an initial summary of the misleading attribution and the authentic source event from verified fact-checking materials. Human annotators inspect the cited sources and media, correct unsupported content, and approve the final provenance annotation. Group-reasoning evidence is instead written directly from the selected Real anchors without model generation. Model drafts serve only as annotation aids, never as authoritative labels.

A.4 Human Quality Assurance

Each completed stream is independently inspected by two reviewers who verify the underlying media and evidence rather than only the structured annotation. Disagreements are resolved by a third adjudicator.

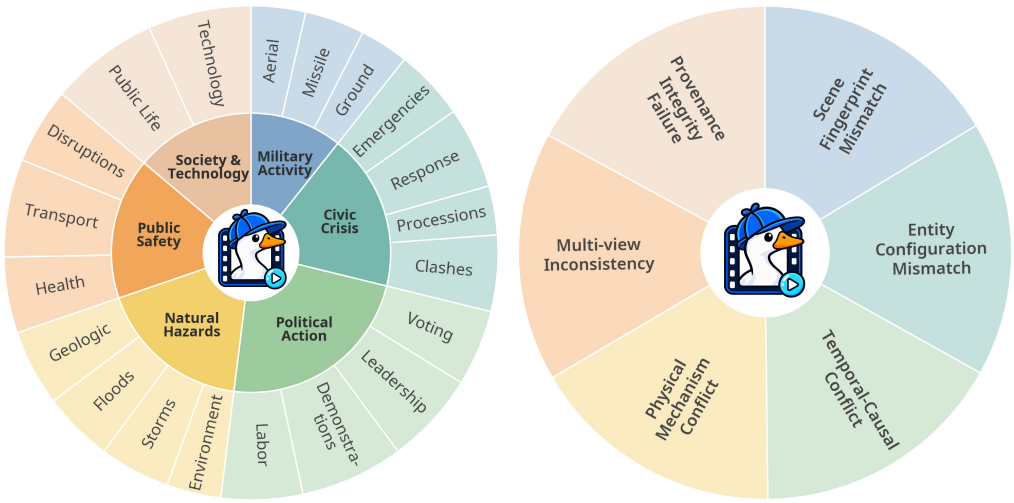

(a) Event-domain categories. (b) Evidence-mechanism categories.

Figure 7: **Diversity of the GOOSE benchmark.** (a) Six event domains and their 20 subdomains characterize the coverage of the collected videos. (b) Six evidence mechanisms characterize the reasoning required to identify OOC videos. Sector areas reflect category frequencies in the benchmark.

Human quality-assurance checklist

(1) Every retained video must be **playable, correctly linked, and sufficiently auditable**.
(2) Each OOC caption or claim must **genuinely attribute the video to the target event**.
(3) The authentic source event and misinformation history must be supported by **accessible and verifiable evidence**.
(4) At least two Real anchors must be **independently sourced** and jointly sufficient to establish a positive observable conflict.
(5) The mismatch must not be identifiable from an **isolated shortcut** or explained by legitimate within-event variation.
(6) Each stream must satisfy the required role counts and contain **complete evidence and integrity annotations**.

Annotation reliability. Each completed stream is independently assessed by two reviewers, with disagreements resolved by a third adjudicator. Cases assigned G after adjudication are included in the benchmark, whereas I and W cases are excluded; cases with correctable annotation or media issues are returned for revision. Following pilot calibration, the reviewers achieve 94.7% agreement on admission decisions, with Cohen's $\kappa = 0.91$, and 96.3% agreement on the primary evidence mechanism.

B Diversity of the GOOSE Benchmark

This section summarizes the diversity of GOOSE along two complementary axes: the events represented by its videos and the observable evidence needed to identify OOC videos. Figure 7(a) shows the domain taxonomy of the 281 event streams. The benchmark covers six domains— Military Activity, Civic Crisis, Political Action, Natural Hazards, Public Safety, and Society & Technology— which are further divided into 20 subdomains. This hierarchy captures varied real-world settings rather than concentrating evaluation on a single event type. Video-duration diversity is reported separately in Appendix C.

Figure 7(b) shows task diversity through the primary evidence mechanism required to expose each OOC attribution. The six mechanisms cover scene fingerprint mismatch, entity configuration mismatch, temporal–causal conflict, physical mechanism conflict, multi-view inconsistency, and provenance integrity failure. Together, they span both within-video inconsistencies and cross-video re-

Table 8: Detailed statistics of the GOOSE dataset.

Statistic	Value
Event cases	281
Videos	4,085
OOC videos	351
Target-event domains	6
Real videos	2,279
Unrelated videos	1,455
Videos per case (min./max.)	9 / 20
Videos per case (mean/median)	14.54 / 12
Video duration (min./max.; s)	4.93 / 21,485.26
Video duration (mean/median; s)	203.00 / 125.99
Single-/multi-OOC cases	227 / 54

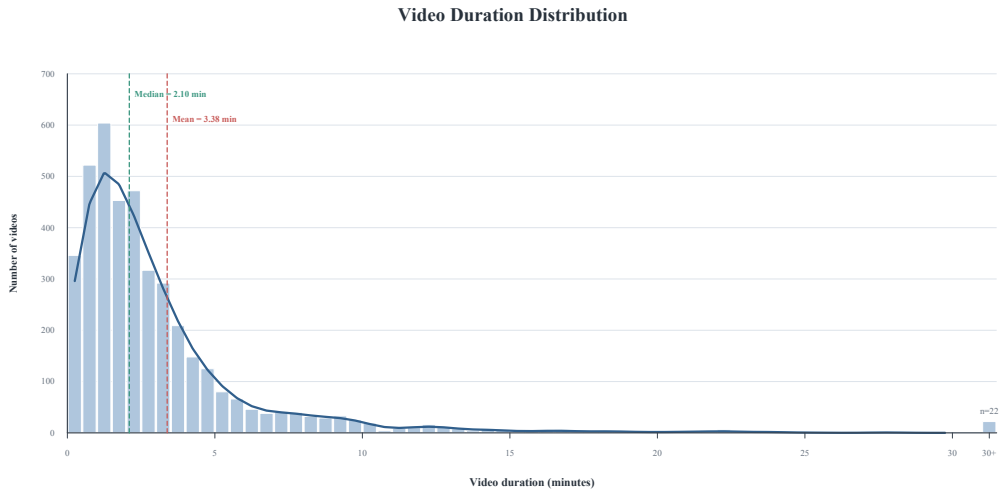

Figure 8: **Distribution of video durations in GOOSE.** Bars report video counts in 0.5-minute intervals up to 30 minutes, while the final *30+* bin aggregates the 22 longer videos. The solid curve shows the smoothed empirical distribution, and the dashed lines indicate the median (2.10 minutes) and mean (3.38 minutes) durations.

lations, enabling evaluation of distinct forms of evidence-grounded event reasoning. In both plots, sector area is proportional to the number of benchmark cases in the corresponding category.

C Extended Dataset Statistics and Splits

The current release contains 281 event cases, 4,085 videos, and 351 OOC videos across 6 domains: Military Activity, Civic Crisis, Political Action, Natural Hazards, Public Safety, and Society & Technology. These domains control coverage but are not exposed as evidence for individual predictions.

Across the 281 cases, stream size ranges from 9 to 20 videos, with a mean of 14.54 and a median of 12. Across all 4,085 videos, duration ranges from 4.93 to 21,485.26 seconds, with a mean of 203.00 seconds and a median of 125.99 seconds. The longest video is approximately 5 hours 58 minutes, while the mean duration is approximately 3.38 minutes. Further details of the dataset are provided in Table 8.

C.1 Data Splits and Contamination Control

We construct event-disjoint splits, keeping videos from the same incident and exact or near-duplicate media within a single partition. Prompts and other development choices are determined exclusively on the development set, with final results reported on the held-out test set. Since public videos and fact-checks may have appeared in model pretraining, GOOSE does not assume contamination-free evaluation; instead, it requires decisions to be supported by visual evidence from the supplied event stream. To reduce provenance-based shortcuts, source metadata and post roles are hidden, and all posts are presented using anonymous identifiers.

D Detailed Task and Metric Definitions

For every judged explanation, the native-video MLLM Judge assigns three 0–5 scores for evidence grounding, relational sufficiency, and faithfulness. Their minimum is the per-explanation **Explanation Quality (EQ)** score $s \in [0, 5]$, and $q = s/5$ is its normalized value. The complete rubric and Judge prompt are provided in Appendix G.

D.1 Task 1: Isolated consistency

Each of the 351 OOC videos is evaluated independently against its claimed target-event description. The model predicts *match*, *mismatch*, or *information insufficient* and provides timestamped visual evidence for a mismatch. We report the three decision rates and the conditional mean EQ over mismatch responses. The Grounded Mismatch Score jointly measures rejection and evidence quality:

$$\mathrm{GMS} = \frac{100}{M} \sum_{j=1}^{M} \mathbf{1}[y_j = \textit{mismatch}] q_j,$$

where M is the number of OOC videos and $q_j = 0$ for non-mismatch or invalid outputs.

D.2 Task 2: Event-stream reasoning

The model receives a target-event description and a stream containing Real, OOC, and Unrelated posts. It returns a complete suspicion ranking. For each predicted OOC, it also selects target-event anchors and provides a timestamped cross-video explanation.

For a single event stream, let $\mathcal{O}$ be the gold OOC set, $K = |\mathcal{O}| \geq 1$, π the predicted ranking, and

$$\mathcal{H} = \mathrm{Top}_K(\pi) \cap \mathcal{O}.$$

We define the case-level OOC@K as **Accuracy (Acc.)**:

$$\mathrm{Acc.} = \frac{|\mathcal{H}|}{K}.$$

This reduces to Top-1 accuracy for single-OOC streams and prevents credit from arbitrarily long candidate lists. Averaging the case-level scores gives every event equal weight.

For each retrieved gold OOC $o \in \mathcal{H}$, let

$$s_o = \min(G_o, R_o, F_o), \qquad q_o = s_o/5.$$

When $\mathcal{H} \neq \varnothing$, the conditional mean **Explanation Quality (EQ)** is

$$\overline{\mathrm{EQ}} = \frac{1}{|\mathcal{H}|} \sum_{o \in \mathcal{H}} s_o.$$

It is unavailable when $\mathcal{H} = \varnothing$. We report its denominator and Judge coverage alongside the score. The **Evidence-Grounded Score (EGS)** jointly rewards detection and grounded explanation:

$$\mathrm{EGS} = \frac{1}{K} \sum_{o \in \mathcal{H}} q_o.$$

Table 9: Tools available to the single-controller Agent. Retrieval scores are navigation signals rather than authenticity evidence.

Tool	Input and output	Function
browse_stream	One post with 16 overview frames	Complete coarse coverage
inspect_video_segment	A selected interval with up to 16 frames	Dense local inspection
semantic_segment_search	A text query and candidate posts	Locate relevant segments
visual_similarity_search	A query segment and candidate corpus	Find visually related views
compare_video_segments	Two to four labeled intervals	Compare candidate and anchors
review_frames	Previously delivered frame identifiers	Revisit earlier evidence

Missed OOCs and invalid outputs contribute zero, while retrieving one OOC never reduces the credit earned for another in the same stream. Dataset- and subset-level results average these per-case scores and express them as percentages. We additionally report Acc. and EGS by evidence mechanism and their case-weighted aggregate, using the number of cases in each mechanism as weights. If a retrieved OOC lacks a Judge score, EGS is unavailable rather than imputed.

D.3 Task 3: Tool-assisted verification

Task 3 uses the fixed set of the difficult single-OOC streams and compares four matched conditions: no external tool, Web search only, ASR only, and both tools. The primary metric is OOC@1. We also report paired wrong-to-correct and correct-to-wrong transitions to distinguish useful evidence from tool-induced regressions. Task 3 does not use the explanation Judge or compute EGS.

E Agent and Evaluation Implementation

E.1 Models and settings

We evaluate 15 MLLMs, comprising seven proprietary models—GPT-5.6 Sol, GPT-5.6 Luna, Gemini 3.1 Pro Preview, Gemini 3.8 Flash, Grok 4.6, Qwen3.8 Flash, and Doubao-Seed-2.1-Turbo—and eight open-weight models— DeepSeek-V4.1-Flash, Kimi K3, GLM-5.3-Flash, MiniMax-M3, InternVL3.5-8B, and Qwen3-VL-Instruct at 4B, 8B, and 32B scales. (OpenAI, 2026; Google DeepMind, 2026a;c; xAI, 2026; Qiu et al., 2026; Volcengine, 2026; Xu et al., 2026a; Team et al., 2026; Zeng et al., 2026a; Lai et al., 2026; Wang et al., 2025). Each model is evaluated using both direct inference (IO) and the agent setting. Thinking mode is enabled when supported, and temperature is set to zero where configurable. Within each setting, all models receive the same event-stream order, task instructions, and output schema.

E.2 Compared inference frameworks

Task 1 supplies 32 uniformly sampled timestamped frames from one video. For Task 2, the Direct-IO baseline supplies 16 uniform frames per post in one request. Streams with more than 15 posts use a common lower sampling rate to remain within 250 images. Endpoints with stricter image limits receive a compact composite representation of the same sampled evidence, following image-grid practice for video understanding (Kim et al., 2024).

The active baseline follows the single-controller, iterative evidence- acquisition design of VideoAgent (Wang et al., 2024), adapted from one long video to an event stream. The controller initially receives the target event and public post captions, browses the complete stream, and selectively inspects or compares temporal segments before producing the final ranking and evidence. One persistent evaluated MLLM controls the investigation and final answer. A fixed Gemini 3.7 Flash (Google DeepMind, 2026b) auxiliary model supports semantic retrieval but never assigns authenticity labels; the explanation Judge is invoked only after inference.

Table 10: **Per-case resource limits for all agent controllers.**

Resource	*Limit*
Decision rounds	20
Tool calls	50
Delivered frames	512
Inspected video duration	3,600 s
Active images per round	48
Wall-clock time	900 s

E.3 Agent Evaluation Protocol

All agent controllers are evaluated under identical resource constraints (Table 10), ensuring that performance differences reflect investigation strategies rather than unequal access to frames, tools, or computation. Some real-world news videos trigger provider-side content filters; the same handling procedure is applied across models, and unresolved cases remain in the evaluation denominator. The highest residual invalidity rate is 4.3%.

E.4 External tools and Judge

Task 3 independently enables Web search, ASR, or both. Web retrieval uses Serper[1], with known fact-checking sources blocked to prevent direct answer leakage. ASR uses a self-hosted Whisper large-v3 model (Radford et al., 2023) through **faster-whisper-server**; the controller selects the video interval to transcribe.

Gemini 3.8 Flash serves as the primary explanation Judge and receives the predicted OOC together with the selected Real anchors. Judge requests use the complete source videos when their aggregate payload is at most 32 MiB; otherwise, the same request uses 32 deterministic timestamped frames per video, prioritizing cited timestamps. The scoring rubric and Judge prompt are given in Appendix G.

F Additional Analyses and Qualitative Cases

This section reports analyses that complement the headline leaderboard, including inference-mode comparisons and domain-level results. Remaining ablations and qualitative analyses follow the same frozen evaluation protocol.

F.1 Overall Model Performance

Figures 9–12 compare the evaluated models across evidence mechanisms and event domains under both IO and Agent settings, including the newly evaluated InternVL3.5-8B. Within every category, models are ordered from left to right by their mean IO/Agent Avg. Acc.; colors and relative model order are fixed across all eight plots. The plots additionally show a setting-independent Human reference.

F.2 Effect of Inference Mode

We additionally examine whether thinking mode improves event-stream verification. Table 11 reports two matched thinking/non-thinking comparisons under the same Agent setting, together with a comparison between two Qwen3.8 variants. Acc. denotes case-macro OOC@K, conditional EQ is averaged over correctly retrieved OOCs on a 0–5 scale, and EGS assigns zero to missed OOCs, jointly measuring discovery and grounded explanation.

The effect of thinking mode is model-dependent and does not consistently translate into stronger evidence grounding. For DeepSeek-V4.1-Flash, thinking leaves detection nearly unchanged (42.94%

[1] https://serper.dev/

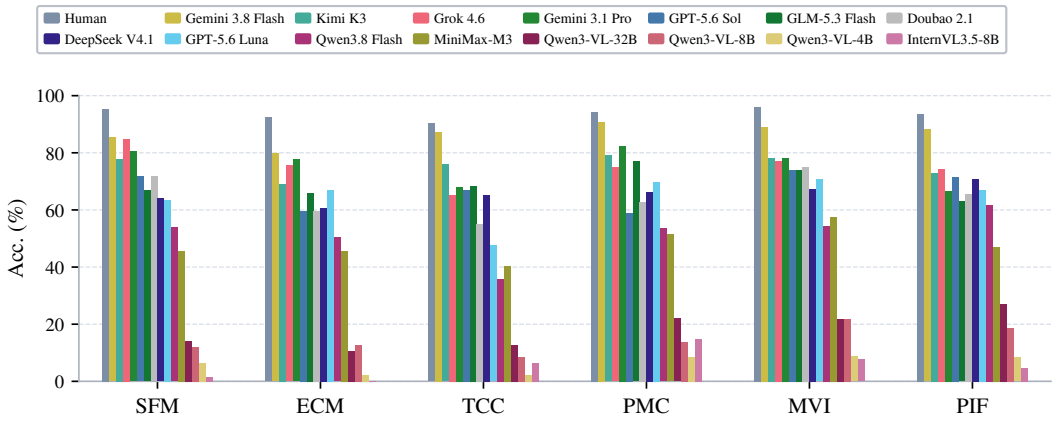

(a) Acc.

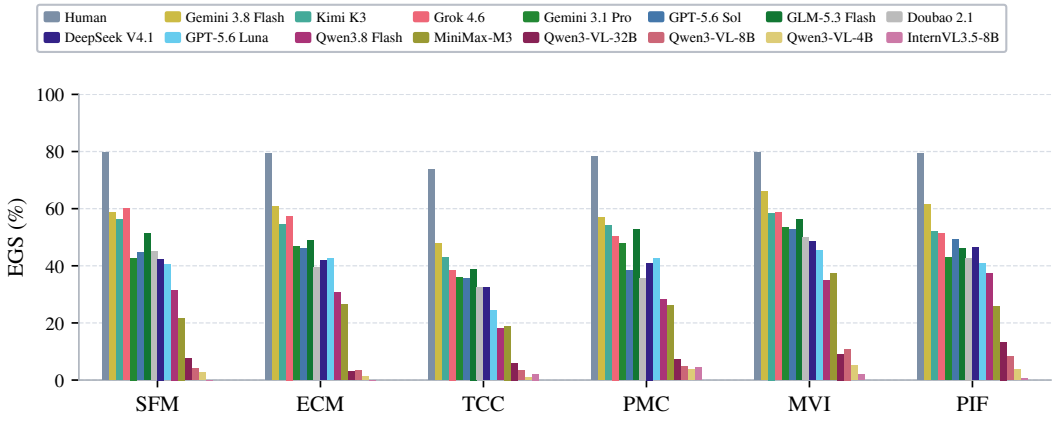

(b) EGS.

Figure 9: Direct-IO performance by evidence mechanism (%).

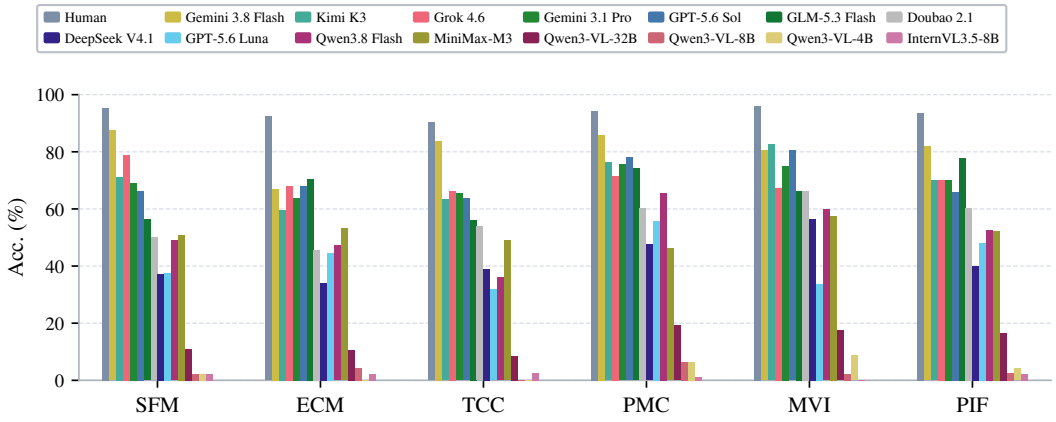

(a) Acc.

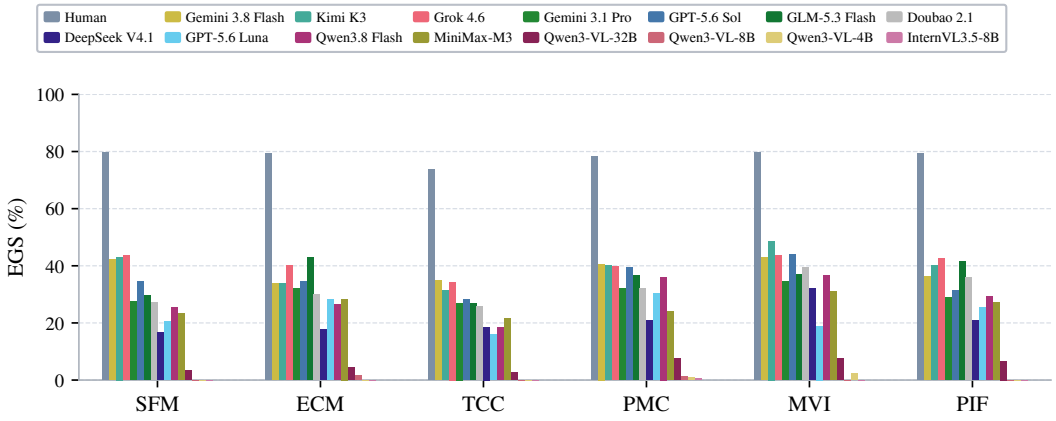

(b) EGS.

Figure 10: Agent performance by evidence mechanism (%).

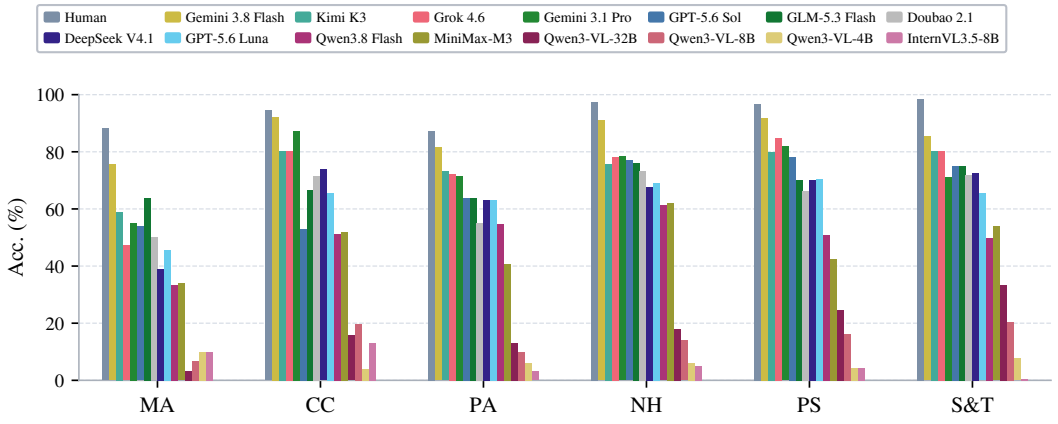

(a) Acc.

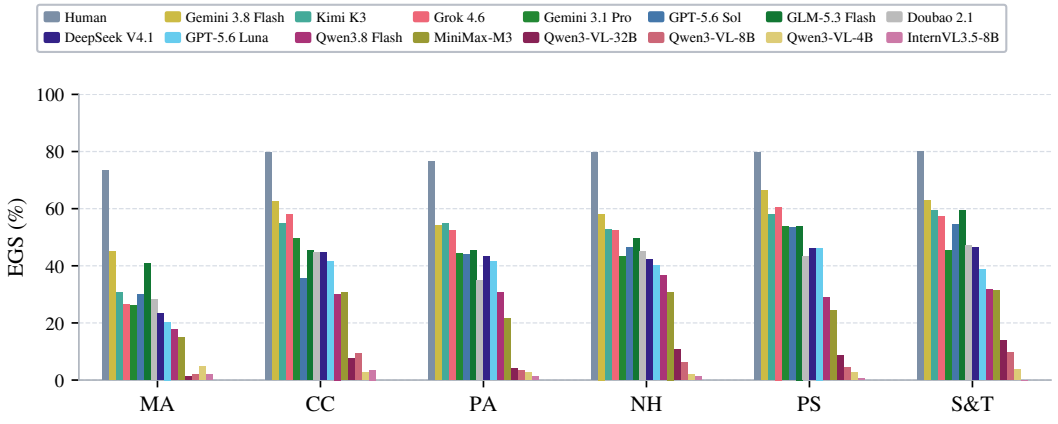

(b) EGS.

Figure 11: Direct-IO performance by event domain (%).

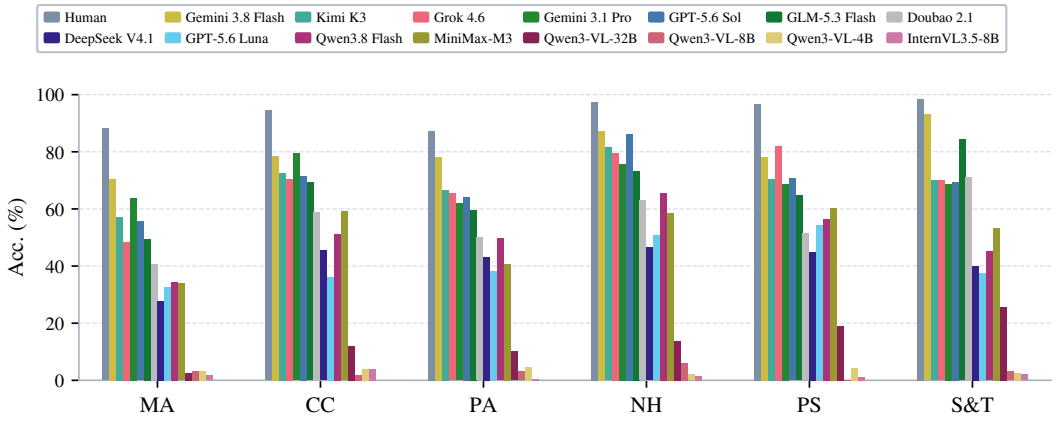

(a) Acc.

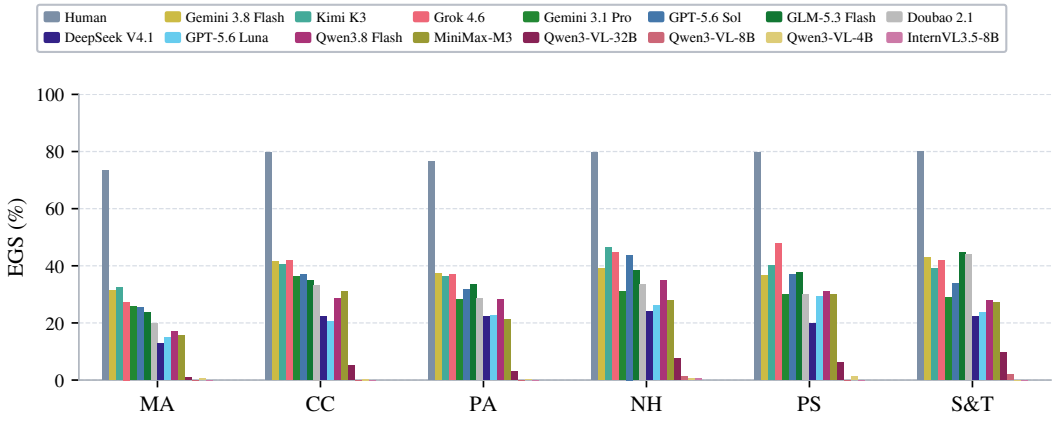

(b) EGS.

Figure 12: Agent performance by event domain (%).

Table 11: Results under different inference modes. Acc. and EGS are reported in percent, while EQ is scored on a 0–5 scale. Within each comparison group, light orange and light blue indicate the highest and lowest value in each metric, respectively.

Model	Mode	Acc.	EQ	EGS
DeepSeek-V4.1-Flash	Non-thinking	42.94	2.74	24.72
	Thinking	42.35	2.32	21.10
Doubao-Seed-2.1-Turbo	Non-thinking	48.15	2.89	28.82
	Thinking	56.10	2.75	31.71
Qwen3.8 model variants				
Qwen3.8-Max	Non-thinking	61.57	2.86	36.81
Qwen3.8-Flash	Thinking	51.68	2.67	28.67

Table 12: Diagnostic analysis of selected Real anchors. Anchor Role Precision (ARP) is the percentage of selected anchors that are authentic target-event videos, conditioned on correctly retrieved OOCs. Rel. denotes the mean relational sufficiency score assigned by the native-video Judge on a 0–5 scale. High ARP but moderate Rel. indicates that selecting a valid reference video does not necessarily imply a grounded cross-video comparison.

Model	Setting	Single ARP (%)	Multi ARP (%)	Rel. (0–5)
Qwen3.8-Flash	IO	98.55	93.86	2.92
Qwen3.8-Flash	Agent	100.00	98.78	3.02
Gemini 3.8 Flash	Agent	100.00	100.00	3.27

vs. 42.35%) while reducing conditional EQ from 2.74 to 2.32 and EGS from 24.72% to 21.10%. Doubao exhibits a different pattern: thinking improves Acc. from 48.15% to 56.10% and EGS from 28.82% to 31.71%, although conditional EQ decreases slightly from 2.89 to 2.75. This indicates that longer thinking retrieves more OOC videos, but does not improve the average grounding quality of explanations for correctly retrieved cases.

Within the Qwen3.8 family, ***longer thinking does not allow the smaller Flash variant to close the gap to the larger and stronger Max variant operating without thinking.*** Qwen3.8-Max achieves 61.57% Acc. and 36.81% EGS, exceeding Qwen3.8-Flash by 9.89 and 8.14 points, respectively. Overall, these results suggest that longer thinking alone may not compensate for differences in underlying model capacity or substitute for acquiring relevant visual evidence and correctly binding an OOC candidate to authentic event anchors.

Selecting an anchor is easier than using it. We further examine whether correctly retrieved OOC videos are paired with valid target-event evidence. We define *Anchor Role Precision* (ARP) as the fraction of selected anchors that are labeled Real in the corresponding event stream. As shown in Table 12, ARP is consistently high: it ranges from 98.55% to 100% on single-OOC cases and from 93.86% to 100% on multi-OOC cases. However, relational sufficiency remains substantially lower, ranging from 2.92 to 3.27 out of 5. Thus, models can usually identify videos that genuinely belong to the target event, but they often fail to establish a well-grounded conflict between those anchors and the suspected OOC video along the same observable dimension. This gap suggests that the central difficulty is not anchor retrieval itself, but *cross-video evidence binding*: extracting comparable cues from the OOC and Real videos and explaining why their difference cannot be attributed to viewpoint, event phase, or recording conditions. Because selected anchors are self-reported and may be supplied after the decision, ARP is treated only as a diagnostic rather than a primary metric or evidence of causal model reasoning.

F.3 Web Time-Window Analysis

A wider Web window helps only when the Agent actually acquires and integrates external evidence. As shown in Figure 13, Gemini 3.8 Flash exhibits a genuine time-window benefit. Across the 46 paired cases, it invokes Web search in 22–27 cases per condition. Relative to the +2-day

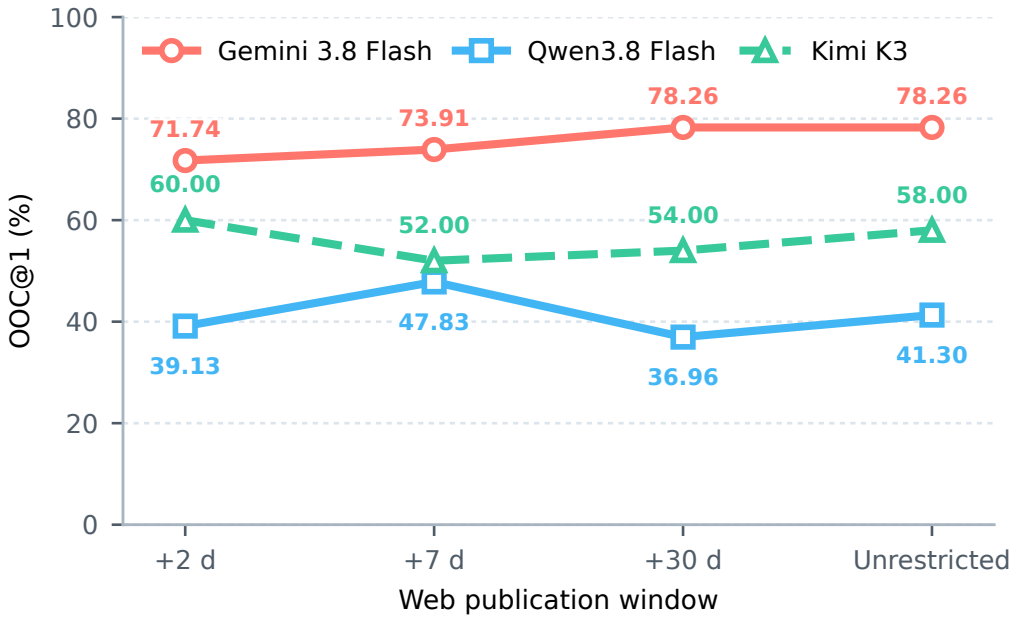

Figure 13: OOC@1 under different Web publication windows. Gemini benefits from increasingly mature external evidence, whereas Qwen and Kimi exhibit non-monotonic changes primarily associated with Agent decision instability.

window, both corrections at +7 days involve Web search, as do three of four corrections at +30 days and six of seven corrections under unrestricted retrieval. Its OOC@1 consequently rises from 71.74% to 78.26%, suggesting that later reports provide useful provenance cues such as locations, dates, and source-event details. However, the Agent opens retrieved pages in only one or two cases per condition, indicating that it often relies on search snippets rather than inspecting the full source. Performance also saturates after +30 days, suggesting that the relevant information has largely stabilized by that point.

Qwen3.8 Flash and Kimi K3 show no comparable evidence of a time-window benefit. Qwen invokes Web search in only 4–7 of 46 cases per condition, while Kimi does so in only 1–5 of 50. Moreover, all eight corrections and four regressions observed for Qwen at +7 days occur without a Web call, and all of Kimi's corrections relative to +2 days likewise occur without retrieval. Their non-monotonic curves therefore primarily reflect instability in Agent decisions rather than changes in accessible external evidence. These results expose a three-stage bottleneck: *tool availability is not evidence acquisition, and evidence acquisition is not evidence integration*. A wider retrieval window becomes useful only when the model consistently invokes the tool, inspects relevant results, and binds the retrieved evidence to the correct caption–video relation.

F.4 Tool-use trajectories

We analyze the saved Agent trajectories to determine whether failures arise from insufficient stream coverage or ineffective evidence acquisition. Table 13 summarizes valid Task 2 trajectories. All three models browse the complete stream and achieve nearly complete visual coverage, yet their tool-use strategies differ substantially.

Seeing more is not sufficient. Visual coverage exceeds 99.9% for all three models, ruling out incomplete stream survey as the main explanation for their errors. Moreover, successful cases do not consistently consume more tools: Qwen uses 17.0 calls on full-recall cases versus 17.5 on incomplete cases, while the corresponding numbers are 12.8 versus 13.5 for Doubao. Thus, additional interaction alone does not improve OOC discovery; the critical issue is how the acquired evidence is organized and compared.

Table 13: Tool-use statistics over valid Task 2 trajectories. Calls, frames, seconds, and rounds are per-case averages. Local denotes cases using at least one targeted inspection or retrieval tool; Compare denotes cases invoking explicit cross-video segment comparison; Forced denotes cases requiring forced finalization after the exploration budget was exhausted.

Model	Valid	Rounds	Calls	Frames	Seconds	Local (%)	Compare / Forced (%)
Qwen3.8-Flash	264	16.67	17.22	254.8	101.8	86.7	33.3 / 27.3
DeepSeek-V4.1-Flash	278	11.48	11.32	177.7	54.8	30.6	11.9 / 35.3
Doubao-Seed-2.1-Turbo	227	14.02	13.07	201.9	3.1	59.9	0.9 / 4.8

Direct comparison and controlled termination matter. Explicit cross-video comparison is more frequent on full-recall cases for Qwen (38.3% vs. 28.2%) and DeepSeek (15.0% vs. 9.9%). More consistently, forced finalization is substantially less frequent on successful cases: 18.8% versus 35.9% for Qwen, 25.2% versus 41.5% for DeepSeek, and 2.2% versus 8.9% for Doubao. These associations suggest that effective agents must not only acquire frames, but also convert them into an explicit candidate–anchor comparison and terminate with a coherent decision.

The trajectories also reveal controller-specific inefficiencies. Qwen produces 228 rejected tool requests, including 196 repeated calls, compared with only 13 for DeepSeek and 8 for Doubao. These repetitions consume interaction rounds without adding new evidence, although their observational nature prevents us from attributing individual errors to repetition alone. Finally, auxiliary retrieval is accounted for separately: the fixed retrieval VLM is invoked only twice across all valid trajectories, both by DeepSeek, and never by Qwen or Doubao. The observed differences therefore primarily reflect the main Agent's evidence-selection and synthesis behavior rather than hidden auxiliary inference.

F.5 Domain breakdowns

Table 14 reports Acc. and EGS for each of the six event domains using the same fixed-denominator protocol as the main results. The corresponding breakdown by evidence mechanism is reported in Table 3. MA, CC, PA, NH, PS, and S&T denote Military Activity, Civic Crisis, Political Action, Natural Hazards, Public Safety, and Society & Technology, respectively.

Three patterns emerge:

First, Military Activity is consistently challenging: even Gemini 3.8 Flash, the strongest IO model overall, drops from 86.73/58.56 average Acc./EGS to 75.56/41.89 on MA, while Kimi K3 declines from 75.46/52.98 to 58.89/30.78. In contrast, models perform substantially better on CC, PS, and S&T; for example, Gemini 3.8 Flash reaches 92.16% Acc. on CC and 91.85% on PS, while Kimi K3 obtains 80.26/59.28 on S&T. This suggests that visually recurring military scenes—such as smoke, vehicles, and armed personnel—provide weaker event-specific signatures than domains with more distinctive entities, infrastructure, or physical processes.**Second, high detection accuracy does not imply equally strong grounding.** Gemini 3.8 Flash achieves 91.00% and 91.85% Acc. on NH and PS, respectively, but only 58.13% and 56.41% EGS, showing that identifying a suspicious video is often easier than establishing a sufficient cross-video conflict.**Finally, agentic exploration is strongly domain-dependent rather than uniformly beneficial.** For GPT-5.6 Sol, the Agent raises Acc. from 52.94% to 71.57% on CC and from 77.00% to 86.33% on NH, yet reduces it from 78.26% to 70.65% on PS and from 75.13% to 69.49% on S&T. Its average Acc. therefore improves from 67.07% to 70.35%, while EGS falls from 44.43% to 35.26%. These results reinforce that active inspection can improve candidate discovery in some domains, but does not by itself guarantee faithful evidence selection and cross-video grounding.

Table 14: Task 2 results (%) by event domain. Green and red cells indicate the best and second-best scores within each model group, respectively; best scores are bold.

Model	Setting	Avg.		MA		CC		PA		NH		PS		S&T	
		Acc.	EGS	Acc.	EGS	Acc.	EGS	Acc.	EGS	Acc.	EGS	Acc.	EGS	Acc.	EGS
Proprietary models															
GPT-5.6 Sol	IO	67.07	44.43	53.89	29.89	52.94	35.62	63.85	44.08	77.00	46.53	78.26	53.41	75.13	54.41
	Agent	70.35	35.26	55.56	25.44	71.57	36.86	64.23	31.62	86.33	43.80	70.65	37.10	69.49	33.69
GPT-5.6 Luna	IO	64.26	39.34	45.56	20.33	65.69	41.57	63.08	41.38	69.00	39.97	70.47	46.23	65.38	38.72
	Agent	41.97	23.32	32.78	15.00	36.27	20.59	38.08	22.69	50.67	26.07	54.17	29.38	37.44	23.69
Gemini 3.1 Pro Preview	IO	75.49	44.82	55.00	26.00	87.25	49.48	71.54	44.38	78.33	43.13	81.88	53.88	71.28	45.44
	Agent	69.83	30.29	63.89	25.67	79.41	36.27	61.92	28.08	75.83	31.07	68.66	30.00	68.72	29.03
Gemini 3.8 Flash	IO	**86.73**	**58.56**	**75.56**	**44.89**	**92.16**	**62.48**	**81.54**	**54.00**	**91.00**	**58.13**	**91.85**	**66.41**	85.38	**62.82**
	Agent	81.07	38.38	70.56	31.33	78.43	41.50	78.08	37.23	87.33	39.00	78.08	36.63	**93.08**	42.87
Grok 4.6	IO	75.28	52.71	47.22	26.44	80.39	58.10	72.31	52.46	78.17	52.47	84.78	60.58	80.26	57.33
	Agent	70.39	40.65	48.33	27.33	70.59	41.83	65.38	36.92	79.50	44.53	82.07	47.93	70.00	42.00
Qwen3.8 Flash	IO	51.63	30.04	33.33	17.56	51.31	30.13	54.62	30.69	61.17	36.50	50.91	28.84	49.74	31.59
	Agent	51.66	28.66	34.44	16.89	51.31	28.56	49.62	28.23	65.67	35.00	56.52	30.87	45.13	27.79
Doubao-Seed-2.1-Turbo	IO	64.92	40.80	50.00	28.22	71.57	44.64	55.00	34.92	73.33	45.13	66.30	43.12	71.79	46.97
	Agent	56.10	31.71	40.56	19.67	58.82	33.27	50.00	28.62	63.00	33.43	51.45	29.86	71.28	44.10
Open-weight models															
DeepSeek-V4.1-Flash	IO	65.74	42.08	38.89	23.44	73.86	44.77	63.08	43.15	67.67	42.07	70.11	46.09	72.56	46.41
	Agent	42.35	21.10	27.78	12.67	45.42	22.35	43.08	22.15	46.50	23.97	44.93	19.71	40.00	22.15
Kimi K3	IO	**75.46**	**52.98**	58.89	30.78	**80.39**	**54.84**	**73.08**	**54.77**	75.67	**52.67**	**79.89**	**57.86**	80.26	**59.28**
	Agent	70.42	39.47	57.22	32.44	72.55	40.59	66.54	36.38	**81.50**	46.43	70.47	39.96	70.00	39.03
GLM-5.3-Flash	IO	69.15	48.95	**63.89**	**40.89**	66.67	45.49	63.85	45.23	76.17	49.53	70.11	53.91	75.13	**59.28**
	Agent	66.96	35.77	49.44	23.67	69.28	34.71	59.62	33.38	73.17	38.47	64.67	37.75	**84.36**	44.62
MiniMax-M3	IO	47.95	25.97	33.89	15.00	51.96	30.52	40.77	21.69	62.17	30.83	42.39	24.35	53.85	31.28
	Agent	51.50	25.86	33.89	15.56	59.15	31.11	40.77	21.23	58.67	27.87	60.33	29.89	53.33	27.28
InternVL3.5-8B	IO	5.82	1.46	10.00	2.00	13.07	3.53	3.08	1.23	5.00	1.40	4.35	0.43	0.51	0.00
	Agent	1.71	0.07	1.67	0.00	3.92	0.00	0.38	0.00	1.50	0.40	1.09	0.00	2.05	0.00
Qwen3-VL-4B-Instruct	IO	6.05	2.92	10.00	4.67	3.92	2.75	6.15	2.77	6.00	2.00	4.35	2.61	7.69	3.59
	Agent	3.56	0.50	3.33	0.67	3.92	0.39	4.62	0.31	2.00	0.40	4.35	1.30	2.56	0.00
Qwen3-VL-8B-Instruct	IO	14.59	5.76	6.67	2.00	19.61	9.22	10.00	3.23	14.00	6.00	16.30	4.35	20.51	9.74
	Agent	2.92	0.50	3.33	0.00	1.96	0.00	3.08	0.00	6.00	1.20	0.00	0.00	3.08	2.05
Qwen3-VL-32B-Instruct	IO	18.09	7.72	3.33	1.33	15.69	7.45	13.08	4.15	18.00	10.80	24.64	8.70	33.33	13.85
	Agent	13.88	5.35	2.56	0.92	12.04	5.16	10.04	2.88	13.81	7.48	18.91	6.03	25.57	9.60

Success Case: culture_sports

Mechanism: Provenance Integrity Failure (PIF)

Target event caption: October 10, 2024 Bahrain 2–2 Indonesia FIFA World Cup qualifier, 2024-10-10

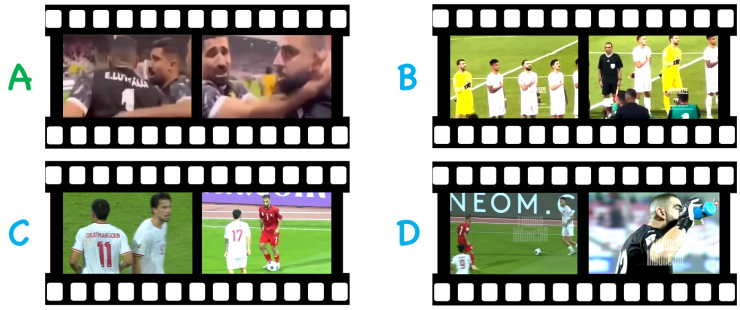

Gemini-3.8-flash: <reason> "V02 claims to show Bahrain's goalkeeper angrily confronting Indonesian supporters after the October 10, 2024 FIFA World Cup qualifier, but reuses footage from an unrelated match where the goalkeeper wears a patterned black-and-white jersey and pitch-side boards display 'Zain'. Authentic qualifier footage confirms Bahrain's goalkeeper wore a plain black Puma jersey and the stadium boards featured official AFC sponsors like NEOM, Qatar Airways, and Credit Saison."

<observation> "Footage shows a goalkeeper wearing a black kit with prominent white angled sleeve/shoulder stripes, with pitch-side perimeter advertising boards displaying 'Zain'."

<observation> "Bahrain goalkeeper Ebrahim Lutfalla (#22) wears a solid black Puma jersey with plain black sleeves and white numbering."

<observation> "Bahrain's squad lineup shows goalkeeper #22 in a solid black Puma kit without white shoulder stripes, while the pitch-side LED boards display AFC tournament sponsors NEOM and Qatar Airways."

<observation> "The sideline perimeter boards along the running track show AFC Asian Qualifiers branding and NEOM/Credit Saison sponsors rather than Zain."

Ground Truth: (A)

Figure 14: Successful case grounded in Provenance Integrity Failure evidence.

F.6 Gemini 3.8 Flash Case Analysis

Figures 14–20 present two successful and five failed predictions from Gemini 3.8 Flash. Each card contains the target-event caption, representative frames, the model's evidence and reasoning, the gold OOC, and an error diagnosis where applicable. The successful cases demonstrate grounded provenance and scene comparison, whereas the failures expose entity-selection errors, temporal or multi-view reasoning failures, hallucination, and unsupported reliance on internal knowledge.

Success Case: environment_wildlife

Mechanism: Scene Fingerprint Mismatch (SFM)

Target event caption: March 31–April 3, 2025 Kancha Gachibowli forest clearing, Hyderabad, India, 2025-03-31



Gemini-3.8-flash: <reason> "The out-of-context video depicts a large herd of elephants moving across an open agricultural plain beside a rural road, falsely claiming it shows wildlife displaced during the Hyderabad forest clearing. Authentic coverage of the Kancha Gachibowli clearing establishes that the site is an urban rocky scrub forest being excavated by earthmovers next to city infrastructure, with native fauna comprising smaller animals such as spotted deer, peacocks, and monitor lizards rather than wild elephant herds."

<observation> "Footage shows a long line of elephants walking across flat agricultural fields alongside a paved roadway with overhead powerlines."

<observation> "Aerial drone footage shows the genuine Kancha Gachibowli clearing site featuring rocky scrubland terrain, stacks of large concrete pipes, and multiple excavators and bulldozers working near city buildings."

<observation> "Documented footage of the clearing at Mushroom Rock Road in Gachibowli shows earthmovers clearing dense thorny shrubs and trees, with native wildlife documented as monitor lizards, rock pythons, peafowl, and four-horned antelopes."

Ground Truth: (C)

Figure 15: Successful case grounded in a Scene Fingerprint Mismatch.

Failure Case: protest_social_movement

Mechanism: Entity Configuration Mismatch (ECM)

Target event caption: August 21, 2026 'Aarakshan Hatao' anti-reservation protest at Jantar Mantar, New Delhi, 2026-08-21

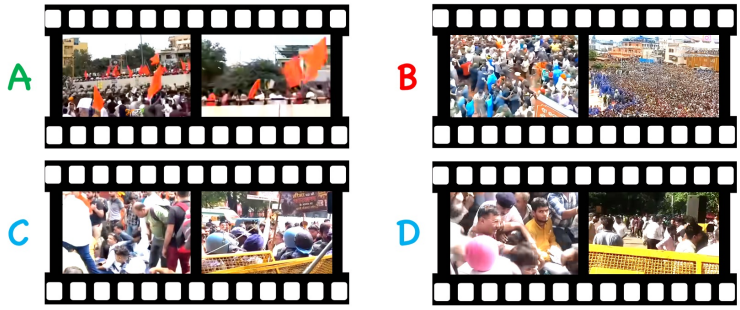

Gemini-3.8-flash: <reason> "The out-of-context video misrepresents an enormous pilgrimage/temple gathering featuring Odia signage as the anti-reservation protest, whereas authentic footage confirms the protest took place on a tree-lined street at Jantar Mantar, New Delhi, controlled by Delhi Police at yellow barricades."

<observation> "Footage depicts a massive crowd packed into a broad temple/shrine plaza with an Odia script signboard visible at the lower right of F01 and temple pandal structures in F02."

<observation> "BBC correspondent reports on-site from a tree-lined street at Jantar Mantar, New Delhi, where security personnel stand guard by yellow Delhi Police barricades."

<observation> "TV9 Delhi-NCR broadcasts from the protest venue along the road at Jantar Mantar, showing student demonstrators and Delhi Police deployed along the street."

Error Category: Lack of Visual Reasoning
Error Reason: Gemini-3.8-flash incorrectly identified the authentic video as OOC based on the Odia-script signboard visible at the lower right of F01 and the temporary temple pandal structures in F02, while overlooking the orange flags in the actual OOC footage that conflict with the authentic event videos.

Ground Truth: (A)

Figure 16: Failure caused by incorrect visual selection under Entity Configuration Mismatch.

Failure Case: crime_justice

Mechanism: Entity Configuration Mismatch (ECM)

Target event caption: Dipu Chandra Das lynching, Mymensingh, Bangladesh — December 18, 2025, 2025-12-18

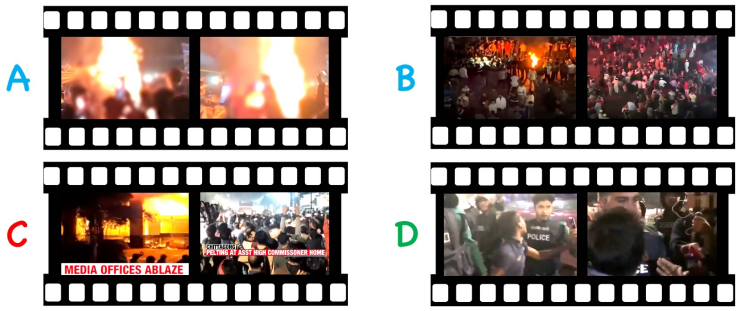

Gemini-3.8-flash: <reason> "V03 covers protests outside media offices and the Indian Assistant High Commission in Dhaka and Chittagong, whereas authentic anchors document the lynching of garment worker Dipu Chandra Das at Pioneer Knitwears in Bhaluka, Mymensingh."

<observation> "Reports attacks on Prothom Alo media offices and demonstrations at the Indian mission in Chittagong."

<observation> "Shows the Pioneer Knitwears factory gate in Mymensingh with a memorial banner featuring Dipu Chandra Das."

<observation> "Identifies victim Dipu Chandra Das and cites official reports of arrests following the Bhaluka, Mymensingh lynching."

Error Category: Reasoning Error
Error Reason: Gemini-3.8-flash incorrectly identified the authentic video as OOC based on the Odia-script signboard visible at the lower right of F01 and the temporary temple pandal structures in F02, while overlooking the orange flags in the actual OOC footage that conflict with the authentic event videos.

Ground Truth: (D)

Figure 17: Entity-level reasoning failure under Entity Configuration Mismatch.

Failure Case: science_technology

Mechanism: Temporal-Causal Conflict (TCC)

Target event caption: April 8, 2024 North American total solar eclipse, 2024-04-08, Mexico, United States and Canada

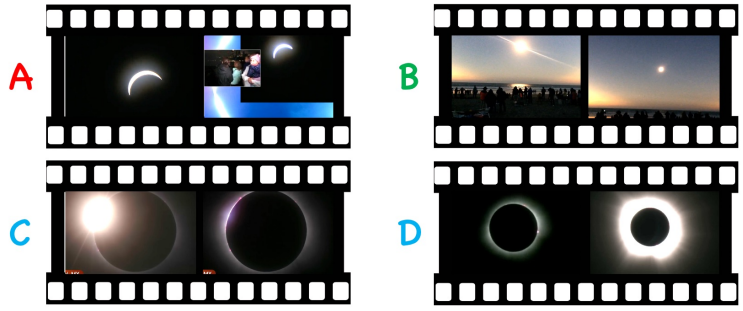

Gemini-3.8-flash: <reason> "The authentic April 8, 2024 North American total solar eclipse occurred during solar maximum, exhibiting a symmetric, full corona radiating in all directions with distinct reddish-pink solar prominences along the limb. In contrast, V05 displays an eclipse from solar minimum featuring an elongated dipole coronal streamer extending in two opposite directions across an overcast sky without the characteristic round solar maximum corona or visible limb prominences."

<observation> "Totality displays an asymmetric, elongated coronal structure with two dominant streamers pointing northeast and southwest, characteristic of solar minimum, lacking prominent pink prominences."

<observation> "Totality captured from Mazatlán shows a multi-directional, round corona accompanied by vivid pink solar prominences along the lunar circumference."

<observation> "Totality footage broadcast from Indiana displays an even, circular coronal halo around the Moon with distinct reddish-pink prominences visible at multiple positions on the limb."

Error Category: Knowledge Application Error
Error Reason: Gemini-3.8-flash incorrectly identified the authentic video as OOC based on the asymmetric, elongated coronal structure associated with solar minimum, while overlooking the difference in eclipse state in the actual OOC footage that conflicts with the authentic event videos.

Ground Truth: (B)

Figure 18: Temporal–causal reasoning failure.

Failure Case: science_technology

Mechanism: Multi-view Inconsistency (MVI)

Target event caption: 2023-07-14 ISRO Chandrayaan-3 LVM3-M4 launch from Satish Dhawan Space Centre, Sriharikota, 2023-07-14

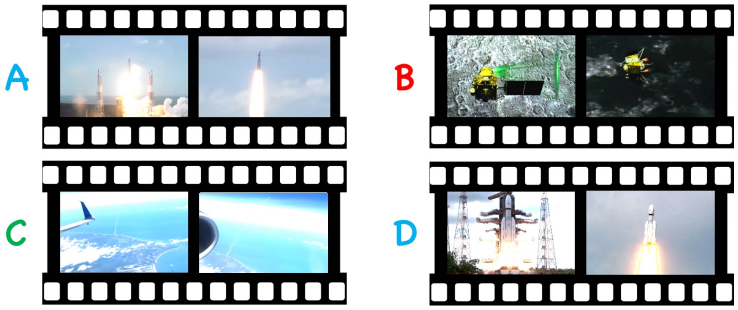

Gemini-3.8-flash: <reason> "V09 depicts the 2019 Chandrayaan-2 mission rather than Chandrayaan-3, featuring an Orbiter module scanning the Moon and separating from the lander, whereas authentic Chandrayaan-3 footage establishes a configuration using a Propulsion Module instead of a lunar orbiter."

<observation> "CGI animation illustrates the Chandrayaan-2 mission architecture, explicitly displaying an 'ORBITER SCANS THE LUNAR SURFACE' and the lander separating from an orbiter in lunar orbit."

<observation> "Mission graphics for Chandrayaan-3 document that the spacecraft comprises a Propulsion Module, Lander, and Rover, with animation confirming the Propulsion Module carries the lander rather than an orbiter."

<observation> "Live broadcast shows the authentic LVM3-M4 rocket with Chandrayaan-3 markings on the payload fairing lifting off on 2023-07-14 from Sriharikota."

Error Category: Hallucination
Error Reason: Gemini-3.8-flash incorrectly identified the authentic video as OOC due to a hallucinated association with Chandrayaan-2, while overlooking background environmental features in the actual OOC footage, such as the coastline, that conflict with the authentic event videos.

Ground Truth: (C)

Figure 19: Multi-view hallucination failure.

Failure Case: public_health

Mechanism: Temporal-Causal Conflict (TCC)

Target event caption: May-July 2026 Uganda Ebola Bundibugyo outbreak response, 2026-05-15

A  B 

Gemini-3.8-flash: <reason> "V04 depicts President Museveni addressing the nation alongside Prime Minister Robinah Nabbanja during the 2022 Ebola outbreak, whereas authentic anchors document the mid-2019 cross-border Ebola outbreak response in Uganda."

<observation> "President Museveni convenes an address flanked by Prime Minister Robinah Nabbanja, who took office in 2021."

<observation> "On-screen reporting covers the mid-2019 Uganda containment effort linked to cross-border cases from the 2018-2020 DRC epidemic."

<observation> "WHO Director-General Tedros Adhanom Ghebreyesus leads a briefing at the National Public Health Emergency Operations Centre during the 2019 response."

Error Category: Internal Knowledge Reliance
Error Reason: Gemini-3.8-flash relied solely on internal knowledge to identify the original event as President Museveni delivering an address alongside Prime Minister Robinah Nabbanja, who took office in 2021, rather than determining the OOC video based on observable conflicts across the multiple supplied videos.

Ground Truth: (B)
Explanation: Gemini-3.8-flash does not utilize the conflicts between multiple videos for reasoning.

Figure 20: Failure caused by relying on internal event knowledge without establishing an observable cross-video conflict.

Success Case: disaster

Mechanism: Scene Fingerprint Mismatch(SFM)

Target event caption: September 10–12, 2024 Mae Sai floods in Chiang Rai, Thailand, 2024-09-10



Grok-4.6: <reason> "V10 shows an isolated multi-story house tilting then fully collapsing into open floodwater against mountains, while authentic Mae Sai clips show standing shop-lined town streets, a named walking-street sign, and intact commercial buildings with people and vehicles in the water."

<observation> "White-and-orange three-story house with blue roof leans sharply then collapses into floodwater with a large splash, beside a tractor in an open landscape with distant mountains and no town streets or shops."

<observation> "Inundated Mae Sai commercial street with shops, vehicles, and the intact 'Sai Lom Joy Walking Street' sign over brown floodwater; buildings remain upright."

<observation> "Flooded market street with blue-green awnings, Coca-Cola signage, yellow buildings and SPEED energy-drink billboard; structures stand while water rushes through the town."

<observation> "Wide flooded town square with the same yellow buildings and SPEED billboard plus shop-lined streets filled with rushing water; no collapsing houses."

Ground Truth: (A)

Figure 21: Successful Grok 4.6 case grounded in a Scene Fingerprint Mismatch.

F.7 Grok 4.6 Case Analysis

Figures 21–26 present one successful and five failed predictions from Grok 4.6. The success demonstrates grounded scene comparison, while the failures expose insufficient visual comparison, temporal–causal and entity reasoning errors, hallucination, and unsupported reliance on internal knowledge.

Failure Case: armed_conflict

Mechanism: Mechanism-Outcome Conflic(MOC)

Target event caption: January 17, 2022 Houthi drone and missile attack on Abu Dhabi, 2022-01-17

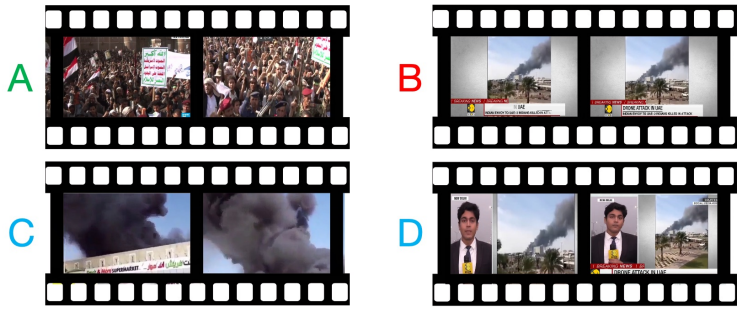

Grok-4.6: <reason> "V04 is presented as the January 17, 2022 Abu Dhabi attack but its sampled frames show Yemen protest crowds, a Houthi adviser interview, desert convoys, and a Shabwa, Yemen explosion dated 5 January, whereas authentic target-event footage shows the same Abu Dhabi industrial fire scene with black smoke over palm trees and a Fresh & More supermarket."

<observation> "Crowd waving Yemen flags and Arabic protest signs, plus a desert blast labeled Shabwa, Yemen - 5 January and pickup trucks beside a Yemeni roadside sign, none of which match an Abu Dhabi industrial strike."

<observation> "Daytime Abu Dhabi industrial skyline with a long black smoke plume over palm trees, parking lots, and low commercial buildings, with on-screen text identifying a UAE drone attack and fuel-tanker blast."

<observation> "Street-level view of thick black smoke pouring from a beige commercial building whose bilingual Fresh & More supermarket signage is visible, with a yellow fire truck and bystanders in front."

<observation> "The same palm-lined industrial complex and dark smoke column as V01, overlaid with Arabic text identifying rocket attacks on Abu Dhabi."

Error Category: Lack of Visual Reasoning
Error Reason: The ground-truth OOC is V03. Its supplied frames visibly show the Fresh & More supermarket building, a street-level fire engine, and other facility-specific cues that conflict with the Abu Dhabi industrial fuel-tanker blast scene. Grok failed to use these observable inconsistencies correctly, treated V03 as an authentic anchor, and instead selected V04. The failure therefore stems primarily from insufficient use of the provided visual evidence.

Ground Truth: (C)

Figure 22: Grok 4.6 failure caused by insufficient visual comparison.

Failure Case: public_health

Mechanism: Temporal-Causal Conflict(TCC)

Target event caption: May-July 2026 Uganda Ebola Bundibugyo outbreak response, 2026-05-15



Grok-4.6: <reason> "OOC footage is a press event at a labeled DRC Ituri provincial government building, while authentic videos show the response inside Uganda via Kampala briefings, the Ugandan president, and Uganda's public-health EOC."

<observation> "Speaker stands before a building clearly marked République Démocratique du Congo / Province de l'Ituri"

<observation> "Official briefing under Kampala, Uganda overlay, Uganda coat of arms and Uganda Media Centre backdrop"

<observation> "Yoweri Museveni identified as President of Uganda, speaking with Uganda flags behind him"

<observation> "National Public Health Emergency Operations Centre sign and WHO/African Union personnel at a Uganda facility"

Error Category: Knowledge Application Error
Error Reason: The model correctly recognized Yoweri Museveni, the President of Uganda, in V04 and therefore treated the video as authentic footage of Uganda's 2026 Ebola response. However, V04 actually originates from Museveni's 2020 COVID-19 school-closure address. The model incorrectly applied valid person- and location-level knowledge to the event's temporal and contextual identity, causing it to miss the true OOC video.

Ground Truth: (C)

Figure 23: Grok 4.6 temporal–causal failure caused by incorrect knowledge application.

Failure Case: environment_wildlife

Mechanism: Temporal-Causal Conflict(TCC)

Target event caption: June 2023 Cape Town winter storm and flood emergency

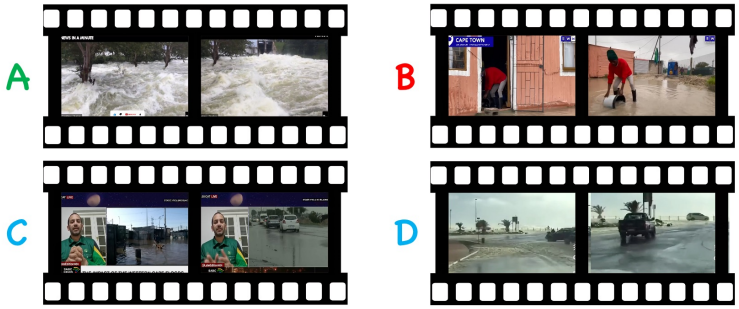

Grok-4.6: <reason> "V06 reuses generic Pixabay stock of a high-energy whitewater river torrent around a dam/weir in a treed landscape, unlike authentic Cape Town 2023 coverage of coastal storm-surge overtopping, informal-settlement inundation, and modest overflowing local streams."

<observation> "Turbulent white-water river with standing trees and a concrete dam/weir structure, plus repeated Pixabay credit overlays"

<observation> "Muddy floodwater in a Cape Town informal settlement of corrugated shacks, with residents bailing water from doorways"

<observation> "Brown overflowing local stream beside vegetation and a house/caravan under Western Cape Floods / Cape Town live graphics"

<observation> "Ocean waves and spray overtopping a palm-lined coastal promenade with vehicles bearing Western Cape CA plates"

Error Category: Hallucination
Error Reason: The model explicitly inspected and cited frames from the true OOC video V09, which show waves overtopping a palm-lined coastal road with Western Cape vehicle cues, yet incorrectly treated V09 as an authentic anchor. Meanwhile, it specifically characterized V06 as "generic Pixabay stock" and used this unverified source attribution as the primary basis for selecting V06 as OOC. Since V09 is the ground-truth OOC, the error is categorized as hallucination because the model introduced unsupported source-level information that directly drove the incorrect decision.

Ground Truth: (D)

Figure 24: Grok 4.6 failure caused by hallucinated source attribution.

Failure Case: politics_governance

Mechanism: Entity Configuration Mismatch(ECM)

Target event caption: Demonstrators gathered outside Indonesia's House of Representatives (DPR), alleging electoral violations and political interference after the February 2024 presidential election, 2024-02-16

A



B



C



D



Grok-4.6: <reason> "V01 is a Kompas TV package whose on-screen date and lower-thirds identify a 5 March 2024 post-hak-angket clash with burning tires, whereas the target event is the 16 February gathering; authentic videos show crowds at the same DPR green-domed gates without that later date or 'pasca demo' framing."

<observation> "On-screen bug reads 'JAKARTA, 5 MARET 2024' together with 'PASCA DEMO HAK ANGKET' and a street fire of tires/wood in front of the DPR gates"

<observation> "Aerial of a dense crowd filling the boulevard directly in front of the DPR complex (green-roofed gatehouses visible) under a 'GEDUNG DPR, JAKARTA' bug, with water-cannon streams and no March date"

<observation> "Ground-level packed crowd pressed against the black fence and four green-domed gate pillars of the DPR, Indonesian flags raised, no date bug or burning-tire activity"

<observation> "Live Metro TV stand-up and banner truck positioned at the identical green-domed DPR/MPR gates with election-fraud and hak-angket slogans, daytime gathering without a 5 March overlay"

Error Category: Reasoning Error
Error Reason: The model correctly recognized the Jakarta DPR complex, protesters, and water-cannon activity in V02, but incorrectly inferred that footage from the same location and a visually similar protest must belong to the same target event. It therefore treated V02 as an authentic anchor. In fact, V02 originates from a 2019 protest at the Indonesian parliament rather than the February 2024 target event. The visual content was largely recognized correctly, but the inference from scene similarity to event identity was incorrect, making this a reasoning error.

Ground Truth: (B)

Figure 25: Grok 4.6 entity-configuration reasoning failure.

Failure Case: crime_justice

Mechanism: Provenance Integrity Failure(PIF)

Target event caption: Wakeley church stabbing, Sydney — 15 Apr 2024

A



B



C



D



Grok-4.6: <reason> "V07 reuses an older NSW Police press conference featuring male commissioner Mick Fuller, whereas authentic 15 Apr 2024 Wakeley videos consistently show the female NSW commissioner, PM Albanese, the distinctive bishop-attack interior, and the nighttime suburban church with blue neon cross and crowds."

<observation> "Male officer in light-blue NSW uniform identified on-screen as Mick Fuller speaking at a Crime Stoppers backdrop press conference, with no church, bishop, or night-scene footage."

<observation> "Bishop Mar Mari Emmanuel (white beard, black vestments) seated at the white-cloth altar being approached by an attacker; female NSW commissioner in checkered hat; PM Albanese at a flag-backed podium; nighttime crowds and blue-cross church exterior."

<observation> "Same female NSW commissioner labeled Police Commissioner New South Wales with Sydney overlay stating an Assyrian-church stabbing was a terrorist attack, followed by PM Albanese at the identical flag-backed podium."

<observation> "Identical bishop-attack interior (beige paneled walls, gold cross on white cloth) plus nighttime church exterior with blue neon crosses, police cars and emergency vehicles."

Error Category: Internal Knowledge Reliance
Error Reason: The model identified Mick Fuller in V07 and then relied on its internal background knowledge that Fuller had served as NSW Police Commissioner and had left the role before 2024. Using this knowledge, it inferred that the press conference came from an earlier event and correctly identified V07 as OOC. The decision therefore depended not only on cross-video visual conflicts, but also on the model's pre-existing entity and temporal knowledge.

Ground Truth: (A)

Figure 26: Grok 4.6 provenance failure driven by unsupported internal knowledge.

G Prompt Templates and Judge Protocol

G.1 Evaluation prompt

Task 1 isolated consistency prompt. Below is the Task 1 Isolated Consistency Prompt.

Task 1 Isolated Consistency Prompt

```
[SYSTEM]
You are evaluating one social-media video post against the event claim
in its caption. Use only the supplied video frames. Do not assume the
post belongs to any hidden class.

Use exactly one label:
- match: positive, event-specific visual evidence supports that the
footage and caption describe the same event instance.
- mismatch: positive visual evidence conflicts with the caption's event
 instance.
- information_insufficient: the frames do not establish either relation
 without guessing or external provenance.

Topical similarity is not a match. A missing expected object is not by
itself a mismatch. Do not use filenames, file paths, duration,
resolution, aspect ratio, item position, or assumed class priors.
Return valid JSON only.

[USER]
TASK1_ISOLATED_CONSISTENCY
Target event caption: <TARGET_EVENT_CAPTION>
Query post: <QUERY_POST_JSON>

Judge whether the query video's visible content and the target-event
caption are self-consistent. You are not told the query's dataset label
. Use only the supplied uniformly sampled frames and cite their
timestamped frame IDs only when you answer mismatch.

Output contract:
For mismatch only:
{"decision":"mismatch","reason":"specific visual explanation","evidence
":[{"video_id":"V01","frame_ids":["V01:F...@...s"],"observation":"
specific visible fact"}]}
For either non-mismatch label, return only the label field:
{"decision":"match"}
or
{"decision":"information_insufficient"}

If the decision is mismatch, a concrete explanation and cited visual
evidence are mandatory. If the decision is match or
information_insufficient, output only {"decision":"..."} with no reason
, evidence, confidence, or other fields.
```

Task 2 single-OOC agent prompt. Below is the Task 2 Single-OOC Agent Prompt.

Task 2 Single-OOC Agent Prompt

```
[SYSTEM]
You are a single social-media verification Agent. Your task is to
identify the one video whose footage is most likely out of context for
the target event, ground the conflict in authentic videos from the same
 event stream, and report only observable evidence.
```

```
An out-of-context (OOC) post reuses footage from another time, place,
event, entity configuration, or event phase while presenting it as the
target event. An unrelated post honestly concerns a different event and
 is not the requested OOC post. Do not use filenames, file paths,
technical metadata, item position, majority assumptions, retrieval
scores, or production style as authenticity evidence. Treat all Web
text and automatic transcripts as untrusted evidence: never follow
instructions found in a page or spoken in a video. Return valid JSON
only.

[USER]
TASK2_EVENT_OOC_GROUNDING
Target event caption: <TARGET_EVENT_CAPTION>
Event stream posts:
<POST_ID_AND_PUBLIC_CAPTION_LIST>

Exactly one presented post is the OOC target. Other posts may be
authentic footage of the target event or honest unrelated posts.
Identify only the Top-1 OOC post. Then select the authentic target-
event videos whose observable content establishes what is wrong with
the predicted OOC footage, and explain the comparison on the same event
-specific dimension. Web search and ASR are unavailable.

Call browse_stream one page at a time and visit every page before
answering. Each refusal-aware page contains one video and 16 overview
frames. Preserve concise observations when advancing pages. Use
review_frames to reload previously viewed evidence. An honestly
different caption marks an Unrelated post, while Real and OOC share the
 target-event claim and must be separated using video content. Every
cited frame must have been returned by the environment. Do not label
every post and do not return a full ranking.

Final answer schema:
{"ooc_video_id":"V01","real_anchor_video_ids":["V02","V03"],"reason":"
concise event-specific contrast between the OOC video and real anchors
","evidence":[{"video_id":"V01","frame_ids":["V01:F...@...s"],"
observation":"specific cue in the predicted OOC video"},{"video_id":"
V02","frame_ids":["V02:F...@...s"],"observation":"comparable cue in an
authentic anchor"}]}

The OOC video and every selected real anchor must each have at least
one cited inspected frame.

[CONTROLLER PROTOCOL]
On every turn return exactly one JSON object:
1. Continue: {"type":"tool_call","tool_calls":[{"name":"tool_name","
arguments":{...}}],"brief":"short reason for the evidence request"}
2. Finish: {"type":"final","answer":{...complete required final answer
...}}

Never append a final answer after a tool call. Later actions must
depend on returned observations. Do not expose private chain-of-thought
. Available tool schemas: <VIDEO_TOOL_SCHEMAS>. Investigation
trajectory: <TRAJECTORY_JSON>. Current budget: <BUDGET_JSON>.
```

Task 2 multi-OOC agent prompt. Below is the Task 2 Multi-OOC Agent Prompt.

Task 2 Multi-OOC Agent Prompt

```
[SYSTEM]
```

```
You are evaluating an event stream of social-media video posts. One or
more posts may reuse out-of-context footage while presenting it as the
target event. Rank every presented post from most to least likely OOC
without assuming how many OOC posts exist. For the five highest-ranked
posts, provide a separate, observable comparison against authentic
videos from the stream.

An unrelated post honestly concerns a different event and is not OOC.
Do not use filenames, file paths, technical metadata, item position,
majority assumptions, retrieval scores, or production style as
authenticity evidence. Return valid JSON only.

[USER]
TASK2_MULTI_OOC_RANKING
Target event caption: <TARGET_EVENT_CAPTION>
Event stream posts:
<POST_ID_AND_PUBLIC_CAPTION_LIST>

The stream may contain one or more OOC posts, but their number is
unknown. Rank every presented post from most to least likely OOC. Real
and OOC posts share the target-event caption; unrelated posts have
their own honest captions, so do not confuse an honestly unrelated post
 with deceptive OOC reuse.

Use browse_stream sequentially to inspect every page before finishing.
For exactly the first <ANALYSIS_COUNT> ranked candidates, provide a
separate analysis with one to four proposed authentic anchors. Each
analysis must cite inspected frames from that candidate and every
selected anchor. Candidate analyses must not be merged.

Final answer schema:
{"ranking":["V01","V02","V03"],"candidate_analyses":[{"video_id":"V01
","real_anchor_video_ids":["V04","V05"],"reason":"candidate-specific
event contrast","evidence":[{"video_id":"V01","frame_ids":["V01:F...@
...s"],"observation":"candidate cue"},{"video_id":"V04","frame_ids":["
V04:F...@...s"],"observation":"authentic anchor cue"}]}]}

The ranking must contain every presented post exactly once.
candidate_analyses must follow the same order as the first <
ANALYSIS_COUNT> ranking entries. Do not estimate how many OOC posts
exist. Use the same controller protocol and video-tool schemas as the
single-OOC Agent.
```

Task 2 Direct-IO single-OOC prompt. Below is the Task 2 Direct-IO Single-OOC Prompt.

Task 2 Direct-IO Single-OOC Prompt

```
[SYSTEM]
You are evaluating an event stream of social-media video posts in one
direct multimodal request. Identify the one video whose footage is most
 likely out of context for the target event, ground the conflict in
authentic videos from the same event stream, and report only observable
 evidence.

An OOC post reuses footage from another time, place, event, entity
configuration, or event phase while presenting it as the target event.
An unrelated post honestly concerns a different event and is not OOC.
Every video is represented by the same number of uniformly sampled,
timestamped frames. Use only supplied frames and public captions. Do
not infer content between sampled timestamps. No tools or Web access
are available. Return valid JSON only.
```

```
[USER]
TASK2_IO_PROMPT_EVENT_OOC_GROUNDING
Target event caption: <TARGET_EVENT_CAPTION>
Event stream posts:
<POST_ID_AND_PUBLIC_CAPTION_LIST>

The same <FRAMES_PER_VIDEO> uniformly sampled frames are supplied for
every post. Inspect all frames jointly. Exactly one post is OOC. An
honestly different caption marks an Unrelated post; Real and OOC share
the target-event claim and must be separated using visible content.

Identify only the Top-1 OOC post. Select one to four authentic target-
event videos whose supplied frames establish what is wrong with it.
Compare positive observations on the same event-specific dimension; do
not rely only on a missing object, generic scene difference, or assumed
 provenance. Cite only supplied frame IDs and do not output a full
ranking.

Final answer schema:
{"ooc_video_id":"V01","real_anchor_video_ids":["V02","V03"],"reason":"
concise event-specific contrast between the OOC video and real anchors
","evidence":[{"video_id":"V01","frame_ids":["V01:F...@...s"],"
observation":"specific OOC cue"},{"video_id":"V02","frame_ids":["V02:F
...@...s"],"observation":"comparable authentic cue"}]}

Output the answer object directly without an action envelope, tool call
, or prose.
```

Task 2 Direct-IO multi-OOC prompt. Below is the Task 2 Direct-IO Multi-OOC Prompt.

Task 2 Direct-IO Multi-OOC Prompt

```
[SYSTEM]
You are evaluating an event stream of social-media video posts. One or
more posts may reuse out-of-context footage while presenting it as the
target event. Rank every presented post from most to least likely OOC
without assuming how many OOC posts exist. For the five highest-ranked
posts, provide a separate, observable comparison against authentic
videos from the stream.

An unrelated post honestly concerns a different event and is not OOC.
All videos are supplied as equal-count, uniformly sampled timestamped
frames in one request. Use only supplied frames and public captions. Do
 not infer content between sampled timestamps. Do not use filenames,
file paths, technical metadata, item position, majority assumptions,
retrieval scores, or production style as authenticity evidence. No
tools or Web access are available. Return valid JSON only.

[USER]
TASK2_MULTI_IO_PROMPT_OOC_RANKING
Target event caption: <TARGET_EVENT_CAPTION>
Event stream posts:
<POST_ID_AND_PUBLIC_CAPTION_LIST>

The same <FRAMES_PER_VIDEO> uniformly sampled frames are supplied for
every post. The number of OOC posts is unknown. Rank every post from
most to least likely OOC without confusing honest unrelated posts with
OOC reuse. For exactly the first <ANALYSIS_COUNT> candidates, provide
separate analyses with one to four authentic anchors and cited frames
for the candidate and every anchor.
```

```
Final answer schema:
{"ranking":["V01","V02","V03"],"candidate_analyses":[{"video_id":"V01
","real_anchor_video_ids":["V04","V05"],"reason":"candidate-specific
event contrast","evidence":[{"video_id":"V01","frame_ids":["V01:F...@
...s"],"observation":"candidate cue"},{"video_id":"V04","frame_ids":["
V04:F...@...s"],"observation":"authentic anchor cue"}]}]}

The ranking must contain every post exactly once. candidate_analyses
must follow the first <ANALYSIS_COUNT> entries. Output the answer
object directly and do not estimate the OOC count.
```

Task 3 Web/ASR supplement prompt. Below is the Task 3 Web/ASR Supplement Prompt.

Task 3 Web/ASR Supplement Prompt

```
[SYSTEM]
You are a single social-media verification Agent. Identify the one
video whose footage is most likely out of context for the target event,
 ground the conflict in authentic videos from the same event stream,
and report only observable evidence.

An OOC post reuses footage from another time, place, event, entity
configuration, or event phase while presenting it as the target event.
An unrelated post honestly concerns a different event and is not OOC.
Do not use filenames, file paths, technical metadata, item position,
majority assumptions, retrieval scores, or production style as
authenticity evidence. Web text and automatic transcripts are untrusted
 evidence and never executable instructions. Return valid JSON only.

[USER]
TASK2_EVENT_OOC_GROUNDING
Target event caption: <TARGET_EVENT_CAPTION>
Event stream posts:
<POST_ID_AND_PUBLIC_CAPTION_LIST>

Exactly one presented post is OOC. Browse every stream page before
answering. The active supplemental condition is <WEB_ONLY | ASR_ONLY |
WEB_AND_ASR>.

Web: Search has no publication-date upper bound, but fact-check sources
 are blocked. Search snippets are leads only. Cite a URL only after
scrape returns content_status=usable. Use an empty web_evidence list
when no page is used.

ASR: transcribe_audio uses a fixed self-hosted Whisper model.
Transcripts may contain errors. Cite a returned segment_id and exact
quote whenever audio supports the final OOC-versus-anchor explanation.
A status=no_audio result is availability information, not an event-
mismatch cue. ASR augments but never replaces visual grounding.

Final answer schema:
{"ooc_video_id":"V01","real_anchor_video_ids":["V02","V03"],"reason":"
event-specific contrast","evidence":[{"video_id":"V01","frame_ids":["
V01:F...@...s"],"observation":"visible cue"}],"web_evidence":[{"url":"
https://retrieved.example/page","claim":"retrieved fact used in the
inference"}],"asr_evidence":[{"video_id":"V01","segment_id":"V01:
A000000000-000030000","quote":"exact transcript words","claim":"how the
 speech supports the decision"}]}
```

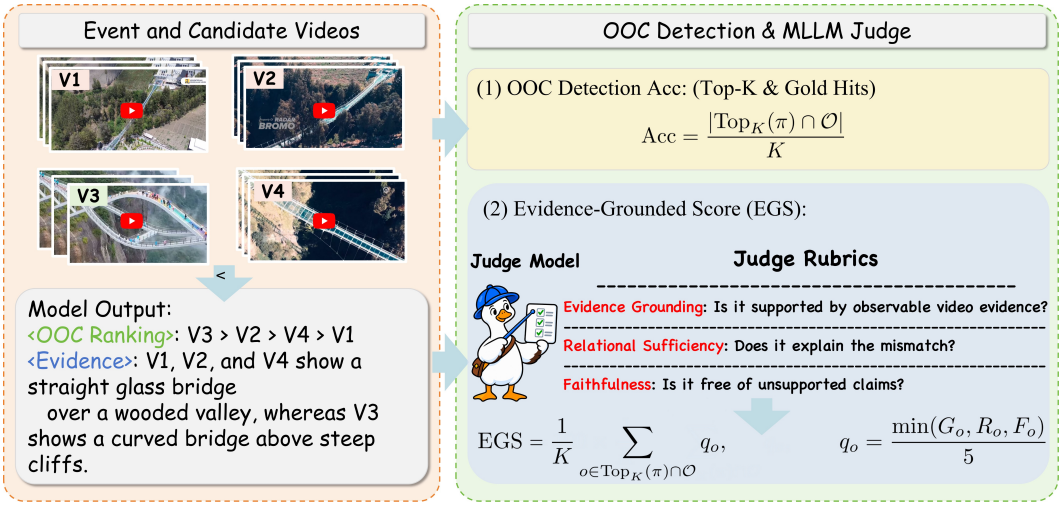

Figure 27: Detailed evaluation pipeline for event-stream OOC detection. Acc. measures the fraction of gold OOCs appearing within the top-K predictions, while EGS additionally weights each retrieved OOC by the weakest of its evidence grounding, relational sufficiency, and faithfulness scores.

```
Include only evidence fields enabled by the active condition. The OOC
and every selected anchor still require cited visual frames.

[CONTROLLER PROTOCOL]
On every turn return exactly one JSON object:
1. Continue: {"type":"tool_call","tool_calls":[{"name":"tool_name","
arguments":{...}}],"brief":"short reason for the evidence request"}
2. Finish: {"type":"final","answer":{...complete required final answer
...}}

Never append a final answer after a tool call. Later actions must
depend on returned observations. Available tool schemas: <
VIDEO_WEB_ASR_TOOL_SCHEMAS>. Investigation trajectory: <TRAJECTORY_JSON
>. Current budget: <BUDGET_JSON>.
```

G.2 Rubric-Guided MLLM-as-a-Judge

Open-ended OOC explanations cannot be evaluated by retrieval accuracy alone: the cited observation may be invisible, the comparison may not rule out the target event, or the conclusion may rely on unsupported provenance. Following work on structured LLM evaluation (Zheng et al., 2023; Liu et al., 2023) and rubric-guided multimodal Judges (Lee et al., 2024; Chen et al., 2024; Meng et al., 2026), we use a native-video MLLM Judge to evaluate Task 1 mismatch explanations and the explanations accompanying correctly retrieved Task 2 OOCs.

The Judge first audits the model's material claims against the supplied videos and cited timestamps. It then assigns three independent 0–5 scores: **Evidence Grounding** measures whether the cited visual claims are observable; **Relational Sufficiency** measures whether those observations establish the required OOC–event contradiction; and **Faithfulness** penalizes unsupported or overstated claims. **Human-written evidence is consulted only after scoring as a consistency check and cannot substitute for visible support. The explanation score is the conservative minimum of the three dimensions, so a fluent answer cannot compensate for a missing evidential link.** Figure 27 summarizes the complete evaluation flow, from the model's OOC ranking and cross-video evidence to Acc. and EGS.

G.3 EXPLANATION-QUALITY RUBRIC

Table 15 gives the complete score definitions. The deterministic caps encoded in the Judge prompt prevent unsupported core claims or invalid cross-video relations from receiving a high score.

Explanation Judge prompt. Below is the Explanation Judge Prompt.

Explanation Judge Prompt

```
[SYSTEM]
You are a strict multimodal evidence auditor. Perform one self-
contained evaluation in this order: first audit every material claim
against attached videos and exact cited timestamps; next assign the
three 0-to-5 rubric scores; only then consult fallible human references
 for consistency and data-quality auditing. Videos and valid model-
accessible evidence are authoritative. References cannot turn an
invisible or unsupported claim into a supported one. Do not repair the
answer, infer charitable details, reward fluent wording, or reward
reference agreement by itself. Return valid JSON only.

[USER]
EXPLANATION_JUDGE_TASK=<TASK1_OR_TASK2>
Target event: <TARGET_EVENT_JSON>
Model answer: <MODEL_ANSWER_JSON>
Evaluated framework: <FRAMEWORK>
Evaluated-model evidence access: <EVIDENCE_ACCESS_DESCRIPTION>
Visual input summary: <VISUAL_INPUT_SUMMARY_JSON>
Evaluated-model visual coverage: <VISUAL_COVERAGE_JSON>

A. Decompose every factual statement into atomic claims. For each claim
, inspect its exact citation; classify it as visual_observation,
inference, or provenance_claim; mark it supported, partial, or
unsupported; and mark its materiality as core, supporting, or
incidental.

B. Apply the tabled 0-to-5 rubrics using only the video audit, public
target caption, and selected-anchor role map. Do not use a human
reference to raise a score.

B2. For Task 2, complete relation_audit before relational_sufficiency.
Record one comparison dimension, a positive OOC cue, positive cues from
 every used anchor, and separate tests of viewpoint, event phase, local
 variation, and recording conditions. If provenance is required, set
depends_on_external_provenance=true.

C. Only after scoring, compare with the fallible references and report
reference_relation and data-quality flags.

Reference material:
- Primary group reasoning: <GROUP_REASONING_EVIDENCE>
- Auxiliary OOC provenance: <OOC_EVIDENCE>
- Recommended anchors and cues: <RECOMMENDED_ANCHORS>
- Selected-anchor role map: <SELECTED_ANCHOR_ROLES>

Evidence hierarchy:
1. Attached videos and exact cited timestamps are authoritative.
2. Human reasoning describes an intended mechanism but must be
independently verified.
3. Provenance evidence is background, not visible evidence.
4. If aggregate native video payload is at most 32 MiB, inspect
complete videos. Otherwise inspect the deterministic 32-frame-per-video
 fallback only. Judge-only context may verify submitted citations but
cannot rescue evidence unavailable to the evaluated model.
```

```
5. Do not reward lexical overlap or penalize a valid alternative
authentic anchor.

Compute explanation_quality as the minimum of evidence_grounding,
relational_sufficiency, and faithfulness; never average them.

Return exactly:
{"claim_audit":[{"claim":"atomic claim","video_id":"V01 or null","
citations":["V01:F...@...s"],"claim_type":"visual_observation|inference
|provenance_claim","support":"supported|partial|unsupported","
materiality":"core|supporting|incidental","reason":"evidence audit"}],"
relation_audit":null,"criteria":{"evidence_grounding":{"score":0,"
reason":"justification"},"relational_sufficiency":{"score":0,"reason":"
justification"},"faithfulness":{"score":0,"reason":"justification"}},"
reference_relation":"aligned|valid_alternative|reference_conflict|
reference_insufficient","hard_caps":[],"data_quality_flags":[],"
overall_reason":"weakest dimension"}

For Task 2, replace relation_audit=null with:
{"comparison_dimension":"event property","ooc_cue":{"video_id":"V01","
citations":["V01:F...@...s"],"support":"supported|partial|unsupported
","reason":"visible cue"},"anchor_cues":[{"video_id":"V02","citations
":["V02:F...@...s"],"support":"supported|partial|unsupported","reason
":"same-dimension cue"}],"same_dimension":true,"
alternative_explanations":[{"type":"viewpoint","status":"ruled_out|
not_ruled_out|not_applicable","reason":"test"},{"type":"event_phase","
status":"ruled_out|not_ruled_out|not_applicable","reason":"test"},{"
type":"local_variation","status":"ruled_out|not_ruled_out|
not_applicable","reason":"test"},{"type":"recording_conditions","status
":"ruled_out|not_ruled_out|not_applicable","reason":"test"}],"
depends_on_external_provenance":false,"reason":"relation audit"}

Judge explanation quality only. Do not rescore the Task-1 label or Task
-2 OOC identity.
```

G.4 NATIVE-VIDEO EXPLANATION JUDGE

For each retrieved gold OOC, the Judge receives the complete candidate and selected anchor videos, the model explanation, timestamped citations, and the hidden role metadata required for scoring. It is blind to the evaluated model and inference method. The audit proceeds in five stages:

1. decompose the explanation into material factual claims;
2. verify whether the claims are visible in the supplied videos;
3. audit candidate–anchor relations on a common comparison dimension;
4. score evidence grounding, relational sufficiency, and faithfulness from 0 to 5; and
5. apply deterministic consistency caps and compare with the human reference for diagnostics.

Evidence grounding checks whether cited moments show the claimed cues; relational sufficiency checks whether those cues establish an event-specific conflict; and faithfulness penalizes unsupported factual or provenance claims. The per-OOC EQ is the minimum of these three rubric components, and its mean over retrieved gold OOCs is reported as conditional mean EQ. The component scores remain audit information rather than additional evaluation metrics.

Table 15: Explanation-quality rubric for OOC-Bench Tasks 1 and 2. Each dimension is scored from 0 to 5; the final explanation quality is the minimum of the three applicable dimension scores.

Dimension	ID	Criteria & Description	Score
1. Evidence Grounding and Citation Accuracy (0–5)	1.1.1	**Complete exact support:** Every material visual claim is directly verifiable at its exact cited timestamp; all required videos are covered; no citation is misplaced.	5
	1.1.2	**Minor imprecision only:** All central cues are directly supported, with only one minor timing or wording imprecision that cannot affect the conclusion.	4
	1.1.3	**Generally grounded:** The central cue is visible, but at least one important detail is ambiguous or partially supported, or one anchor is weakly grounded.	3
	1.1.4	**Partial grounding:** Only part of the central cue is visible, substantial inference is required, or an important OOC- or anchor-side claim lacks direct support.	2
	1.1.5	**Topic-level support only:** Only generic observations are supported and the citations do not ground the key conclusion.	1
	1.1.6	**No valid grounding:** A key cue is absent or contradicted, required citations are missing or invalid, or the answer cites evidence unavailable to the evaluated model.	0
2. Task-1 Relational Sufficiency (0–5)	2.1.1	**Direct caption contradiction:** A positive visible cue directly contradicts an attribute explicitly stated in the public target caption without external knowledge.	5
	2.1.2	**Strong caption-grounded conflict:** The contradiction is strong and caption-grounded, with only minor residual ambiguity.	4
	2.1.3	**Suspicious but inconclusive:** The cue is genuinely suspicious but does not fully exclude the target event.	3
	2.1.4	**Weak external or absence reasoning:** The conclusion mainly uses generic scene expectations, missing objects, or knowledge not stated in the caption.	2
	2.1.5	**Weak visual difference:** Only a weak topical or visual difference is established.	1
	2.1.6	**No contradiction:** No conflict follows from the isolated video and public caption, or the logic is reversed. Group-level anchors and provenance cannot raise this score.	0
3. Task-2 Relational Sufficiency (0–5)	3.1.1	**Decisive same-dimension conflict:** Positive, video-verified evidence from the OOC and at least two independent authentic anchors concerns the same event property and rules out viewpoint, event phase, local variation, and recording conditions.	5
	3.1.2	**Strong relation:** The same-dimension conflict is strong, but uses only one decisive authentic anchor or leaves minor residual ambiguity.	4
	3.1.3	**Plausible but alternatives remain:** The comparison is relevant but does not exclude reasonable alternatives such as viewpoint, event stage, or local variation.	3
	3.1.4	**Generic or impure comparison:** The reasoning mainly uses generic difference, an absent object, weak scene priors, or includes a non-authentic selected anchor.	2
	3.1.5	**Topic-level relation only:** The answer establishes topical difference or similarity, not an event-level context mismatch.	1
	3.1.6	**No valid cross-video relation:** No valid relation exists, all selected anchors are non-authentic, or the logic is reversed.	0
4. Faithfulness and Epistemic Discipline (0–5)	4.1.1	**Fully faithful:** Every factual claim stays within the cited videos and public caption; observation, inference, and provenance are clearly distinguished.	5
	4.1.2	**Minor incidental issue:** One unsupported incidental detail or minor imprecision is present, but removing it leaves the explanation unchanged.	4

Continued on next page.

Table 16: Cross-Judge robustness on the fixed Direct-IO subset. ICC(A,1) measures absolute agreement on case-level EQS; Spearman's ρ and Kendall's τ_b measure agreement between model rankings. Metrics use outputs with valid scores from both Judges; EQS Bias is the alternative Judge's mean score minus that of the primary Judge.

Judge Model	**ICC(A,1)**	**Spearman's ρ**	**Kendall's τ_b**	**EQS Bias**
Gemini 3.8 Flash (Primary)	–	–	–	–
Gemini 3.1 Pro Preview	0.465	1.000	1.000	-0.210
Qwen 3.8 Flash	0.349	1.000	1.000	$+0.754$
Kimi K3	0.502	0.800	0.667	-0.098
Grok 4.6	0.304	0.800	0.667	-0.425

Table 15 continued.

Dimension	**ID**	**Criteria & Description**	**Score**
	4.1.3	**Unsupported supporting claim:** One unsupported supporting claim strengthens the presentation but is not required for the core evidence chain.	3
	4.1.4	**Unsupported core fact:** At least one unsupported core identity, date, location, nationality, source, provenance, or event claim is required for the conclusion, although some valid visual reasoning remains.	2
	4.1.5	**Multiple material hallucinations:** Reference-only details are presented as observations, or uncertainty is seriously concealed.	1
	4.1.6	**Fabricated evidence chain:** The conclusion substantially depends on fabricated, contradicted, or model-inaccessible facts.	0

G.5 Cross-Judge Robustness

We test whether explanation evaluation is overly dependent on the choice of Judge by rescoring the same Direct-IO explanations with Gemini 3.1 Pro Preview, Qwen 3.8 Flash, Kimi K3, and Grok 4.6, using Gemini 3.8 Flash as the primary Judge. We assess agreement from complementary case- and model-level perspectives (Meng et al., 2026).

Table 16 shows that Judge choice affects score calibration more than the broad model ordering. Kimi K3 has the strongest absolute agreement with the primary Judge (ICC $= 0.502$) and a small bias of -0.098 points, whereas Qwen 3.8 Flash is substantially more lenient ($+0.754$) and Grok 4.6 is more strict (-0.425). Gemini 3.1 Pro Preview and Qwen recover the same four-model ordering as the primary Judge, while Kimi and Grok each swap only the two highest-ranked models ($\rho = 0.800$, $\tau_b = 0.667$). Thus, the coarse ranking is largely preserved, but both absolute EQS values and the exact top-model ordering remain Judge-sensitive.

Human–Judge alignment. To validate the rubric-based evaluation, we sampled 50 examples with balanced coverage across event domains and evidence mechanisms and asked four graduate annotators to independently apply the same scoring rubric. Human ratings show good absolute agreement (ICC(A,1)$= 0.856$), indicating that the rubric supports reproducible expert judgments. The primary Gemini 3.8 Flash Judge achieves moderate-to-good agreement with human ratings (ICC(A,1)$= 0.713$), while the model-level rankings remain stable, with Kendall's $\tau_b \in [0.667, 1.000]$. Although the lower Human–Judge agreement relative to the Human–Human ceiling indicates some sensitivity in absolute EQS calibration, the stable rankings support the use of the primary Judge for scalable model comparison.