

INFACT: A Diagnostic Benchmark for Induced Faithfulness and Factuality Hallucinations in Video-LLMs

Junqi Yang^{1,2}, Yuecong Min¹, Jie Zhang^{1*}, Shiguang Shan¹, Xilin Chen¹

¹State Key Laboratory of AI Safety,
Institute of Computing Technology, Chinese Academy of Sciences

²School of Advanced Interdisciplinary Sciences, UCAS

yangjunqi23@mailsucas.ac.cn, {minyuecong, zhangjie, sgshan, xlchen}@ict.ac.cn

Abstract

Despite rapid progress, Video Large Language Models (Video-LLMs) remain unreliable due to hallucinations, which are outputs that contradict either video evidence (faithfulness) or verifiable world knowledge (factuality). Existing benchmarks provide limited coverage of factuality hallucinations and predominantly evaluate models only in clean settings. We introduce INFACT, a diagnostic benchmark comprising 9,800 QA instances with fine-grained taxonomies for faithfulness and factuality, spanning real and synthetic videos. INFACT evaluates models in four modes: Base (clean), Visual Degradation, Evidence Corruption, and Temporal Intervention for order-sensitive items. Reliability under induced modes is quantified using Resist Rate (RR) and Temporal Sensitivity Score (TSS). Experiments on 14 representative Video-LLMs reveal that higher Base-mode accuracy does not reliably translate to higher reliability in the induced modes, with evidence corruption reducing stability and temporal intervention yielding the largest degradation. Notably, many open-source baselines exhibit near-zero TSS on factuality, indicating pronounced temporal inertia on order-sensitive questions.

1 Introduction

Video Large Language Models (Video-LLMs) (OpenAI, 2025; Gemini, 2025; Bai et al., 2025a; Wang et al., 2025; Liu et al., 2025b) have made rapid progress in video understanding in recent years, demonstrating impressive capabilities across broad tasks (Li et al., 2024; Fu et al., 2025; Shangguan et al., 2025; Shafique et al., 2025; Kulkarni et al., 2024; Lin et al., 2026). Despite these advances, the reliability of Video-LLMs in downstream applications is compromised by hallucinations (Li et al., 2025a; Zhang et al., 2024), generating content that contradicts the provided

video evidence (faithfulness) or verifiable world knowledge (factuality).

Recent benchmarks have begun to probe hallucinations in Video-LLMs, e.g., by targeting event/motion-centric failures (Zhang et al., 2024; Kong et al., 2025), using controlled contrastive setups (Li et al., 2025a), or studying prior-driven shortcuts (Bae et al., 2025). However, existing efforts predominantly emphasize video-verifiable inconsistencies, leaving factuality hallucinations substantially under-explored. Moreover, high performance in clean scenarios does not guarantee low hallucination rates, as models may exploit shortcuts such as language priors or static cues. This motivates the evaluation of hallucinations beyond clean settings through controlled evidence perturbations.

To bridge these gaps, we introduce INFACT, a diagnostic benchmark for evaluating Video-LLMs hallucinations regarding faithfulness and factuality in both clean and noisy scenarios. Specifically, INFACT establishes fine-grained taxonomies and comprises 9,800 QA instances spanning real and synthetic videos, covering varying temporal dynamics for faithfulness and diverse knowledge categories for factuality. Furthermore, it supports four evaluation modes: Base (I), Visual Degradation (II), Evidence Corruption (III), and Temporal Intervention (IV) for order-sensitive items.

The non-Base modes apply controlled video perturbations while keeping questions fixed. Modes II–III are invariant-label settings and are evaluated by Resist Rate (RR), which measures whether correct Base decisions remain stable under visual degradation or corrupted evidence. Mode IV disrupts the temporal structure required for correctness and is evaluated using Temporal Sensitivity Score (TSS), which measures whether the model’s predictions changes after shuffling or reversal. Our evaluation reveals that reliability under induced conditions is not uniform: models tend to be more fragile when exposed to misleading evidence than to purely per-

*Corresponding author.

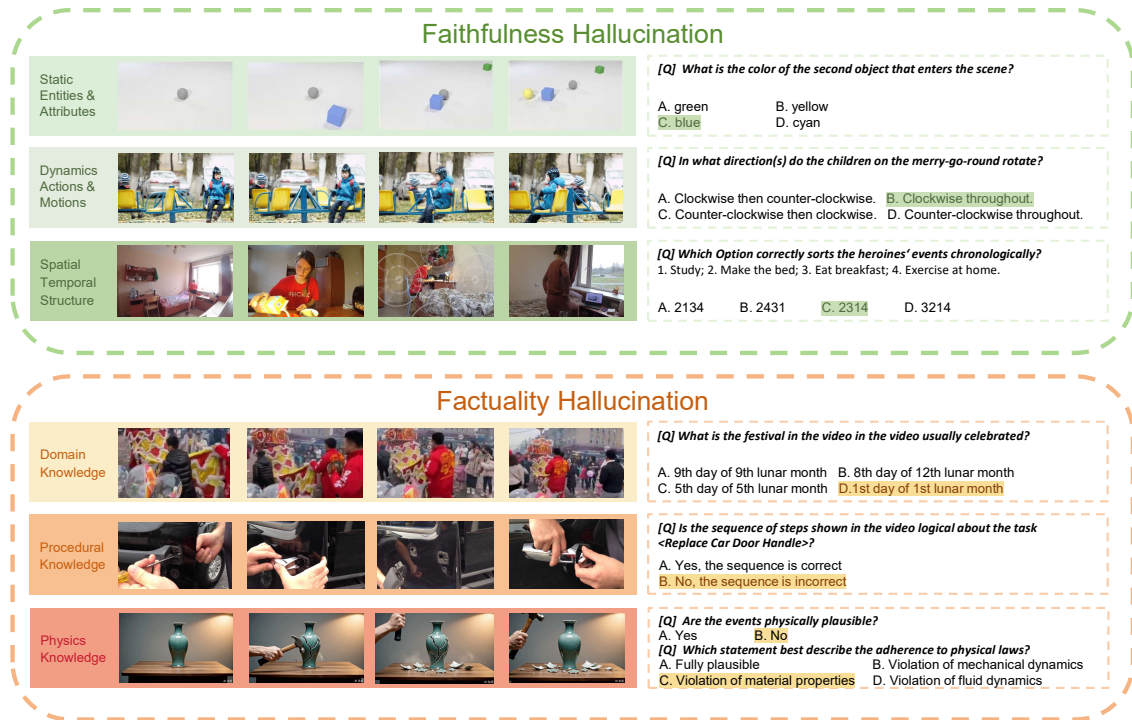


Figure 1: **Examples of Faithfulness and Factuality Hallucinations.** **Top: Faithfulness** items are verified by video evidence, covering Static Entities & Attributes, Dynamic Actions & Motions, and Spatio-Temporal Relations. **Bottom: Factuality** items require consistency with world knowledge and cover Domain Knowledge (Know-WHAT), Procedural Knowledge (Know-HOW), and Physical Knowledge (Know-WHY).

ceptual degradation. Moreover, temporal interventions reveal that many models remain largely insensitive to order disruption in factuality questions, indicating to a gap in temporal grounding beyond clean-scenario accuracy.

Our contributions are three-fold:

1. We introduce INFAC, a diagnostic benchmark comprising 9,800 QA instances spanning real and synthetic videos, with fine-grained taxonomies covering both faithfulness and factuality hallucinations.
2. We propose a four-mode evaluation protocol (Base, Visual Degradation, Evidence Corruption, Temporal Intervention) with paired reliability metrics (RR and TSS) for systematically measuring invariant-label stability and temporal sensitivity.
3. We conduct a systematic evaluation of 14 representative Video-LLMs, revealing their stability under invariant-label perturbations and temporal inertia on order-sensitive items.

2 Related Works

2.1 Hallucination evaluation in Video-LLMs

A growing line of work has benchmarked hallucinations in Video-LLMs from different perspectives (Chen et al., 2025). Existing benchmarks (Zhang et al., 2024; Fu et al., 2025; Li et al., 2025a; Bae et al., 2025; Sung-Bin et al., 2025; Wang et al., 2024b; Kong et al., 2025; Zhou et al., 2025) predominantly focus on faithfulness hallucinations, typically through controlled benchmark constructions that manipulate events, motion, video similarity, or cross-modal consistency to test grounding in the input video. For instance, EventHallusion (Zhang et al., 2024) focuses on event-level dynamics and relations, MHBench (Kong et al., 2025) targets motion-related errors with adversarial triplets, and VidHalluc (Li et al., 2025a) constructs visually distinct yet semantically similar video pairs to expose fragile grounding. Related controlled settings further investigate shortcut behaviors driven by priors or spurious correlations, such as narrative priors in NOAH (Lee et al., 2025), action-scene correlations in UNSCENE (Bae et al., 2025), and cross-modal inconsistency settings in

Table 1: **Comparison of INFAC T with recent hallucination benchmarks for video understanding.** INFAC T covers video-grounded faithfulness and world-knowledge-grounded factuality with fine-grained taxonomies. It also supports controlled evaluation modes for visual degradation, evidence corruption, and temporal intervention within a unified protocol.

Benchmark	# Ques. / # Videos	Faithfulness	Factuality	Source	Visual Degradation		Evidence Corruption		Temporal Intervention	
					Motion Blur	Gaussian Noise	Caption Injection	Adversarial Noise	Shuffle	Reverse
VIDHAL(Choong et al., 2024)	- / 400	✓(5)	✗	Real	✗	✗	✗	✗	✗	✗
EventHallusion(Zhang et al., 2024)	- / 400	✓(3)	✗	Real	✗	✗	✗	✗	✗	✗
VideoHallucer(Wang et al., 2024b)	1,800 / 948	✓(3)	✓(3)	Real	✗	✗	✗	✗	✗	✗
VIDHALLUC(Li et al., 2025a)	9,295 / 5,002	✓(3)	✗	Real	✗	✗	✗	✗	✗	✗
VideoHallu(Li et al., 2025c)	3,233 / 3,233	✓(2)	✓(4)	Synthetic	✗	✗	✗	✗	✗	✗
OURS	9,800 / 9,800	✓(12)	✓(12)	Real & Synthetic	✓	✓	✓	✓	✓	✓

AVHBench (Sung-Bin et al., 2025). Some benchmarks broaden data sources or verification scope: VideoHallu (Li et al., 2025c) introduces synthetic videos, while VidHallucer (Wang et al., 2024b) distinguishes cases by whether the target claim can be verified from the video.

However, as summarized in Table 1, most existing benchmarks concentrate on hallucinations that can be judged against the input video (faithfulness), whereas factuality hallucinations requiring verifiable world knowledge are much less explored.

2.2 Existing Video Understanding Benchmarks

Existing video understanding benchmarks primarily evaluate Video-LLMs in clean settings with capability-oriented scoring. Broad-spectrum suites such as MVBench (Li et al., 2024), VideoMME (Fu et al., 2025), and TOMATO (Shangguan et al., 2025) target general task coverage, while ViMUL-Bench (Shafique et al., 2025) and CityGuesser-style QA (Kulkarni et al., 2024) emphasize knowledge-heavy settings.

However, clean-input capability scores do not directly measure reliability: high accuracy can mask shortcut-based success, in which models rely on language priors, static cues, or dataset biases rather than video-dependent evidence (Liu et al., 2025a; Yan et al., 2025; Zhang et al., 2025; Liu et al., 2024b; Bae et al., 2025).

To reduce shortcut effects, several benchmarks adopt controlled designs such as conflicting videos or temporal multiple-choice setups (Liu et al., 2024b; Cores et al., 2025). Although useful for diagnosing temporal dependence, these protocols are not formulated as hallucination evaluations: they do not separate evidence bases (video vs. verifiable world knowledge) and do not probe reliability under explicit evidence-corruption conditions.

3 INFAC T

In this section, we introduce INFAC T, a fine-grained benchmark designed to evaluate Video-LLMs faithfulness and factuality in both clean and noisy scenarios. We first detail the taxonomy of hallucinations in § 3.1 and the data construction process in § 3.2, describing the aggregation of 9,800 questions from public video-QA datasets, instructional resources, and synthetic collections. To investigate the root causes of model failure, § 3.3 introduces three hallucination induction modes devised to probe reliability under visual degradation, evidence corruption, and temporal intervention. Finally, we introduce two kinds of evaluation metrics to quantify reliability in § 3.4.

3.1 Taxonomy

We categorize hallucinations by the evidentiary basis required for verification, where faithfulness requires alignment with visual content and factuality necessitates consistency with world knowledge.

Faithfulness hallucinations occur when model outputs contradict explicit visual evidence in the video. Following the common object/attribute/relation hierarchy (Liu et al., 2024a) used in static vision-language evaluation, we extend this framework to videos by incorporating temporal dynamics. This yields three hierarchical levels organized by increasing spatiotemporal complexity (Table A5). **Level 1 (Static Entities & Attributes)** corresponds to *Object* and *Attribute* tiers. This level requires local perception to resolve *Entity Recognition*, *Unique Entity Counting*, *Temporal Attributes Recognition*, *Static Attributes Recognition*, and *Scene Text Recognition*. **Level 2 (Dynamic Actions & Motions)** requires dynamic perception to aggregate visual features over time to resolve *Action Recognition*, *Repetitive Action Counting* and *Motion Attributes Recognition*.

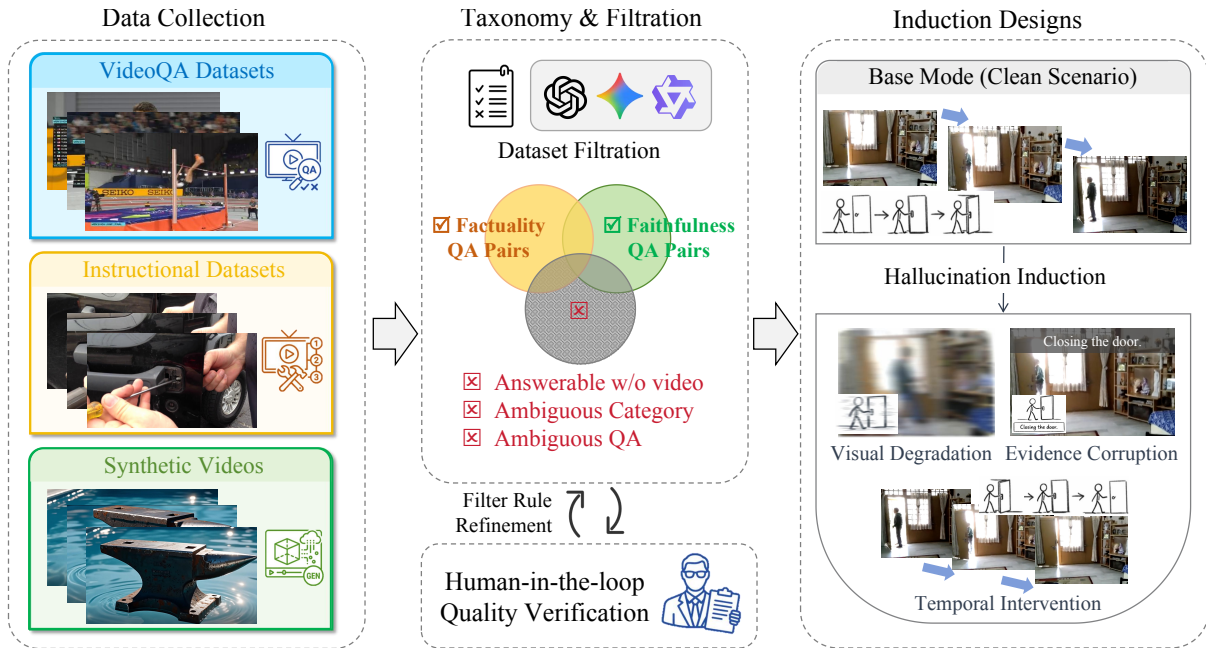


Figure 2: **Overview of the INFACCT construction process.** **Left:** Candidate videos and QA pairs are collected from multiple sources, including video QA datasets, instructional datasets, and synthetic videos. **Middle:** Samples are organized into fine-grained faithfulness and factuality dimensions, and filtered to remove ambiguous or non-video-grounded items, followed by human-in-the-loop quality verification. **Right:** The resulting benchmark supports four evaluation modes: Base, Visual Degradation, Evidence Corruption, and Temporal Intervention.

Level 3 (Spatio-Temporal Relations) corresponds to relation-level reasoning based on global spatio-temporal structure, including *Spatial Relation Recognition*, *Temporal Relation Recognition*, *State Transition Detection*, and *Temporal Localization*.

Factuality hallucinations occur when model outputs contradict world knowledge, requiring information beyond what is present in the video content alone. We structure factuality into three categories based on the knowledge required (Table A6). **Domain Knowledge (Know-WHAT)** evaluates consistency with verifiable world knowledge across diverse domains, including *Cultural Event Recognition*, *Historical Background Identification*, *Geospatial Localization*, and *Entertainment-Related Recognition*. **Procedural Knowledge (Know-HOW)** assesses instructional validity by verifying whether the described sequence of steps respects prerequisite relations and causal dependencies. This category covers *Electronic*, *Mechanical*, *Domestic*, and *Clinical* domains. **Physical Knowledge (Know-WHY)** probes the adherence of a model to fundamental physical laws, including *Mechanical*, *Fluid Mechanics*, *Material Property*, and *Spatio-Temporal continuity*. To ensure rigor, We restrict factuality to objectively verifiable knowledge and explicitly exclude subjective or am-

biguous claims. Figure 3 summarizes the distribution of samples across the fine-grained taxonomy for both faithfulness and factuality.

3.2 Data Construction

Faithfulness Data. To construct a rigorous benchmark for hallucination evaluation, we curate high-quality samples from established video understanding benchmarks, including MVBench (Li et al., 2024), Video-MME (Fu et al., 2025), and TOMATO (Shangguan et al., 2025), which provide diverse questions suitable for video evidence verification. To guarantee data quality, we employ a two-stage alignment process. First, we manually map the categories of source benchmarks to our proposed taxonomy. Subsequently, we implement a "taxonomy-aligned filtration" pipeline that utilizes LLM-assisted consensus labeling. This ensures strict adherence to the spatio-temporal complexity levels defined in Appendix Table A5, thereby refining the quality of constructed benchmark.

Factuality Data. Factuality samples are constructed to cover three verifiable knowledge requirements: Domain, Procedural, and Physical Knowledge. For **Domain Knowledge**, candidates are drawn from knowledge-intensive video QA resources such as CityGuessr68k (Kulkarni et al.,

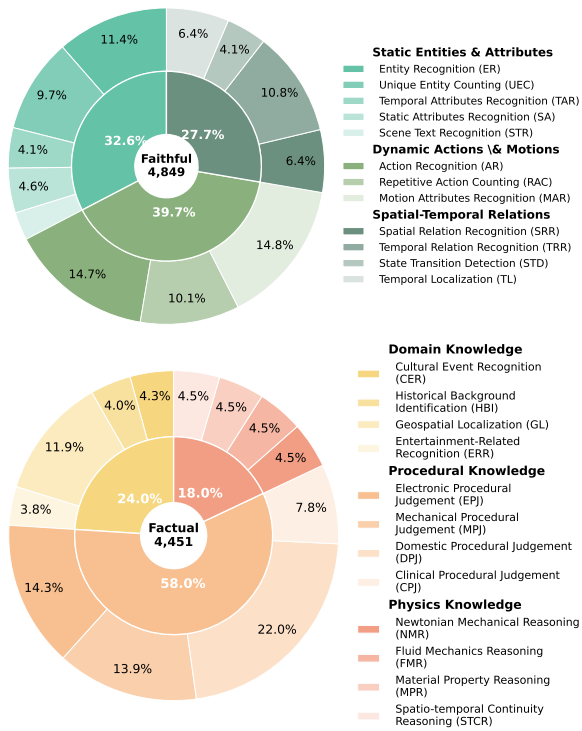


Figure 3: **Dataset composition of INFAC.** Distribution over the fine-grained taxonomy for faithfulness (top) and factuality (bottom).

2024) and ViMULBench (Shafique et al., 2025) to ensure comprehensive coverage across diverse domains such as culture, history, geography, and entertainment. For **Procedural Knowledge**, instructional videos are sourced from COIN (Tang et al., 2019) and MedVidQA (Gupta et al., 2023), where step annotations provide temporal boundaries for each procedure segment. To create counterfactuals, we randomly shuffle these segments to disrupt the temporal order required by prerequisite relations and causal dependencies. The resulting videos support *binary* procedural judgement (correct vs. incorrect). For **Physical Knowledge**, physically implausible videos are synthesized using text-to-video generation models (e.g., Sora (OpenAI, 2024), Wan2.5 (Wan AI, 2025), and Gemini Veo 3 (Google DeepMind, 2025)). These samples violate physical laws, such as gravity-defying motion, serving as a test for physical reasoning capabilities. Unlike training-time synthetic data pipelines for vision–language models (Liu et al., 2024c), the synthetic videos here serve solely as evaluation probes for factuality hallucinations.

Dataset Filtration and Quality Review. An LLM-ensemble filter is first applied to remove samples with linguistic ambiguity and those solvable

without video evidence (e.g., answerable from text-only priors). Since source annotations and edge cases in taxonomy mapping may still introduce noise, a **human-in-the-loop** verification stage is further conducted. Rather than discarding individual problematic samples only, annotators also identify recurring error patterns and trace them back to their causes (e.g., a question template or a mapping rule). This enables iterative updates to the upstream prompts and filtering logic and batch correction of affected samples. The process is repeated until the benchmark satisfies the quality criteria.

3.3 Hallucination Induction Designs

INFAC establishes four distinct evaluation modes to probe evidence-consistent behavior under controlled conditions. Mode I (Base) establishes the baseline performance on clean data. Modes II & III (Robustness) apply label-preserving transformations to test whether models maintain correct answers under visual degradation or conflicting evidence. Mode IV (Sensitivity) applies label-changing interventions to test sensitivity to disrupted temporal structure.

Mode I: Base (clean scenario). The Base mode evaluates the model on the original, unaltered video and question pairs. It serves as the reference standard for establishing upper-bound performance. Crucially, to isolate the impact of interventions from intrinsic item difficulty, we employ a *paired comparison protocol*: behavior under all induced modes is assessed relative to the model’s Base outcome on the same specific item.

Mode II: Visual Degradation. This mode probes hallucinations induced by perceptual uncertainty. We introduce three types of visual perturbations: (1) **Gaussian noise** and (2) **Motion blur**, both applied uniformly at the frame level, and (3) **Video Compression**, which simulates platform-specific re-encoding artifacts by reducing the bitrate to a lossy threshold. While these operators reduce low-level perceptual clarity, they preserve the high-level semantic content required to answer the question. Consequently, a reliable model should maintain the correct answer despite the visual noise, rather than introducing unsupported details when the evidence becomes harder to perceive.

Mode III: Evidence Corruption. Since VideoLLMs often integrate auxiliary textual signals (e.g, subtitles) with visual data, reliability hinges

on the ability to prioritize authentic visual evidence over misleading external cues. We employ three operators to simulate untrusted conditioning while keeping the ground-truth label invariant. **Caption injection** introduces a cross-modal conflict by overlaying subtitles on randomly sampled video segments. The injected subtitles mix (i) content-irrelevant sentences and (ii) LLM-generated misleading statements conditioned on the question and the ground-truth answer option, constructed to form a plausible but incorrect textual cue that conflicts with the video evidence required for correctness (e.g., a video showing *opening a door* is paired with the subtitle *closing the door*). **Subtitle corruption** injects noisy ASR-like subtitles with subtitle–video desynchronization, simulating the unreliable OCR/ASR outputs that models may encounter when processing captioned or subtitled content in the wild. **Adversarial noise** targets the visual encoder by applying a transfer-based black-box perturbation generated with MI-FGSM (Dong et al., 2018). To ensure the attack generalizes across different architectures, perturbations are produced using a proxy ensemble of visual encoders (InternVL3-8B (Zhu et al., 2025), Qwen3VL-8B (Bai et al., 2025a), and VideoMAE (Tong et al., 2022)) to reduce reliance on a single proxy model. Since the ground-truth label is intended to remain unchanged, the desired behavior is *stability* against misleading conditioning.

Mode IV: Temporal Intervention. This mode probes temporal evidence consistency. Specifically, whether the model’s correctness genuinely stems from understanding event order and state transitions. Unlike the previous modes, here the temporal structure is intentionally destroyed, rendering the original ground truth invalid. To evaluate this, we curate an *order-sensitive* subset from Action Dynamics, Spatio-Temporal Structure, and Procedural Knowledge, where the chronological sequence is strictly required for correctness. We apply temporal intervention to these videos via frame-level shuffling or reversal. If a model retains the Base prediction after intervention, it indicates temporal insensitivity and reliance on order-invariant cues.

3.4 Evaluation Metrics

Accuracy is reported under the Base (clean-scenario) setting to measure performance on faithfulness and factuality. To quantify reliability under induced conditions, we additionally report metrics

tailored to invariant-label settings and temporal interventions. Modes II–III (Visual Degradation and Evidence Corruption) are evaluated by Resist Rate (RR), while Mode IV (Temporal Intervention) is evaluated by Temporal Sensitivity Score (TSS).

Resist Rate (RR). For visual degradations (\mathcal{T}_{deg}) and evidence corruptions (\mathcal{T}_{cor}), the ground-truth label is intended to remain unchanged. We measure reliability by whether a model preserves its correct Base prediction under each operator. For an operator $p \in \mathcal{T}_{\text{pert}} = \mathcal{T}_{\text{deg}} \cup \mathcal{T}_{\text{cor}}$, we define:

$$RR_p = \frac{\sum_i \mathbb{I}(f(V_i, q_i) = y_i^{\text{gt}}) \cdot \mathbb{I}(f(p(V_i), q_i) = y_i^{\text{gt}})}{\sum_i \mathbb{I}(f(V_i, q_i) = y_i^{\text{gt}})},$$

where (V_i, q_i) denotes the i -th video-question pair, y_i^{gt} is the corresponding ground-truth answer, $f(\cdot, \cdot)$ denotes a Video-LLM, and $\mathbb{I}(\cdot)$ is the indicator function. We report operator-wise scores (e.g., RR_{gau} , RR_{mb} , RR_{adv} , RR_{cap}), and compute RR_{deg} and RR_{cor} as the mean over degradation and corruption operators, respectively.

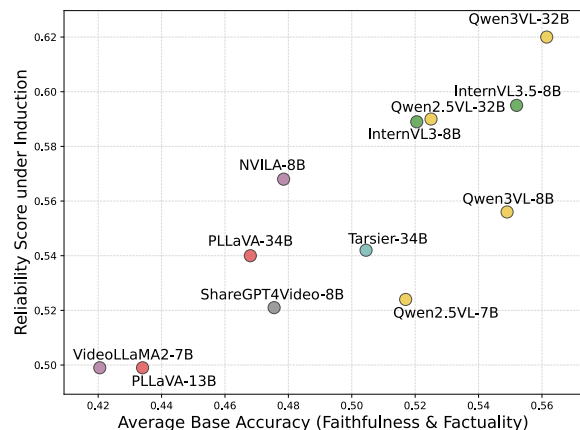


Figure 4: **Base accuracy vs. average reliability score under inductions.** Base accuracy is measured in Mode I and averaged over faithfulness and factuality. The average reliability score under induction aggregates RR over Modes II–III and TSS over Mode IV.

Temporal Sensitivity Score (TSS). Temporal interventions (shuffling or reversal) test whether model decisions are genuinely grounded in temporal structure. TSS is computed on the **order-sensitive subset** $\mathcal{S}_{\text{order}}$, where temporal order is essential for correctness. Unlike the robustness tests used for RR, these interventions effectively invalidate the original ground-truth label. Consequently, a temporally grounded model should diverge from its initial decision after intervention. We define

Table 2: **Faithfulness results on INFACT.** Text-only: question-only accuracy. Base: Mode I. RR: Modes II–III. TSS: Mode IV (order-sensitive subset). Models are grouped by availability and sorted by Avg Score.

Model	Text-only	Base	Evidence Corruption				Visual Degradation				Temporal Intervention			Avg Score
			RR_{adv}	RR_{cap}	RR_{sub}	RR_{cor}	RR_{cmp}	RR_{gau}	RR_{mb}	RR_{deg}	TSS_{shu}	TSS_{rev}	TSS	
PLLaVA-13B (Xu et al., 2024)	0.285	0.458	0.782	0.678	0.922	0.794	0.777	0.813	0.699	0.763	0.010	0.014	0.012	0.523
VideoLLaMA2-7B(Cheng et al., 2024)	0.261	0.434	0.751	0.564	0.910	0.742	0.818	0.725	0.749	0.764	0.113	0.091	0.102	0.536
ShareGPT4Video-8B (Chen et al., 2024)	0.292	0.462	0.719	0.688	0.894	0.767	0.840	0.762	0.754	0.785	0.101	0.094	0.098	0.550
PLLaVA-34B (Xu et al., 2024)	0.287	0.505	0.757	0.618	0.920	0.765	0.772	0.874	0.848	0.831	0.070	0.073	0.072	0.556
Tarsier-34B (Wang et al., 2024a)	<u>0.302</u>	0.568	0.785	0.644	0.960	0.796	0.759	0.846	0.612	<u>0.739</u>	0.158	0.208	0.183	0.573
Qwen2.5VL-7B(Bai et al., 2025b)	0.278	0.531	0.630	0.659	0.973	0.754	0.804	0.727	0.640	0.724	0.230	0.258	0.244	0.574
NVILA-8B (Liu et al., 2024d)	0.293	0.533	0.804	0.669	0.963	0.812	0.797	0.853	0.809	0.820	0.122	0.150	0.136	0.589
Qwen3VL-8B(Bai et al., 2025a)	0.295	0.557	0.676	0.715	0.941	0.777	0.829	0.785	0.744	0.786	0.224	0.189	0.207	0.590
Qwen2.5VL-32B(Bai et al., 2025b)	0.286	0.538	0.817	0.724	0.949	0.830	0.790	0.824	0.773	0.796	0.198	0.204	0.201	0.609
InternVL3-8B (Zhu et al., 2025)	0.301	0.576	<u>0.777</u>	0.709	0.979	0.822	0.828	0.847	0.817	<u>0.831</u>	0.169	0.212	0.191	0.615
InternVL3.5-8B (Wang et al., 2025)	0.294	<u>0.606</u>	0.796	0.721	<u>0.978</u>	<u>0.832</u>	0.841	0.792	<u>0.850</u>	0.828	0.171	0.241	0.206	0.622
Qwen3VL-32B (Bai et al., 2025a)	0.287	0.602	0.809	0.718	0.975	0.834	0.802	0.836	0.815	0.818	0.248	0.289	0.269	0.640
GPT-5.1 (OpenAI, 2025)	0.305	<u>0.687</u>	<u>0.856</u>	<u>0.823</u>	0.958	<u>0.879</u>	<u>0.940</u>	0.801	<u>0.854</u>	<u>0.865</u>	<u>0.355</u>	0.281	<u>0.318</u>	<u>0.687</u>
Gemini3-flash(Gemini, 2025)	0.295	0.784	0.886	0.914	0.962	0.921	0.942	<u>0.859</u>	0.892	0.898	0.581	0.492	0.537	0.785

Table 3: **Factuality results on INFACT.** Same evaluation protocol and metrics as Table 2.

Model	Text-only	Base	Evidence Corruption				Visual Degradation				Temporal Intervention			Avg Score
			RR_{adv}	RR_{cap}	RR_{sub}	RR_{cor}	RR_{cmp}	RR_{gau}	RR_{mb}	RR_{deg}	TSS_{shu}	TSS_{rev}	TSS	
VideoLLaMA2-7B(Cheng et al., 2024)	0.272	0.407	0.676	0.590	0.917	0.728	0.818	0.723	0.772	0.771	0.000	0.000	0.000	0.500
ShareGPT4Video-8B (Chen et al., 2024)	0.294	0.489	0.715	0.519	0.883	0.706	0.840	0.816	0.798	0.818	0.000	0.000	0.000	0.508
PLLaVA-13B (Xu et al., 2024)	0.279	0.410	0.788	0.585	0.919	0.764	0.777	0.947	0.705	0.810	0.000	0.000	0.000	0.525
NVILA-8B (Liu et al., 2024d)	0.291	0.424	0.818	0.541	0.961	0.773	0.797	0.888	0.759	0.815	0.000	0.000	0.000	0.529
PLLaVA-34B (Xu et al., 2024)	0.265	0.431	0.782	0.603	0.920	0.768	0.772	0.897	0.812	0.827	0.000	0.000	0.000	0.532
InternVL3-8B (Zhu et al., 2025)	0.295	0.465	0.738	0.612	0.951	0.767	0.828	0.914	0.813	<u>0.852</u>	0.000	0.000	0.000	0.540
InternVL3.5-8B(Wang et al., 2025)	0.298	0.498	0.775	0.619	0.955	0.783	0.841	0.921	0.851	0.871	0.000	0.000	0.000	0.551
Tarsier-34B (Wang et al., 2024a)	0.289	0.441	0.909	0.625	0.960	0.831	0.759	0.930	0.820	0.836	0.000	0.000	0.000	0.556
Qwen2.5VL-32B(Bai et al., 2025b)	0.296	0.512	0.804	0.624	0.930	<u>0.786</u>	0.790	<u>0.889</u>	0.764	0.814	0.105	0.101	0.103	0.568
Qwen2.5VL-7B(Bai et al., 2025b)	0.281	0.503	0.670	0.575	<u>0.968</u>	0.738	0.804	0.784	0.774	0.787	0.212	0.262	0.237	0.587
Qwen3VL-8B (Bai et al., 2025a)	0.296	0.541	0.707	0.606	0.929	0.747	0.829	0.793	0.768	0.797	<u>0.339</u>	<u>0.334</u>	<u>0.337</u>	0.627
Qwen3VL-32B (Bai et al., 2025a)	<u>0.305</u>	0.521	0.809	0.641	<u>0.963</u>	0.804	0.802	0.904	0.821	0.842	0.231	0.254	0.243	0.630
GPT-5.1 (OpenAI, 2025)	0.308	0.678	0.901	0.814	<u>0.963</u>	0.893	0.940	0.911	0.876	0.909	0.391	0.414	0.402	0.735
Gemini3-flash(Gemini, 2025)	<u>0.305</u>	0.752	0.917	0.839	0.969	0.908	0.942	0.902	0.881	0.908	0.741	0.696	0.719	0.845

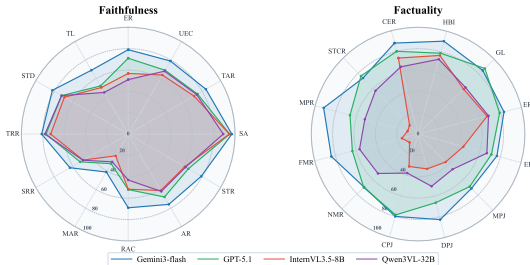


Figure 5: **Comparison of four representative models on fine-grained evaluation dimensions.** The left radar plot shows performance on faithfulness dimensions, while the right radar plot shows performance on factuality dimensions. Each axis corresponds to a fine-grained category in the INFACT taxonomy, and each curve represents one representative model. Higher values indicate better performance on the corresponding dimension.

TSS as the rate at which model ceases to predict the original ground-truth label y_{gt} after intervention. For an intervention operator $p \in \mathcal{T}_{iv}$ (e.g., shuffling or reversal), we define:

$$TSS_p = \frac{\sum_{i \in S_{order}} \mathbb{I}(f(V_i, q_i) = y_{gt}) \cdot \mathbb{I}(f(p(V_i), q_i) \neq y_{gt})}{\sum_{V_i \in S_{order}} \mathbb{I}(f(V_i, q_i) = y_{gt})}$$

We report TSS_{shu} and TSS_{rev} , and TSS is their mean. A low TSS indicates *temporal inertia*: the model tends to retain the Base decision even when the supporting temporal evidence is destroyed, suggesting a reliance on static priors rather than temporal logic.

4 Experiments

4.1 Setups

Fourteen models are evaluated on INFACT, including two proprietary systems (GPT-5.1 (OpenAI, 2025), Gemini3-flash (Gemini, 2025)) and twelve open-source baselines: VideoLLaMA2-7B (Cheng et al., 2024), PLLaVA-13B/34B (Xu et al., 2024), ShareGPT4Video-8B (Chen et al., 2024), Qwen2.5-VL-7B/32B (Bai et al., 2025b), Qwen3-VL-8B/32B (Bai et al., 2025a), Tarsier-34B (Wang et al., 2024a), NVILA-8B (Liu et al., 2024d), InternVL3-8B (Zhu et al., 2025), and InternVL3.5-8B (Wang et al., 2025). All models are evaluated zero-shot with default prompts using 16 uniformly sampled frames per video (Shang-guan et al., 2025; Rawal et al., 2025).




Figure 6: **Effect of the number of sampled frames.** **Left:** Base accuracy under $\{2, 8, 16, 24, 32\}$ uniformly sampled frames. **Right:** induced metrics (RR_{deg} , RR_{cor} , TSS) under $\{8, 16, 24, 32\}$ frames.

4.2 Results and Analysis

Base accuracy vs. reliability under induced modes.

Figure 4 reveals a strong association between Base accuracy (Mode I) and induced-mode reliability (Modes II–IV), averaged over faithfulness and factuality (Pearson $r=0.978$; Spearman $\rho=0.969$). However, rankings are not fully preserved: models with similar Base accuracy can diverge under controlled perturbations. Tables 2–3 suggest that these gaps are driven by differences in stability (RR) and temporal sensitivity (TSS), so clean accuracy alone is insufficient to characterize evidence-consistent behavior.

Stability under Visual Degradation & Evidence Corruption Inductions (RR).

Tables 2–3 show that evidence corruption degrades RR more than visual degradation, and caption injection is typically more damaging than transfer-based adversarial noise. For instance, VideoLLaMA2-7B drops from $RR_{adv}=0.751$ to $RR_{cap}=0.564$ on faithfulness, and Tarsier-34B drops from $RR_{adv}=0.909$ to $RR_{cap}=0.625$ on factuality. This pattern suggests that stability is more fragile when misleading auxiliary cues are introduced than when the visual signal is merely degraded.

Sensitivity under Temporal Intervention (TSS).

Tables 2–3 report TSS on the order-sensitive subset, where frame shuffling or reversal disrupts the temporal structure required for correctness. Several open-source baselines exhibit *temporal inertia*, most visibly on factuality where multiple models yield $TSS=0$, with predictions unchanged after shuffling or reversal on Base-correct order-sensitive items. In contrast, Qwen3VL-8B and Gemini3-flash attain 0.337 and 0.719, indicating that temporal intervention separates models by their reliance on temporal evidence. Among open-source baselines, Qwen2.5/3-VL achieves comparatively higher TSS (Tables 2–3), highlighting the potential

role of more explicit time-aligned spatiotemporal positional encoding (e.g., mRoPE-style designs).

Performance across Fine-grained Dimensions.

For finer diagnosis, we further analyze four representative models across fine-grained faithfulness and factuality dimensions (Figure 5). Across these models, factuality dimensions tend to lag behind faithfulness dimensions, especially in procedural judgement and physical reasoning. On the faithfulness side, the remaining weaknesses concentrate on motion- and structure-centric dimensions (e.g., temporal localization and motion attributes), suggesting that fine-grained temporal aggregation remains challenging under our evaluation setting.

Effect of Frame Sampling Rate.

We conduct a frame-ablation study on 500 videos randomly sampled from INFACT (Figure 6). Base accuracy is evaluated with $\{2, 8, 16, 24, 32\}$ uniformly sampled frames and quickly saturates, where 16 frames already capture most of the gain (Figure 6, left). Under induced modes with $\{8, 16, 24, 32\}$ frames, RR_{deg} and RR_{cor} vary modestly and TSS shows no consistent upward trend (Figure 6, right), suggesting that increasing the number of sampled frames alone does not consistently improve reliability under induced modes within this range, particularly for temporal sensitivity.

5 Conclusion

We present INFACT, a fine-grained benchmark for evaluating faithfulness and factuality hallucinations in Video-LLMs under both Base mode and controlled evidence perturbations, including invariant-label inductions and temporal interventions on order-sensitive items. Across 14 models, factuality remains challenging for procedural judgement and physical reasoning, whereas faithfulness errors concentrate in motion- and structure-centric skills. Models also exhibit limited stability under

invariant-label perturbations and pronounced temporal inertia under temporal intervention.

Limitations

The induced modes use controlled operators as proxies for deployment perturbations and do not cover the full space of corruptions or unreliable conditioning. Temporal Intervention is tested on an order-sensitive subset using shuffling and reversal; it probes temporal reliance but does not localize the cues driving model decisions.

Acknowledgements

This work is partially supported by the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant XDB0680202, the Key Research and Development Program of Xinjiang Uyghur Autonomous Region under Grant 2024B03026, and Beijing Nova Program under Grant 20230484368.

References

- Kyungho Bae, Jinhyung Kim, Sihaeng Lee, Soonyoung Lee, Gunhee Lee, and Jinwoo Choi. 2025. Mash-vm: Mitigating action-scene hallucination in video-llms through disentangled spatial-temporal representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13744–13753.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025a. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Lin Chen, Xilin Wei, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Bin Lin, Zhenyu Tang, Li Yuan, Yu Qiao, Dahua Lin, Feng Zhao, and Jiaqi Wang. 2024. [Sharegpt4video: Improving video understanding and generation with better captions](#). *Preprint*, arXiv:2406.04325.
- Zhiyuan Chen, Yuecong Min, Jie Zhang, Bei Yan, Jiahao Wang, Xiaozhen Wang, and Shiguang Shan. 2025. A survey of multimodal hallucination evaluation and detection. *International Journal of Computer Vision*.
- Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, and Lidong Bing. 2024. [Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms](#). *arXiv preprint arXiv:2406.07476*.
- Wey Yeh Choong, Yangyang Guo, and Mohan Kankanhalli. 2024. [Vidhal: Benchmarking temporal hallucinations in vision llms](#). *arXiv preprint arXiv:2411.16771*.
- Daniel Cores, Michael Dorkenwald, Manuel Mucientes, Cees G. M. Snoek, and Yuki M. Asano. 2025. [Lost in time: A new temporal benchmark for videollms](#). *Preprint*, arXiv:2410.07752.
- Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. 2018. [Boosting adversarial attacks with momentum](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9185–9193.
- Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, Peixian Chen, Yanwei Li, Shaohui Lin, Sirui Zhao, Ke Li, Tong Xu, Xiawu Zheng, Enhong Chen, Caifeng Shan, and 2 others. 2025. [Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24108–24118.
- Gemini. 2025. [Gemini-3-flash](#). <https://blog.google/products/gemini/gemini-3-flash/>.
- Google DeepMind. 2025. [Veo](#). <https://deepmind.google/models/veo/>. Accessed: 2026-03-10.
- Deepak Gupta, Kush Attal, and Dina Demner-Fushman. 2023. [A dataset for medical instructional video classification and question answering](#). *Scientific Data*, 10(1):158.
- Ming Kong, Xianzhou Zeng, Luyuan Chen, Yadong Li, Bo Yan, and Qiang Zhu. 2025. [Mhbench: Demystifying motion hallucination in videollms](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(4):4401–4409.
- Parth Parag Kulkarni, Gaurav Kumar Nayak, and Mubarak Shah. 2024. [Cityguessr: City-level video geo-localization on a global scale](#). In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LXIII*, page 293–311, Berlin, Heidelberg. Springer-Verlag.
- Kyuhoo Lee, Euntae Kim, Jinwoo Choi, and Buru Chang. 2025. [Noah: Benchmarking narrative prior driven hallucination and omission in video large language models](#). *Preprint*, arXiv:2511.06475.

- Chaoyu Li, Eun Woo Im, and Pooyan Fazli. 2025a. **Vid-halluc: Evaluating temporal hallucinations in multi-modal large language models for video understanding**. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13723–13733.
- Chenglin Li, Feng Han, Yikun Wang, Ruilin Li, Shuai Dong, Haowen Hou, Haitao Li, Qianglong Chen, Feng Tao, Jingqi Tong, Yin Zhang, and Jiaqi Wang. 2025b. **Videopro: Adaptive program reasoning for long video understanding**. *Preprint*, arXiv:2509.17743.
- Kunchang Li, Yali Wang, Yanan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Lou, Limin Wang, and Yu Qiao. 2024. **Mvbench: A comprehensive multi-modal video understanding benchmark**. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22195–22206.
- Zongxia Li, Xiyang Wu, Guangyao Shi, Yubin Qin, Hongyang Du, Fuxiao Liu, Tianyi Zhou, Dinesh Manocha, and Jordan Lee Boyd-Graber. 2025c. **Videohallu: Evaluating and mitigating multi-modal hallucinations on synthetic video understanding**. *arXiv preprint arXiv:2505.01481*.
- Honglin Lin, Zheng Liu, Yun Zhu, Chonghan Qin, Juekai Lin, Xiaoran Shang, Conghui He, Wentao Zhang, and Lijun Wu. 2026. **Mmfinereason: Closing the multimodal reasoning gap via open data-centric methods**. *Preprint*, arXiv:2601.21821.
- Hanchao Liu, Wenyan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024a. **A survey on hallucination in large vision-language models**. *Preprint*, arXiv:2402.00253.
- Jiazhen Liu, Yuhang Fu, Ruobing Xie, Runqian Xie, Xingwu Sun, Fengzong Lian, Zhanhui Kang, and Xirong Li. 2025a. **Phd: A chatgpt-prompted visual hallucination evaluation dataset**. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19857–19866.
- Shiyu Liu, Xinyi Wen, Zhibin Lan, Ante Wang, and Jinsong Su. 2026. **Countering the over-reliance trap: Mitigating object hallucination for lvlms via a self-validation framework**. *Preprint*, arXiv:2601.22451.
- Yuanxin Liu, Shicheng Li, Yi Liu, Yuxiang Wang, Shuhuai Ren, Lei Li, Sishuo Chen, Xu Sun, and Lu Hou. 2024b. **Tempcompass: Do video llms really understand videos?** In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 8731–8772. Association for Computational Linguistics.
- Zheng Liu, Hao Liang, Bozhou Li, Wentao Xiong, Chong Chen, Conghui He, Wentao Zhang, and Bin Cui. 2024c. **Synthvlm: Towards high-quality and efficient synthesis of image-caption datasets for vision-language models**. *Preprint*, arXiv:2407.20756.
- Zheng Liu, Mengjie Liu, Jingzhou Chen, Jingwei Xu, Bin Cui, Conghui He, and Wentao Zhang. 2025b. **Fusion: Fully integration of vision-language representations for deep cross-modal understanding**. *Preprint*, arXiv:2504.09925.
- Zhijian Liu, Ligeng Zhu, Baifeng Shi, Zhuoyang Zhang, Yuming Lou, Shang Yang, Haocheng Xi, Shiyi Cao, Yuxian Gu, Dacheng Li, Xiuyu Li, Yunhao Fang, Yukang Chen, Cheng-Yu Hsieh, De-An Huang, An-Chieh Cheng, Vishwesh Nath, Jinyi Hu, Sifei Liu, and 8 others. 2024d. **Nvila: Efficient frontier visual language models**. *Preprint*, arXiv:2412.04468.
- Guangtao Lyu, Xinyi Cheng, Chenghao Xu, Qi Liu, Muli Yang, Fen Fang, Huilin Chen, Jiexi Yan, Xu Yang, and Cheng Deng. 2025. **Revealing perception and generation dynamics in lvlms: Mitigating hallucinations via validated dominance correction**. *Preprint*, arXiv:2512.18813.
- Guangtao Lyu, Qi Liu, Chenghao Xu, Jiexi Yan, Muli Yang, Xueting Li, Fen Fang, and Cheng Deng. 2026. **Revealing and enhancing core visual regions: Harnessing internal attention dynamics for hallucination mitigation in lvlms**. *Preprint*, arXiv:2602.15556.
- OpenAI. 2024. **Sora: Creating video from text**. <https://openai.com/index/sora/>. Accessed: 2026-03-10.
- OpenAI. 2025. **Gpt-5.1: A smarter, more conversational chatgpt**. <https://openai.com/index/gpt-5-1/>.
- Ruchit Rawal, Reza Shirkavand, Heng Huang, Gowthami Somepalli, and Tom Goldstein. 2025. **Argus: Hallucination and omission evaluation in video-llms**. *Preprint*, arXiv:2506.07371.
- Bhuiyan Sanjid Shafique, Ashmal Vayani, Muhammad Maaz, Hanoona Abdul Rasheed, Dinura Disanayake, Mohammed Irfan Kurpath, Yahya Hmaiti, Go Inoue, Jean Lahoud, Md. Safirur Rashid, Shadid Intisar Quasem, Maheen Fatima, Franco Vidal, Mykola Maslych, Ketan Pravin More, Sanoojan Baliah, Hasindri Watawana, Yuhao Li, Fabian Farestam, and 10 others. 2025. **A culturally-diverse multilingual multimodal video benchmark and model**. *Preprint*, arXiv:2506.07032.
- Ziyao Shangguan, Chuhan Li, Yuxuan Ding, Yanan Zheng, Yilun Zhao, Tesca Fitzgerald, and Arman Cohan. 2025. **Tomato: Assessing visual temporal reasoning capabilities in multimodal foundation models**. In *International Conference on Representation Learning*, volume 2025, pages 7593–7734.
- Kim Sung-Bin, Oh Hyun-Bin, Lee Jung-Mok, Arda Senocak, Joon Son Chung, and Tae-Hyun Oh. 2025. **Avhbench: A cross-modal hallucination benchmark for audio-visual large language models**. In *International Conference on Representation Learning*, volume 2025, pages 24244–24271.
- Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou.

2019. Coin: A large-scale dataset for comprehensive instructional video analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. 2022. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *Advances in Neural Information Processing Systems*.
- Wan AI. 2025. Wan 2.5: Native audio like veo3 + 1080p video generation. <https://wan25.ai/>. Accessed: 2026-03-10.
- Jiawei Wang, Liping Yuan, Yuchen Zhang, and Hao-miao Sun. 2024a. Tarsier: Recipes for training and evaluating large video description models. *Preprint*, arXiv:2407.00634.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, and 1 others. 2025. Internv13. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*.
- Yuxuan Wang, Yueqian Wang, Dongyan Zhao, Cihang Xie, and Zilong Zheng. 2024b. Videohalluciner: Evaluating intrinsic and extrinsic hallucinations in large video-language models. *arXiv preprint arXiv:2406.16338*.
- Lin Xu, Yilin Zhao, Daquan Zhou, Zhijie Lin, See Kiong Ng, and Jiashi Feng. 2024. Pllava: Parameter-free llava extension from images to videos for video dense captioning. *arXiv preprint arXiv:2404.16994*.
- Bei Yan, Zhiyuan Chen, Yuecong Min, Jie Zhang, Jiahao Wang, Xiaozhen Wang, and Shiguang Shan. 2025. Shale: A scalable benchmark for fine-grained hallucination evaluation in vlms. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 13442–13449.
- Jiacheng Zhang, Yang Jiao, Shaoxiang Chen, Na Zhao, Zhiyu Tan, Hao Li, Xingjun Ma, and Jingjing Chen. 2024. Eventhallusion: Diagnosing event hallucinations in video llms. *arXiv preprint arXiv:2409.16597*.
- Jie Zhang, Zhongqi Wang, Mengqi Lei, Zheng Yuan, Bei Yan, Shiguang Shan, and Xilin Chen. 2025. Dysca: A dynamic and scalable benchmark for evaluating perception ability of vlms. In *International Conference on Learning Representations (ICLR)*.
- Qianchi Zhang, Hainan Zhang, Liang Pang, Hongwei Zheng, and Zhiming Zheng. 2026. Stable-rag: Mitigating retrieval-permutation-induced hallucinations in retrieval-augmented generation. *Preprint*, arXiv:2601.02993.
- Wenrui Zhou, Mohamed Hendy, Shu Yang, Qingsong Yang, Zikun Guo, Yuyu Luo, Lijie Hu, and Di Wang. 2025. Flattery in motion: Benchmarking and analyzing sycophancy in video-llms. *Preprint*, arXiv:2506.07180.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, and 1 others. 2025. Internv13: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*.

Open-Source Video-LLMs	HF Checkpoint
VideoLLaMA2-7B	DAMO-NLP-SG/VideoLLaMA2-7B
Qwen2.5VL-7B	Qwen/Qwen2.5-VL-7B-Instruct
Qwen3VL-8B	Qwen/Qwen3-VL-8B-Instruct
ShareGPT4Video-8B	Lin-Chen/sharegpt4video-8b
NVILA-8B	Efficient-Large-Model/NVILA-8B
InternVL3-8B	OpenGVLab/InternVL3-8B
InternVL3.5-8B	OpenGVLab/InternVL3_5-8B
PLLaVA-13B	ermu2001/pllava-13b
Qwen2.5VL-32B	Qwen/Qwen2.5-VL-32B-Instruct
Qwen3VL-32B	Qwen/Qwen3-VL-32B-Instruct
PLLaVA-34B	ermu2001/pllava-34b
Tarsier-34B	omni-research/Tarsier-34b

Table A1: Simplified configurations for open-source multimodal foundation models used in the evaluation.

A Appendix

A.1 Detailed Taxonomy and Examples

In this appendix, we provide the comprehensive definitions and detailed examples for our proposed hallucination taxonomy. Table A5 details the three-level hierarchy for **Faithfulness Hallucination**. Table A6 presents the categorization for **Factuality Hallucination**.

A.2 TSS: Order-Sensitive vs. Non-Order-Sensitive

To validate that TSS provides meaningful separation, Table A2 compares TSS computed on the order-sensitive subset ($\mathcal{S}_{\text{order}}$) and non-order-sensitive items ($\mathcal{S}_{\text{nonorder}}$) for models with non-trivial temporal responsiveness. For all models, TSS on $\mathcal{S}_{\text{order}}$ substantially exceeds TSS on $\mathcal{S}_{\text{nonorder}}$, confirming that temporal intervention is more impactful on genuinely time-dependent items and that the order-sensitivity annotation provides a useful signal.

Model	TSS ($\mathcal{S}_{\text{order}}$)	TSS ($\mathcal{S}_{\text{nonorder}}$)
Qwen2.5-VL-7B	20.5%	4.5%
Qwen2.5-VL-32B	19.5%	6.0%
Qwen3-VL-8B	28.0%	9.0%
Qwen3-VL-32B	25.0%	8.0%
InternVL3-8B	14.0%	6.0%
InternVL3.5-8B	22.5%	5.5%
Tarsier-34B	21.5%	7.5%

Table A2: **TSS on order-sensitive vs. non-order-sensitive items.** Models with non-trivial temporal responsiveness show meaningful separation between subsets.

A.3 Human Validity Audit for Invariant-Label Operators

A key assumption behind Modes II–III is that the perturbation operators are *label-preserving*, i.e., although the input is degraded or corrupted, the correct answer should remain unchanged. To verify that this assumption is aligned with human judgment, we conduct a human validity audit over all invariant-label operators.

For each operator, we randomly sample 200 perturbed video–QA pairs and ask three independent annotators, blinded to the original gold labels, to first judge whether the perturbed instance remains answerable and, if so, to select the correct multiple-choice option. We aggregate annotations by majority vote and report two measurements: *answerability*, the fraction of perturbed instances that remain well-posed, and *label preservation*, the fraction of answerable instances whose majority-vote label matches the original gold answer. This protocol provides a direct human check of the invariant-label assumption used by RR.

Table A3 shows that the single-operator perturbations used in Modes II–III are largely consistent with human judgments: answerability ranges from 98.5% to 100.0%, and label preservation ranges from 99.5% to 100.0%. These results support the use of RR as a reliability metric under visual degradation and evidence corruption.

Caption injection is additionally analyzed at the subtitle-subtype level, because it mixes two distinct forms of textual interference: content-irrelevant subtitles and intentionally misleading subtitles. As shown in Table A4, irrelevant captions fully preserve answerability and labels, whereas misleading

captions yield slightly lower answerability (94.0%) and label preservation (94.5%), which is expected given their adversarial design. Even so, the rates remain sufficiently high to justify treating caption injection as an approximately label-preserving operator in our evaluation.

Operator	Answerability	Label-pres.
Compression (cmp)	99.5%	100.0%
Subtitle corruption (sub)	99.0%	99.5%
Gaussian noise (gau)	100.0%	100.0%
Motion blur (mb)	100.0%	99.5%
Adversarial noise (adv)	98.5%	99.5%

Table A3: **Human validity audit for invariant-label operators.** Each operator is audited on 200 perturbed video–QA pairs with three blinded annotators and majority-vote aggregation. High answerability and label-preservation rates indicate strong agreement between the automatic invariant-label evaluation and human judgments.

Subtype	Answerability	Label-pres.
Irrelevant	100.0%	100.0%
Misleading	94.0%	94.5%

Table A4: **Caption injection validity by subtype.** Irrelevant captions fully preserve answerability and labels. Misleading captions show slightly lower rates, consistent with their adversarial design.

A.4 Factuality vs. Knowledge Gaps

To probe whether factuality errors are predominantly hallucination-like (high-confidence wrong) or knowledge-gap-like (low-confidence/uncertain wrong), we conduct a behavioral diagnostic on 369 factuality questions across 8 open-source models, yielding 2,952 model-question pairs (8×369). Each model is evaluated with standard prompting (clean run) and additionally asked to report a self-assessed confidence score (diagnostic run). Wrong instances are classified as: *hallucination-like* if the reported confidence is $\geq 70\%$, *knowledge-gap-like* if the confidence is $\leq 40\%$ or the model expresses explicit uncertainty, and *other* otherwise.

Among 888 analyzable wrong instances, 797 (89.75%) exhibit hallucination-like behavior, while only 39 (4.39%) show knowledge-gap-like patterns. This suggests that factuality failures in current Video-LLMs are overwhelmingly overconfident, producing wrong answers with high self-reported certainty rather than acknowledging uncertainty.

We note that this is a behavioral diagnostic rather than a causal attribution—knowledge gaps may also manifest as overconfident guessing.

A.5 Implementation Details for Operators

Compression (cmp). Videos are re-encoded using FFmpeg with a target bitrate retaining approximately 15.19% of the original (reducing storage from 2.83 GB to 0.43 GB on average), simulating the lossy compression commonly applied by video-sharing platforms.

Subtitle Corruption (sub). Noisy ASR-like subtitles are generated by injecting character-level errors (substitution, deletion, insertion) into ground-truth transcriptions, with intentional subtitle–video desynchronization (random temporal shifts of 0.5–2.0 seconds). This simulates the unreliable OCR/ASR outputs and timing misalignment encountered in deployment.

Gaussian Noise (gau). Gaussian noise is applied uniformly at the frame level to simulate sensor noise or low-light recording conditions. We add zero-mean Gaussian noise to the RGB channels of each frame. The perturbation variance is controlled to keep the noise bounded, ensuring that the high-level semantic content required to answer the question is preserved while degrading low-level perceptual clarity.

Motion Blur (mb). Motion blur is applied uniformly across all frames to simulate camera shake or fast-moving subjects. We apply a linear motion blur filter with randomized kernel sizes and angles to the spatial dimensions of each frame. This operation reduces visual sharpness and introduces perceptual ambiguity without altering the underlying dynamic events.

Caption Injection (cap). To create cross-modal conflicts, we overlay synthesized subtitles onto randomly sampled video segments. As analyzed in Appendix A.3, the injected subtitles consist of a 4:1 mixture of content-irrelevant sentences and LLM-generated misleading statements. The misleading statements are explicitly conditioned on the question and the ground-truth answer option to form a plausible but incorrect textual cue (e.g., pairing a video of opening a door” with the text closing the door”), testing the model’s ability to prioritize authentic visual evidence over misleading text.

Adversarial Noise (adv). We apply a transfer-based black-box perturbation targeting the visual encoder, generated using the MI-FGSM (Dong et al., 2018) algorithm. To ensure the adversarial noise generalizes across different Video-LLMs architectures rather than overfitting to a single proxy, the perturbations are crafted using a proxy ensemble of varied visual encoders (InternVL3-8B (Zhu et al., 2025), Qwen3VL-8B (Bai et al., 2025a), and Video-MAE (Tong et al., 2022)).

Temporal Intervention (shu & rev). Applied exclusively to the order-sensitive subset ($\mathcal{S}_{\text{order}}$), these operators disrupt the temporal structure essential for correctness. **Shuffling (shu)** randomly permutes the sequence of frames across the video, disrupting both local motion and global event order. **Reversal (rev)** strictly inverts the chronological order of the frames from end to start, reversing the direction of actions and state transitions. Both operators render the original ground-truth label invalid by destroying the necessary causal dependencies.

A.6 Extended Related Work

Beyond hallucination benchmarking, a complementary line of work studies *long-video reasoning* and *hallucination mitigation* in (video) large vision-language models. We summarize representative efforts here and discuss their relation to INFACT.

Long-video reasoning. VideoPro (Li et al., 2025b) proposes adaptive program reasoning for long video understanding, decomposing queries into executable programs that selectively attend to relevant temporal segments. Such program-level routing is orthogonal to our diagnostic goal: INFACT measures *whether* Video-LLMs remain faithful and factual under controlled perturbations, while methods like VideoPro aim to *improve* reasoning efficiency on long inputs. Our order-sensitive subset and temporal intervention protocol could serve as a stress test for such reasoning pipelines.

Permutation-induced hallucinations. Stable-RAG (Zhang et al., 2026) identifies and mitigates hallucinations induced by the *permutation* of retrieved passages in retrieval-augmented generation, showing that order sensitivity is a non-trivial source of unreliability even when the underlying evidence is unchanged. This closely parallels our Temporal Intervention mode (shu/rev), which exposes how Video-LLMs over-rely on positional priors rather

than actual temporal evidence. Together, these findings suggest that *order-induced hallucination* is a cross-modal phenomenon worthy of unified study.

Hallucination mitigation in LVLMs. A growing body of work mitigates hallucinations by intervening on model internals or outputs. Lyu et al. (2026) identify and enhance core visual regions through internal attention dynamics, while Lyu et al. (2025) diagnose perception-generation imbalance and propose a validated dominance correction to rebalance modality contributions. Liu et al. (2026) instead target the over-reliance trap via a self-validation framework that re-checks object-level claims against the visual input. These mitigation strategies are complementary to INFACT: our benchmark provides a fine-grained diagnostic signal (RR under invariant-label perturbations, TSS on order-sensitive items) that such methods can be evaluated against, moving beyond clean-accuracy reporting toward reliability-aware evaluation.

Task & Definition	Example
1. Static Entities & Attributes	
1.1 Entity Recognition <i>Identify an entity label supported by video evidence.</i>	<i>Which human organ is visible in the video?</i> (A) Stomach (B) Liver (C) Lungs (D) Kidneys
1.2 Unique Entity Counting <i>Count the number of distinct entity instances.</i>	<i>How many babies does the lion mother have in the video?</i> (A) 4 (B) 3 (C) 2 (D) 5
1.3 Temporal Attributes Recognition <i>Attribute query with an explicit temporal/event anchor.</i>	<i>What is the color of the second object that enters the scene?</i> (A) Brown (B) Cyan (C) Gray (D) Purple
1.4 Static Attributes Recognition <i>Attribute query without explicit temporal/event anchor.</i>	<i>What color is the T-shirt worn by the boy?</i> (A) Yellow (B) White (C) Black (D) Blue
1.5 Scene Text Recognition <i>Detecting and reading text that appears in the video.</i>	<i>What is written on the first made keychain?</i> (A) Google (B) YouTube (C) Facebook (D) Spotify
2. Dynamic Actions & Motions	
2.1 Action Recognition <i>Discriminate fine-grained / near-confusable action.</i>	<i>Which description correctly matches the actions?</i> (A) Attacking (B) Chasing (C) Running (D) Stealing
2.2 Repetitive Action Counting <i>Count occurrences of actions across the video.</i>	<i>How many times does the hammer's handle hit the floor throughout the video?</i> (A) 4 (B) 3 (C) 1 (D) 0 (E) 5 (F) 2
2.3 Motion Attributes Recognition <i>Identify motion attributes (direction, rotation, speed, trajectory) of an entity.</i>	<i>In what direction(s) is the wheel rotating?</i> (A) Counter-clockwise throughout (B) Clockwise throughout (C) No rotation (D) Counter-clockwise then clockwise.
3. Spatio-Temporal Relations	
3.1 Spatial Relation Recognition <i>Infer spatial relations or spatial locations in the video.</i>	<i>Where is the hidden object at the end of the game?</i> (A) Under 1st left (B) Under 3rd left (C) Under 2nd from left
3.2 Temporal Relation Recognition <i>Determine the temporal relations among events.</i>	<i>What happened after the person sat on the sofa/couch?</i> (A) Opened the box (B) Opened the door (C) Sat on the floor (D) Put down the clothes.
3.3 State Transition Detection <i>Determine whether a state variable changes over time.</i>	<i>Is the bag empty at the end?</i> (A) Yes (B) No (C) The person doesn't interact with a bag
3.4 Temporal Localization <i>Localize when an action/state/event holds within the video.</i>	<i>When in the video does the action occur?</i> (A) Beginning (B) Middle (C) End (D) Throughout

Table A5: Faithfulness Hallucination Taxonomy.

Category & Definition	Example
1. Domain Knowledge (Know-WHAT)	
1.1 Cultural Event Recognition <i>Knowledge of traditions, festivals, artistic heritage, and customs.</i>	<i>When is the festival in the video usually celebrated?</i> (A) 9th day of 9th lunar month (B) 8th day of 12th lunar month (C) 5th day of 5th lunar month (D) 1st day of 1st lunar month
1.2 Historical Background Identification <i>Knowledge of past events, chronologies, archaeology, and historical figures.</i>	<i>When was the Site shown in the video discovered?</i> (A) 1928-1929 (B) 1925-1926 (C) 1920-1922 (D) 1934-1938
1.3 Geospatial Localization <i>Knowledge of geospatial locations, landmarks, and regional characteristics.</i>	<i>Which city is this video most likely recorded in?</i> (A) San Diego (B) Phoenix (C) Tucson (D) Las Vegas
1.4 Entertainment-Related Recognition <i>Knowledge of movies, TV, music, celebrities, and sports.</i>	<i>What genre of movie is portrayed in the video?</i> (A) Reality TV (B) Documentary (C) Thriller (D) Romance
2. Procedural Knowledge (Know-HOW)	
2.1 Electronic Procedural Judgement <i>Logical steps for repairing computers and electronic hardware.</i>	<i>Is the sequence of steps about task <Replace Laptop Screen> shown logical?</i> (A) Yes, the sequence is correct (B) No, the sequence is incorrect
2.2 Mechanical Procedural Judgement <i>Sequence and tools for vehicle maintenance and mechanical repair.</i>	<i>Is the sequence of steps about task <Replace Car Door Handle> shown logical?</i> (A) Yes, the sequence is correct (B) No, the sequence is incorrect
2.3 Domestic Procedural Judgement <i>Workflow for household DIY, plumbing, and furniture repair.</i>	<i>Is the sequence of steps about task <Replace Light Socket> shown logical?</i> (A) Yes, the sequence is correct (B) No, the sequence is incorrect
2.4 Clinical Procedural Judgement <i>Medical protocols for examinations, first aid, and clinical operations.</i>	<i>Is the sequence of steps about task <Measure the Jaw Opening> shown logical?</i> (A) Yes, the sequence is correct (B) No, the sequence is incorrect
3. Physical Knowledge (Know-WHY)	
3.1 Newtonian Mechanical Reasoning <i>Newtonian laws: gravity, friction, collision, and momentum.</i>	<i>Analyze the physical dynamics in the video, which statement best describes the adherence to physical laws?</i> (A) Fully Plausible (B) Violation of Mechanical Dynamics (C) Violation of Material Properties (D) Violation of Fluid Dynamics
3.2 Fluid Mechanics Reasoning <i>Behavior and interaction of liquids, smoke, and fire.</i>	<i>Analyze the physical dynamics; which statement best describes the adherence to laws?</i> (A) Fully Plausible (B) Violation of Mechanics (C) Violation of Fluids (D) Violation of Materials
3.3 Material Property Reasoning <i>Realism of deformation, hardness, and fracture of solids.</i>	<i>Analyze the physical dynamics; which statement best describes the adherence to laws?</i> (A) Fully Plausible (B) Violation of Mechanics (C) Violation of Fluids (D) Violation of Materials
3.4 Spatio-temporal Continuity Reasoning <i>Continuity and permanence of objects under occlusion or movement.</i>	<i>Analyze the physical dynamics; which statement best describes the adherence to laws?</i> (A) Fully Plausible (B) Violation of Object Permanence (C) Violation of Continuity

Table A6: **Factual Hallucination Taxonomy.** The framework evaluates world knowledge across three dimensions: Domain Knowledge (Know-WHAT), Procedural Knowledge (Know-HOW), and Physical Knowledge (Know-WHY).