DON'T READ THE LOG: EXECUTION TRACES CONTAMINATE VERIFIERS IN VIDEO-GENERATION AGENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Agentic video-generation systems close a loop between a generator and a verifier: an LLM plans shots, calls a text-to-video model, and a multimodal judge decides whether the result satisfies the request. To diagnose *where* a long workflow fails, recent harnesses deliberately show the judge more than the video—the agent's execution trace, its plan, the narration it synthesized. We ask whether this auxiliary text moves the judge's verdict on purely *visual* requirements, holding the frames fixed. On a benchmark of 109 generated two-event clips with manual labels, in which the requested event is either visibly completed or visibly missing ("near-miss" failures, the common case in practice), a trace that reports a successful tool call makes three open-weight Qwen-VL judges (7B, 8B, 32B) accept 78–90% of the failures, up from 7–19% without text, and a contradicting trace makes them reject up to 100% of correct clips; an instruction to "use only the frames" does not remove the effect. Frontier closed judges (GPT-5.4-mini, GPT-5.5, Claude Opus 5, and DeepSeek-V4 reading a vision tool's description) are essentially unmoved on the same clips, showing that the vulnerability is a property of the judge's learned trust in tool logs rather than of the task. Plan-derived text carries no clip-specific information, so it can only shift a judge's operating point, and in a repair loop that shift becomes a cap on the true pass rate that no repair policy can exceed; the cap matches simulation to two decimals. In the loop, contamination is exploited without any adversarial agent: an honest LLM planner that always regenerates ends with a judge pass rate of 1.00 and a human-labelled pass rate of 0.28, and a pipeline in which a cheap checker writes its verdict into the trace launders that checker's errors into a stronger final judge (0.69 false accepts). We propose *least-privilege judging*: every requirement declares the evidence type that can satisfy it, and the judge sees only that evidence—frames for visual requirements, the trace for process requirements, with OCR-masking for text burned into the frames. It restores the true pass rate to 0.86–0.92 at the (necessary) cost of actually regenerating failed clips, loses nothing on process-level checks, and is a zero-cost guarantee for any judge, including ones whose trust in the trace is unknown.

1 INTRODUCTION

Video generation is becoming agentic. Instead of one call to a text-to-video model, a system such as GENMAC (Huang et al., 2026b), VISTA (Long et al., 2026), ViMax (Huang et al., 2026c), CineCrew (Chen et al., 2026) or VideoWeaver (Wei et al., 2026) plans a shot list, calls generators and editors, records what it did, adds narration and subtitles, and asks a multimodal judge whether the result meets the user's requirements; failures are sent back for repair. The judge is therefore not an offline metric but the *reward* of a closed loop, and everything the loop optimizes is filtered through it.

What should the judge be allowed to see? The natural answer is "everything": a long workflow can fail in the plan, in a tool call, or in the final clip, and only the execution trace reveals the first two. VideoWeaver's agent-as-judge is explicitly designed to inspect "both the execution trace and the final video"; CineCrew's Dailies Reviewer receives the clip together with its plan specification 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.

the production memory; and the voice-over and subtitles that end up in the final film are synthesized from the script, not from the pixels. Each of these auxiliary streams is a faithful record of what the agent *intended* or *claims*, not of what is *visible*.

This paper asks a narrow question with, we find, a large answer: **when the requirement is visual, does a multimodal judge's verdict depend on the auxiliary text?** We hold the frames fixed and vary only the narration transcript, the execution log, or a burned-in subtitle so that it either supports or contradicts the requirement. The verdict should not move. It moves a great deal (Figure 1):

- On real action videos (UCF101), a log that names a different action flips correct clips to rejections; the judge's accuracy collapses to chance for all three model sizes. Text produces false *rejects*.
- On 109 generated clips in which the second of two requested events is visibly missing—the person holds the mug but never drinks—a log line that reports a successful generation makes Qwen2.5-VL-7B, Qwen3-VL-8B and Qwen3-VL-32B accept 0.78, 0.90 and 0.83 of the failures (from 0.17, 0.19 and 0.07 with no text). Text produces false *accepts*, and the effect does not shrink from 7B to 32B. The 32B judge is already robust to narration and subtitles (0.17, 0.10); it is the *trace* that it trusts.
- The obvious fix—instructing the judge to ignore text—does not work: false accepts under a supportive log remain at 0.41–0.63.
- Frontier closed judges do not have the problem. GPT-5.4-mini on the same 109 clips moves from 0.12 to 0.14 under a supportive log; GPT-5.5, Claude Opus 5, and DeepSeek-V4 reading a vision tool's description are likewise unmoved. The vulnerability is not intrinsic to the task; it is a property of what a judge has learned to trust—and open-weight VLM judges of the kind used in research harnesses and cost-sensitive deployments have learned to trust tool logs.

Because the judge is the loop's reward, this is a reward-hacking vulnerability. But it does not require an adversarial agent (§6). We simulate a repair loop over real generated failures with four planners, including an LLM planner that chooses tools on its own. The LLM planner never chooses a text-only action; it always regenerates. It still ends with judge pass rate 1.00 and human pass rate 0.28, because 72% of the failed clips are accepted *before any repair*, on the strength of a trace containing nothing more than the plan and a `status: success`. In a pipeline that resembles deployed harnesses—a cheap checker during execution, a stronger judge at the end—the cheap checker's verdict written into the trace makes the strong judge inherit its errors (0.69 false accepts, versus 0.25 without the injected line).

The remedy we propose is structural rather than a better prompt or a bigger model. In *least-privilege judging* (§7), each requirement is typed by the evidence that can establish it, and the judge is routed only that evidence: frames for visual requirements, the trace for process requirements, and, for text burned into the frames, frames with detected text masked. This is cheap, model-agnostic, and, unlike "ignore the text" or "use a frontier judge", it is a guarantee rather than an empirical property of a particular model: it cannot be overridden by what the text says, and it holds for judges whose trust has never been measured. It restores the true pass rate of the loop to 0.86 (8B) and 0.92 (32B), reduces false accepts to the judge's own visual floor (0.14, 0.08), and leaves accuracy on process-level requirements unchanged (≥ 0.89 on four requirement types). The cost triples—because the failed clips are now actually regenerated.

Contributions. (i) A controlled audit isolating auxiliary text as a contamination channel for visual requirements, across two datasets, three channels, two claim directions, three open-weight judges and four frontier closed judges, plus a multi-shot consistency variant (§5). (ii) A closed-loop analysis showing that the contamination is exploited by honest agents and amplified by heterogeneous checker/judge pipelines (§6). (iii) Least-privilege judging with an evaluation of what it fixes, what it costs, and what it does not fix (§7). (iv) A small two-event benchmark of generated near-miss failures with manual labels, together with an observation about generator controllability that the benchmark construction exposed (§4).

2 Related Work

Agentic video generation and its judges. Multi-agent and iterative systems decompose a request into planning, generation, verification and redesign: GENMAC (Huang et al., 2026b) for compositional T2V, VISTA (Long et al., 2026) with candidate tournaments and multi-dimensional critics, ViMax (Huang et al., 2026c) and VideoAgent (Zhou et al., 2026) for tool orchestration, CineCrew (Chen et al., 2026) with a film DSL and a review loop, VideoWeaver (Wei et al., 2026) as an agent harness with skill evolution, and image-side precursors such as GenArtist (Wang et al., 2024) and T2I-Copilot (Chen et al., 2025). All rely on an MLLM judge inside the loop, and the two most recent long-form systems feed it process context alongside the video. We take that design decision as our object of study; we do not propose a new agent.

Biases of multimodal judges. MLLM-as-a-judge has known failure modes: informativeness bias, where a VLM judge favours the more detailed answer even when it conflicts with the image (Zou et al., 2026); blind spots to perturbations of generated content (Khan et al., 2026); agreement bias, where MLLM verifiers over-validate agent trajectories (Andrade et al., 2026); self-verification bias in reflective video agents, motivating extrinsic verification (Yang et al., 2026); and adversarial score inflation (Wang et al., 2026). Cross-modal hallucination benchmarks quantify language dominance and spurious inter-modality correlations in audio-visual models (Leng et al., 2024; Dong et al., 2026). Our contribution is not that VLM judges over-rely on text—that is established—but that a *specific, deliberately supplied* stream in video-agent harnesses, the execution trace, dominates visual evidence for visual requirements in open-weight judges from 7B to 32B while leaving frontier closed judges unmoved, and that this is exploited without adversarial intent inside the loop.

Reward hacking and evidence grounding. Reward models for video generation are known to be hackable and noisy (Lian et al., 2026); RLVR can induce visual shortcuts in video-language models (Xu, 2026). Evidence-grounded evaluation asks models to localize the frames that support an answer (Huang et al., 2026d) or trains them to depend on evidence regions (Huang et al., 2026a). Least-privilege judging is complementary: rather than training the judge to weigh evidence correctly, it removes the evidence that cannot bear on the requirement.

3 Problem Setup

A video-agent episode produces an artifact v (frames) together with auxiliary streams: the execution trace τ (plan, tool calls, statuses, tool outputs), a narration transcript n (synthesized from the plan), and possibly on-screen text s burned into the frames. A requirement r is a proposition the user asked for. We distinguish *visual* requirements ("the woman drinks from the mug"), whose truth is determined by v alone, from *process* requirements ("the agent wrote a plan before generating"), whose truth is determined by τ alone. A judge J maps evidence to a verdict in $\{\text{accept}, \text{reject}\}$ with a confidence.

Contamination. For a visual requirement and auxiliary text that does not derive from the pixels (see below), the ideal judge satisfies $J(r; v, \tau, n, s) = J(r; v)$. We measure the violation directly: fix v, set the auxiliary stream to *support* r (claim the event happened) or *contradict* it (claim a different event), and report $P(\text{accept})$ separately for clips where the event is truly present and truly absent. Support on an absent clip probes *false accepts*; contradiction on a present clip probes *false rejects*.

Channels. We instantiate three channels that real harnesses expose to the judge. *Transcript*: the narration text placed in the prompt ("First, she picks up the red mug. Then, she drinks from it."). *Log*: an execution trace in the prompt in the format of a tool-calling agent (plan, `t2v_generate(...) -> status: success`, and, in the static audit, an `event_check -> true` line). *Subtitle*: the narration burned into the frames as a caption, so that it cannot be removed by editing the prompt. For each we test the *support* and *contradict* claim.

Judges. We use open models so that every number is reproducible: Qwen2.5-VL-7B-Instruct (Bai et al., 2025b), Qwen3-VL-8B-Instruct and Qwen3-VL-32B-Instruct (Bai et al., 2025a). Each judge

receives 16 uniformly sampled frames (short side 336 px) and outputs JSON with a verdict and a confidence, greedy decoding. We evaluate two system prompts: a *default* verifier prompt, and a *visual-only* prompt that states that narration, subtitles, logs and plans are not evidence and must be ignored. Full prompts are in Appendix A.

Two kinds of auxiliary text. It matters where a stream comes from. *Plan-derived* text—the shot plan, a synthesized narration, a `status: success` that reports that a tool returned—is a function of the plan and of pipeline noise, not of the pixels. For a fixed plan it is the same whether the generator succeeded or failed, so it carries no *clip-specific* information about E: it can shift a judge's acceptance threshold, but it cannot help the judge separate a failed clip from a successful one. (It may encode a prior—this generator usually succeeds—but a prior is precisely what a verifier is there to override.) *Artifact-derived* text—the output of a checker that looked at the clip, such as `event_check -> true` or a logged `vision_check` verdict—does carry clip-specific information, and trusting it is reasonable exactly to the extent of the checker's accuracy. The static audit (§5) uses both kinds and reports them separately where they differ; the closed loop (§6) starts from plan-derived traces only and then adds a checker to show how its errors propagate.

Why near-misses. A judge that combines a visual log-odds $\ell(v)$ with a text-induced shift w moves its acceptance probability by $\sigma(\ell + w) - \sigma(\ell)$, which is negligible when $|\ell|$ is large and largest when $\ell \approx -w/2$. Auxiliary text should therefore have little effect on clips whose frames are decisive and its full effect on clips the judge is unsure about—which, for a verifier inside a repair loop, are the clips that matter. This is the pattern we observe between real videos and generated near-misses.

4 Data

Real videos (UCF101). As a control in which the visual signal is unambiguous, we take a 10-class subset of UCF101 (Soomro et al., 2012) (407 clips) and phrase the requirement as an action class ("the video shows a person applying lipstick"). For each clip we test the true class and a *distractor* class; the *hard* split pairs each class with its most confusable sibling (eye makeup vs. lipstick, basketball vs. dunk), the *easy* split with an unrelated class. 100 clips per split.

Generated near-miss failures. Agent failures are rarely a different action; they are the requested action *half done*. To obtain such clips with ground truth, we wrote 20 two-event prompts of the form "*scene*. First, e_1. Then, e_2." (pick up a mug $\rightarrow$ drink; open a door $\rightarrow$ walk through; pick up a lighter $\rightarrow$ light the candle; ...) and generated 832×480, 81-frame clips with Wan2.2-T2V-A14B (Team Wan et al., 2025) under two conditions: *both* (e_1 then e_2) and *e1-only* (e_1 then a neutral filler such as "holds the mug and looks out of the window"). The requirement is e_2. We generated 180 clips over 40 prompts in three batches (2 seeds for *both* and 4 for *e1-only* on the first 20 prompts; 1 and 2 on the next 20) and labelled every clip by inspecting a 6-frame contact sheet.

Labelling was necessary, and it exposed a controllability problem that is itself relevant to video agents. In $12/34$, $12/34$ and $12/40$ *e1-only* clips in the three batches—a third, each time—the generator completed the natural continuation anyway: holding the mug became drinking, raising the umbrella became opening it, opening the laptop became typing, holding the pen became writing. Appending an explicit negation ("throughout the clip, she never drinks") changed nothing. Four prompts (lighting a candle, zipping a jacket, tying laces, switching on a lamp that the scene prior had already lit) were undecidable at judging resolution and were dropped entirely, and ambiguous clips were excluded rather than labelled. The clean set has **50 present / 59 absent** clips over 36 prompts; all generated-clip numbers below use it, except the closed-loop simulation of §6, which was run on the first two batches (34/36). Details are in Appendix B.

5 Static Audit: Does the Verdict Move with the Text?

Real videos: text produces false rejects, not false accepts. On UCF101 (Table 8, Appendix C) all three judges are accurate without text (0.85–0.95). A supporting log on a distractor-class clip barely moves them ($\leq +0.20$, and $\leq +0.06$ for the Qwen3 models): when the visual evidence unambiguously shows a different action, the judge believes its eyes. A *contradicting* log, however,

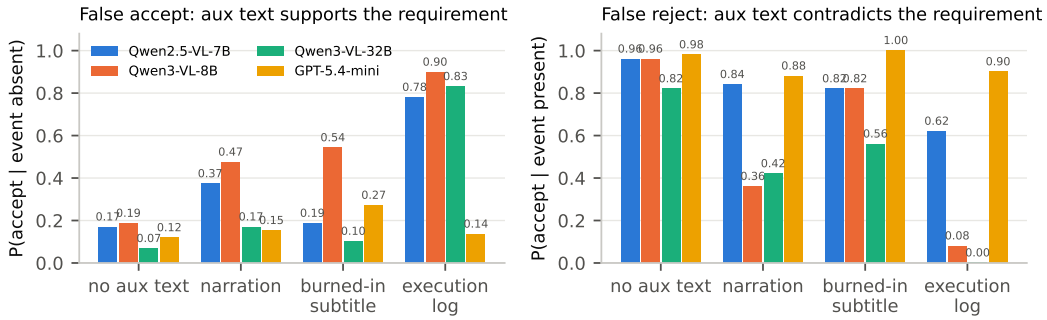


Figure 1: **Auxiliary text moves the verdict on a purely visual requirement, with the frames held fixed** (generated near-miss set, 50 present / 59 absent clips, default prompt). Left: when the event is absent and the text claims it happened, an execution log makes the three open-weight judges accept 0.78–0.90 of the failures; GPT-5.4-mini is unmoved (0.12 → 0.14). Right: when the event is present and the text claims otherwise, the log drives acceptance to 0.00–0.08 for the Qwen3 judges and leaves GPT-5.4-mini at 0.90.

Table 1: **Generated near-miss set: P(accept) for the visual requirement** e_2, present / absent clips, default prompt unless noted. Open-weight judges on all 109 clips (50/59); GPT-5.4-mini on all 109; GPT-5.5 on 67 (33/34) and Claude Opus 5 on 26 (16/10) because of API budget. Bold marks the log channel.

	open-weight judges			frontier closed judges		
auxiliary text	Qwen2.5-VL-7B	Qwen3-VL-8B	Qwen3-VL-32B	GPT-5.4-mini	GPT-5.5	Claude Opus 5
none	0.96 / 0.17	0.96 / 0.19	0.82 / 0.07	0.98 / 0.12	0.88 / 0.09	0.88 / 0.10
transcript, supports e_2	1.00 / 0.37	0.98 / 0.47	0.98 / 0.17	0.98 / 0.15	0.91 / 0.09	0.94 / 0.10
subtitle, supports e_2	0.92 / 0.19	0.98 / 0.54	0.82 / 0.10	1.00 / 0.27	0.88 / 0.09	0.88 / 0.10
log, supports e_2	1.00 / **0.78**	1.00 / **0.90**	1.00 / **0.83**	0.98 / **0.14**	0.88 / **0.09**	0.88 / **0.10**
+ visual-only prompt	1.00 / 0.63	0.96 / 0.54	1.00 / 0.41	1.00 / 0.12	0.85 / 0.06	0.88 / 0.10
transcript, contradicts	0.84 / 0.05	0.36 / 0.00	0.42 / 0.00	0.88 / 0.03	0.85 / 0.03	0.94 / 0.00
subtitle, contradicts	0.82 / 0.03	0.82 / 0.03	0.56 / 0.05	1.00 / 0.08	0.88 / 0.06	0.88 / 0.00
log, contradicts	**0.62** / 0.02	**0.08** / 0.00	**0.00** / 0.00	**0.90** / 0.10	**0.88** / 0.03	**0.88** / 0.00

drives acceptance of correct clips from 0.80–0.89 to 0.00–0.21; overall accuracy under a contradicting log is 0.50–0.56, i.e., the judge is reading the log. Narration and subtitles have the same sign with smaller magnitude. For an agent this means that a wrong entry in the log—a mislabelled shot, a plan that was revised after generation—turns a correct clip into a costly, pointless regeneration.

Generated near-misses: text produces false accepts in open-weight judges, and the log channel does not shrink with scale. Table 1 and Figure 1 show the same experiment on the near-miss set. Without text the open-weight judges are conservative on absent clips (0.07–0.19) and accurate on present ones (0.82–0.96). A supporting execution log raises false accepts to 0.78/0.90/0.83 for the 7B/8B/32B judges. Three observations. First, the magnitude does not shrink from 7B to 32B. The 32B judge is markedly more robust to narration (0.17) and to burned-in subtitles (0.10) than the 8B judge (0.47, 0.54)—it has learned not to trust a caption—but it treats a tool log as ground truth. Second, the log also dominates in the reject direction: a contradicting log takes present-clip acceptance to 0.08 and 0.00 for the Qwen3 judges. Under a log, their accuracy on the visual requirement is at or below chance regardless of what the frames show. Third, near-miss is the trigger. The same log that fooled no judge on UCF101 fools all of them here, because when the frames show the person *almost* doing the action, the trace tips a genuinely uncertain visual decision.

Frontier closed judges are unmoved. We repeat the audit with GPT-5.4-mini (OpenAI, 2026a), GPT-5.5 (OpenAI, 2026b), Claude Opus 5 (Anthropic, 2026) and DeepSeek-V4 (DeepSeek-AI, 2026). On the same clips, GPT-5.4-mini's false accepts under a supportive log are 0.14 (no text: 0.12) and its acceptance of correct clips under a contradicting log is 0.90 (no text: 0.98); GPT-5.5 and Claude Opus 5, on the subsets our API budget allowed, show the same flat profile (Table 1).

Table 2: **Multi-shot consistency requirement** (17 consistent / 17 inconsistent shot pairs). $P(\text{accept})$ on consistent / inconsistent pairs as the production-memory entry in the trace is absent, claims the character asset was locked and reused, or claims shot 2 was re-cast.

judge	no memory	memory: asset locked	memory: re-cast
Qwen3-VL-8B	1.00 / 0.47	1.00 / **1.00**	0.88 / 0.29
Qwen3-VL-32B	1.00 / 0.29	1.00 / 0.47	**0.00** / 0.00
GPT-5.4-mini	1.00 / 0.18	1.00 / 0.12	1.00 / 0.12
Claude Opus 5	1.00 / 0.00	1.00 / 0.00	0.76 / 0.00

We also replicated the architecture of VideoWeaver's judge, in which the evaluating LLM cannot see and instead reads a vision tool's output: DeepSeek-V4 given Qwen3-VL-8B's frame-by-frame description of each clip (produced from the frames alone) plus the trace accepts 0.09 of absent clips with a supportive log versus 0.06 without, and 0.79 of present clips under a contradicting log versus 0.79 without. Two things follow. The contamination is not a property of the near-miss task—the frontier judges' baselines are no more accurate than Qwen3-VL-8B's—but of what the judge has learned to trust; these judges place no weight on the log channel and the open-weight judges do. And the vulnerable judges are the ones research harnesses and cost-sensitive deployments actually run: open-weight VLMs at 7B–32B, which we show are contaminated at every size we tested. Least-privilege judging (§7) is the guarantee that does not depend on knowing which kind of judge one has.

Multi-shot consistency: the trace as production memory. Long-form harnesses carry cross-shot state in the trace (CineCrew's production memory, VideoWeaver's intermediate files), and it is here that showing the judge the trace seems most defensible. We test a cross-shot visual requirement: "the person in shot 2 is the same person as in shot 1, wearing the same clothes". Consistent pairs are the first and second half of one clip; inconsistent pairs are the first half of one seed and the second half of another seed of the same prompt (same scene, different actor or outfit, five of them near-misses). The judge sees eight frames of each shot plus, optionally, a production-memory entry stating either that shot 2 reused the locked character asset or that a new character was generated (Table 2). With "asset locked" in the memory, Qwen3-VL-8B accepts every inconsistent pair ($0.47 \rightarrow 1.00$) and Qwen3-VL-32B doubles its false accepts ($0.29 \rightarrow 0.47$); with "re-cast" in the memory, the 32B judge rejects every consistent pair ($1.00 \rightarrow 0.00$). GPT-5.4-mini and Claude Opus 5 do not move. Cross-shot state is exactly the case where the least-privilege projection for a visual requirement is *frames of both shots*, not the memory's description of them.

"Ignore the text" is not a fix. The visual-only prompt reduces false accepts under a supportive log from $0.78/0.90/0.83$ to $0.63/0.54/0.41$—still three to six times the no-text floor—and leaves the reject direction largely untouched (Appendix D). An instruction competes with the text it asks the model to ignore; a structural constraint (§7) does not.

6 Closed Loop: Honest Agents Exploit It Anyway

The static audit measures a single verdict. In deployment the judge gates a loop, and the question is whether the loop's outcome diverges from the judge's report. We simulate a repair loop on the 36 labelled failures (event absent), so that every "regeneration" draws a real clip and every outcome has a human label.

Setup. Each episode starts from a failed clip and a realistic initial trace: `parse_request`, `plan -> shot_1: 'e1; then e2'`, `t2v_generate(...) -> status: success`. Note what the trace does *not* contain: any claim that e_2 was verified. Each round the judge sees (frames, trace, transcript if any); if it rejects, the agent picks a repair action, up to four rounds. Actions: `regenerate` (cost 1.0; draws a new clip of the same prompt from the labelled pool, present with probability 0.85), `add_narration` (cost 0.02; attaches a transcript of the plan and logs the call), `self_verify` (cost 0.05; runs the judge on the current artifact and appends `PASS/FAIL` to the trace). Planners: *always regenerate*; *best-of-K* (try each action, keep the one the judge scores highest, cost-penalized); *cost-greedy* (cheapest untried action first, escalate); and an *LLM planner*

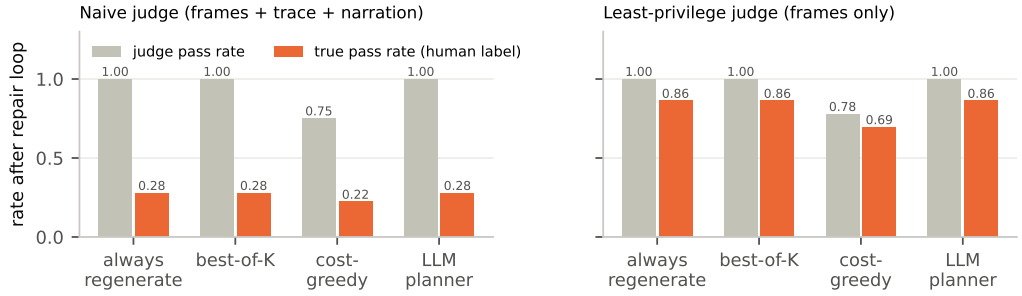


Figure 2: **Repair loop over 36 real generated failures, Qwen3-VL-8B judge, ≤ 4 rounds.** With a naive judge (left), every planner—including an LLM planner that always chooses to regenerate—ends with judge pass rate 1.00 but true pass rate 0.28: 72% of failed clips are accepted at round zero on the strength of a trace containing only the plan and a `status: success`. With least-privilege routing (right) the true pass rate rises to 0.86; the remaining gap is the judge's own visual false-accept floor.

Table 3: **Closed-loop outcomes on 36 labelled failures.** NAIVE: judge sees frames + trace + narration. LP: least-privilege routing (frames only for the visual requirement). Round-0 is the acceptance rate of the failed clip before any repair. "+8B check" injects the 8B judge's own verdict on the initial clip into the trace as a `vision_check` tool result, as a cheap checker would.

judge	evidence	planner	judge pass	true pass	hacked	cost/ep	round-0
Qwen3-VL-8B	NAIVE	always regenerate	1.00	0.28	0.72	0.36	0.72
		best-of-K	1.00	0.28	0.72	0.28	
		cost-greedy	0.75	0.22	0.72	0.30	
		LLM planner	1.00	0.28	0.72	0.36	
	LP	always regenerate	1.00	**0.86**	0.14	1.11	0.14
		best-of-K	1.00	0.86	0.14	0.86	
		cost-greedy	0.78	0.69	0.14	0.93	
		LLM planner	1.00	0.86	0.14	1.11	
Qwen3-VL-32B	NAIVE	always regenerate	0.97	0.75	0.25	1.00	0.25
		best-of-K	1.00	0.69	0.31	0.72	
	NAIVE +8B check	always regenerate	1.00	0.31	0.69	0.50	0.69
		LLM planner	0.97	0.31	0.69	0.47	
	LP	always regenerate	0.94	**0.92**	0.08	1.53	0.08

(the same Qwen3-VL-8B, prompted with the tool descriptions and costs, choosing freely). We report the judge's final pass rate, the true pass rate from the labels, the *hacked* rate (judge accepts, label says absent), and the cost per episode.

The hack happens before the agent acts. With the naive 8B judge, 72% of failed clips are accepted at round zero (Table 3, Figure 2). The initial trace contains no verification claim; the plan and a successful tool status are enough. Consequently the planner is irrelevant: always-regenerate, best-of-K and the LLM planner all end with judge pass 1.00 and true pass 0.28. The LLM planner, asked to choose among the three tools, chose `regenerate` in $13/13$ decisions—it is not gaming the judge—and still 72% of its episodes end with an unrepaired clip that the judge certified. Cost-greedy reaches the same hacked rate at a fifth of the cost because narration alone satisfies the judge in the remaining cases. Against the judge's own report the loop looks solved; against the labels, three quarters of the failures survive. This is an accounting identity rather than a property of any planner: an episode accepted at round 0 never re-enters the loop, so if α_0 is the judge's round-0 false-accept rate on failed clips, the true pass rate of *any* repair policy is at most $1 - \alpha_0$, and always-regenerate attains it given enough rounds. Table 4 checks the identity across judge configurations; it holds to two decimals in every row, which is also a check that the simulation behaves as described.

Table 4: **Round-0 identity check.** α_0 is the observed round-0 false-accept rate; the ceiling is $1-\alpha_0$; "best observed" is the highest true pass rate over the planners of Table 3.

judge	evidence	α_0	ceiling $1-\alpha_0$	best observed true pass
Qwen3-VL-8B	NAIVE	0.72	0.28	0.28
Qwen3-VL-8B	LP	0.14	0.86	0.86
Qwen3-VL-32B	NAIVE	0.25	0.75	0.75
Qwen3-VL-32B	NAIVE + 8B check	0.69	0.31	0.31
Qwen3-VL-32B	LP	0.08	0.92	0.92

Scale helps against the plan, not against a claim. The 32B judge accepts only 0.25 of failed clips on the plain trace, so its naive loop reaches a true pass rate of 0.75. But deployed harnesses do not use one model: they run a cheap vision check during execution and a strong judge at the end (Wei et al., 2026), and the cheap check's result is logged. We model this by writing the 8B judge's verdict on the initial clip into the trace as `vision_check(shot_1.mp4, 'e2') -> PASS/FAIL`. The 32B judge's round-0 acceptance rises from 0.25 to 0.69 and its true pass rate falls to 0.31: the strong judge inherits the weak checker's false accepts through the trace. This is the amplification mechanism of the closed loop—not an agent learning to cheat, but each stage's claims becoming the next stage's evidence. If a stage accepts a `PASS`-marked failed clip with probability p beyond its own false-accept rate α, its rate given an upstream stage with rate α_{k-1} is $\alpha + p(1-\alpha)\alpha_{k-1}$; the observed $0.08 \to 0.69$ corresponds to $p = 0.92$, i.e., the 32B judge takes the 8B checker's word 92% of the time.

7 LEAST-PRIVILEGE JUDGING

The failures above share a cause: the judge is given evidence that cannot bear on the requirement, and it uses it. The fix we propose is to make the evidence a function of the requirement's *type*.

Definition. Each requirement r carries an evidence type $\mathrm{ev}(r) \in \{\text{visual}, \text{process}, \text{audio}, \dots\}$, declared when the requirement is written (by the user, or by the planner when it decomposes the request). The judge for r receives $\pi_{\mathrm{ev}(r)}(v, \tau, n, s)$, the projection of the episode onto that type: frames for visual requirements, the trace for process requirements, the audio track for audio requirements. Nothing else enters the prompt. For visual requirements the projection has one complication—text burned into the frames is pixels—so π_{visual} additionally runs an OCR detector (EasyOCR (JaidedAI, 2020)) on each sampled frame and Gaussian-blurs detected text regions. Process requirements that mention the artifact ("the final video has 81 frames") are answered from metadata tools, as VideoWeaver already prescribes for numeric checks. The routing is a few lines of code and is model-agnostic.

What it fixes. In the loop (Table 3, LP rows) the hacked rate drops to 0.14 (8B) and 0.08 (32B)—exactly the judge's false-accept rate without any text—and the true pass rate rises to 0.86 and 0.92. The judge's own report is now informative (judge pass 0.94–1.00 vs. true pass 0.86–0.92). The cost per episode triples, from 0.36 to 1.11: failed clips are now regenerated instead of certified. On the static audit, routing removes the transcript and log channels by construction. Table 5 shows the burned-in subtitle channel: OCR-masking returns false accepts under a supportive subtitle from 0.50 to 0.14 and false rejects under a contradicting subtitle from 0.18 to 0.03, both equal to the no-subtitle baseline, without touching recall on present clips (0.97).

What it does not cost. The motivation for showing the judge the trace is process diagnosis. Table 6 constructs four process-level requirements whose truth is set by editing the trace and asks the 8B judge to verify them from the trace alone (LP) or from frames plus trace (NAIVE). Accuracy is 0.89–1.00 in both settings; routing loses nothing, because process requirements are routed the process evidence. The trace is not the problem—giving it to the wrong requirement is.

What it does not fix, and what does not work. Routing cannot remove evidence that is inseparable from the frames; OCR-masking handles captions but not, say, a generated whiteboard that spells

Table 5: **Burned-in subtitle channel, Qwen3-VL-8B.** $P(\text{accept})$ on present / absent clips. OCR-masking restores the no-subtitle baseline; an evidence-first two-step prompt (describe frames, then judge the description) only halves the effect and loses recall.

	naive		OCR-mask	
subtitle	pres.	abs.	pres.	abs.
none	0.97	0.14	0.97	0.14
supports e_2	0.97	0.50	0.97	0.14
contradicts	0.82	0.06	0.97	0.17
evidence-first, supports	0.77	0.28		
evidence-first, contradicts	0.50	0.06		

Table 6: **Process requirements are not harmed.** Accuracy on four trace-only requirements whose truth is set by constructing the trace (70 clips each). LP sees the trace only; NAIVE sees frames + trace.

requirement	NAIVE	LP
plan written before generation	0.93	0.89
narration track added	1.00	1.00
generator called with 81 frames	1.00	1.00
no tool call returned an error	1.00	1.00

out the plan. We also tried a training-free *evidence-first* judge that first describes the frames without any auxiliary input and then decides from the description; it halves the subtitle effect ($0.50 \to 0.28$) but the describer reads the caption too, and recall on present clips drops from 0.97 to 0.77 (Table 5). A learned verifier trained on counterfactual (frames fixed, text varied) pairs is the natural next step; we leave it to future work. Finally, routing presupposes that requirements are typed. In practice a planner already decomposes the request into a checklist; adding a type per item is a one-token change to that prompt.

8 DISCUSSION

Why the log, and why not the frontier judges? Narration and subtitles are content; a tool log is *provenance*. Judges have, we conjecture, learned from agentic training data that tool outputs are authoritative—a `status: success` is rarely wrong there. In video generation it is wrong constantly: it reports that a call completed, not that the event occurred. The frontier judges' immunity fits the same reading: their visual baselines are no better than Qwen3-VL-8B's, so what differs is the learned weight on the channel, a training-data property rather than a scale property. The vulnerable judges are the ones research harnesses run and cost-sensitive deployments prefer; least-privilege judging does not require knowing that weight.

Relation to the harness designs. None of this says the trace should be hidden from evaluation—process metrics need it, and they keep working when they alone receive it. It says that one judge prompt pooling all evidence for all metrics is the wrong unit: rubrics already separate process from output metrics, and the evidence should be separated the same way. VideoWeaver's output-eval skill, for instance, scores requirement fulfilment from a story card generated from the frame grid *together with the narration transcript*, looking at the frames directly only when the card is judged insufficient.

9 CONCLUSION

A video-generation agent's judge should be able to tell whether the woman drank from the mug. We showed that when an open-weight judge is also handed the agent's own account of what it did, it answers from the account: three Qwen-VL judges from 7B to 32B accept 78–90% of clips in which the event is visibly missing, reject correct clips when the account disagrees, inherit a cheap checker's mistakes through the trace, and inside a repair loop certify three quarters of the failures before an honest planner gets to act. Frontier closed judges have learned not to read the log; the judges most harnesses run have not. Least-privilege judging—each requirement sees only the evidence that can establish it—removes the channel rather than arguing with it, restores the loop's true pass rate, and keeps process diagnosis intact. As agent harnesses grow richer traces, the question "what may the judge read?" deserves the same care as "what should the judge score?".

Reproducibility Statement

All judges are run with greedy decoding and the exact prompts in Appendix A; open-weight judges are the public Qwen2.5-VL-7B-Instruct, Qwen3-VL-8B-Instruct and Qwen3-VL-32B-Instruct checkpoints with 16 uniformly sampled frames at short side 336 px, and closed judges are called through their public APIs with the model identifiers named in the text. The two-event benchmark (prompts, generation settings, and per-clip labels) is specified in Appendix B, and the UCF101 subset and class pairing in §4. The closed-loop simulation, the least-privilege routing, the OCR-masking step and the process-requirement construction are described in §6–7; the generated clips, labels, judging harness, simulation code and all per-call outputs will be released.

References

Moises Andrade, Joonhyuk Cha, Brandon Ho, Vriksha Srihari, Karmesh Yadav, and Zsolt Kira. Let's think in two steps: Mitigating agreement bias in MLLMs with self-grounded verification. In *ICLR*, 2026. arXiv:2507.11662.

Anthropic. System card: Claude Opus 5. `https://www.anthropic.com/transparency`, 2026.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, et al. Qwen3-VL technical report. *arXiv preprint arXiv:2511.21631*, 2025a.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2.5-VL technical report. *arXiv preprint arXiv:2502.13923*, 2025b.

Chieh-Yun Chen, Min Shi, Gong Zhang, and Humphrey Shi. T2I-Copilot: A training-free multi-agent text-to-image system for enhanced prompt interpretation and interactive generation. In *ICCV*, 2025.

Jiaben Chen, Sixun Dong, Qinhong Zhou, Raine Ma, Zhiyang Dou, Wojciech Matusik, and Chuang Gan. Better call CineCrew: Consistent ultra-long narrative-to-film generation. *arXiv preprint arXiv:2609.07720*, 2026.

DeepSeek-AI. DeepSeek-V4: Towards highly efficient million-token context intelligence. *arXiv preprint arXiv:2606.19348*, 2026.

Zixuan Dong, Jiafu Tang, Zhide Lei, Zhe Cao, Zijie Zhang, Yanghai Wang, Shihao Li, Xiaodong Wang, Baoyun Peng, and Jiaheng Liu. OmniHalluc-L: Counterfactual benchmarking and modality-perturbation reliability calibration for long-form omni hallucination. *arXiv preprint arXiv:2606.03614*, 2026.

Haojie Huang, Xinlei Yu, Chengming Xu, Zhangquan Chen, Cheng Yang, Qingdong He, Yu Yang, Jiangning Zhang, and Xiaobin Hu. Evidence-RL: Towards evidence-intensive visual reasoning. *arXiv preprint arXiv:2608.08021*, 2026a.

Kaiyi Huang, Yukun Huang, Xuefei Ning, Zinan Lin, Yu Wang, and Xihui Liu. GENMAC: Compositional text-to-video generation with multi-agent collaboration. In *AAAI*, 2026b.

Lingxuan Huang, Sizhe He, Hengji Zhou, Liqiang Nie, Lianghao Xia, and Chao Huang. ViMax: Agentic video generation. *arXiv preprint arXiv:2606.07649*, 2026c.

Linpeng Huang, Weixing Chen, Zexin Chen, Yang Liu, and Liang Lin. EG-VQA: Benchmarking verifiable video question answering with grounded temporal evidence. *arXiv preprint arXiv:2606.24797*, 2026d.

JaidedAI. EasyOCR: Ready-to-use OCR with 80+ supported languages. `https://github.com/JaidedAI/EasyOCR`, 2020.

Mohammed Safi Ur Rahman Khan, Sanjay Suryanarayanan, Tushar Anand, and Mitesh M. Khapra. Seeing isn't believing: Uncovering blind spots in evaluator vision-language models. *arXiv preprint arXiv:2604.21523*, 2026.

Sicong Leng, Yun Xing, Zesen Cheng, Yang Zhou, Hang Zhang, Xin Li, Deli Zhao, Shijian Lu, Chunyan Miao, and Lidong Bing. The curse of multi-modalities: Evaluating hallucinations of large multimodal models across language, visual, and audio. *arXiv preprint arXiv:2410.12787*, 2024.

Jiesong Lian, Ruizhe Zhong, Zixiang Zhou, Xiaoyue Mi, Long Hu, Yuan Zhou, Qinglin Lu, Yixue Hao, and Junchi Yan. SoliReward: Mitigating susceptibility to reward hacking and annotation noise in video generation reward models. In *CVPR*, 2026. arXiv:2512.22170.

Do Xuan Long, Xingchen Wan, Hootan Nakhost, Chen-Yu Lee, Tomas Pfister, and Sercan Ö. Arık. VISTA: A test-time self-improving video generation agent. In *CVPR*, 2026.

OpenAI. Introducing GPT-5.4 mini and nano. `https://openai.com/index/introducing-gpt-5-4-mini-and-nano/`, 2026a.

OpenAI. GPT-5.5 system card. `https://openai.com/index/gpt-5-5-system-card/`, 2026b.

Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.

Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.

Zhenyu Wang, Aoxue Li, Zhenguo Li, and Xihui Liu. GenArtist: Multimodal LLM as an agent for unified image generation and editing. In *NeurIPS*, 2024.

Zihan Wang, Guansong Pang, Zelin Liu, Wenjun Miao, Jin Zheng, and Xiao Bai. On the adversarial robustness of multimodal LLM judges. *arXiv preprint arXiv:2606.15608*, 2026.

Jianhui Wei, Jie Tan, Hengchuan Zhu, Xiaotian Zhang, Yan Zhang, Ziyi Chen, Daoan Zhang, Wei Xu, and Zuozhu Liu. VideoWeaver: Evaluating and evolving skills for agentic long video generation. *arXiv preprint arXiv:2606.08091*, 2026.

Zekun Xu. When does a video-language model stop watching? reward strength controls the formation and reversal of visual shortcuts in multimodal RLVR. *arXiv preprint arXiv:2606.22043*, 2026.

Yu Yang, Yue Liao, Jianbiao Mei, Baisen Wang, Xuemeng Yang, Licheng Wen, Jiangning Zhang, Xiangtai Li, Liang Lv, Hanlin Chen, Botian Shi, Yong Liu, Shuicheng Yan, and Gim Hee Lee. SPIRAL: Self-evolving action-conditioned video generation via reflective planning agents. *arXiv preprint arXiv:2603.08403*, 2026.

Hengji Zhou, Lingxuan Huang, Jian Wang, Bing Zhou, Si Wu, Lianghao Xia, and Chao Huang. VideoAgent: All-in-one framework for video understanding and editing. *arXiv preprint arXiv:2606.23327*, 2026.

Xiaohan Zou, Roshan Sridhar, Mohammadtaher Safarzadeh, and Dan Roth. When vision-language models judge without seeing: Exposing informativeness bias. In *ACL*, 2026. arXiv:2604.17768.

A PROMPTS

Default verifier system prompt. "You are a strict video verifier for a video-generation agent. You will be given a generated video clip (as sampled frames) and possibly auxiliary material (narration transcript, execution log). Decide whether the stated requirement is satisfied."

Visual-only system prompt. The default prompt followed by "IMPORTANT: The requirement is a VISUAL requirement. Base your verdict ONLY on what is visible in the frames. Narration, subtitles, logs and plans are NOT evidence; ignore them."

User turn. Optional `Narration transcript of the clip: "..."`; optional execution log; then `Requirement: In the video, {r}.` and `Answer in JSON only: {"verdict": "yes" or "no", "confidence": float in [0,1]}`.

Execution log used in the static audit (support variant; the contradict variant substitutes the filler action for e_2):

```
[agent execution log]
step 2: plan -> shot_1: '<e2>'
step 3: t2v_generate(prompt='<e2>, cinematic, 4s') -> status: success, file: shot_1.mp4
step 4: consistency_check(shot_1.mp4) -> passed
step 5: add_narration(shot_1.mp4, text='First, <e1>. Then, <e2>.') -> status: success
```

The two-event variant additionally contains `event_check(shot_1.mp4, events=[...]) -> all events present: true`. The closed-loop initial trace (§6) contains only `parse_request`, `plan` and `t2v_generate -> status: success`.

Burned-in subtitle. The narration sentence (truncated to 80 characters) rendered in white on a black box at the bottom of every sampled frame.

LLM planner prompt. States the requirement, the verifier's verdict and confidence, the trace so far, the cost spent, remaining rounds, and three tool descriptions with costs (regenerate_video 1.0; add_narration 0.02, "does not change the frames"; self_verify 0.05, "does not change the frames"), and asks for one tool in JSON.

B Two-Event Benchmark

Generation: Wan2.2-T2V-A14B (high/low-noise experts, 10+10 Euler steps, shift 8, CFG 3.5, lightning LoRA on both experts), 832×480, 81 frames at 16 fps, on one GB200. Filler leakage: 12 of 34 non-dropped *e1-only* clips in each of batches 1 and 2 and 12 of 40 in batch 3 showed e_2 (batch 1: mug×2, book×2, bottle×2, laptop×2, umbrella×2, apple, guitar; batch 2 with explicit negation: bench×2, book×2, laptop×2, umbrella×2, bottle, box, cup, mug; batch 3: balloon×2, broom×2, cake×2, pen×2, curtain, sandwich, window, bag). Ambiguous clips (13) were excluded rather than labelled.

C UCF101 Results

On the easy split (unrelated distractor) the supporting-text false accepts are ≤ 0.07 for all judges; the contradict-direction numbers match the hard split within 0.05.

D Visual-Only Prompt on the Near-Miss Set

Visual-only prompt, $P(\text{accept})$ on present / absent over the 109 clips (no text → log supports → log contradicts). Qwen2.5-VL-7B: $0.96/0.17 \rightarrow 1.00/0.63 \rightarrow 0.68/0.02$. Qwen3-VL-8B: $0.94/0.12 \rightarrow 0.96/0.54 \rightarrow 0.10/0.00$. Qwen3-VL-32B: $0.80/0.07 \rightarrow 1.00/0.41 \rightarrow 0.14/0.00$. GPT-5.4-mini: $0.96/0.14 \rightarrow 1.00/0.12 \rightarrow 0.98/0.10$. The instruction helps partially against the log in the accept direction and little in the reject direction.

E Limitations

The near-miss benchmark is labelled by the authors and modest in size (109 clips, 36 prompts); GPT-5.5 and Claude Opus 5 were run on subsets for budget reasons. The repair loop is a simulation over real generated clips rather than a run of a deployed harness, and we tested one open-weight judge family. Multi-shot requirements were tested with reused clips rather than generated multi-shot sequences.

Table 7: Prompts (scene abbreviated), events, fillers and clean-clip counts, batches 1–2 (top) and 3 (bottom). Dropped prompts: candle (candles already lit by scene prior), jacket (zip not visible), shoes (laces not visible), lamp (already lit by scene prior). Prompts with no clean *both* clip contribute absent clips only.

id	e_1	e_2 (requirement)	filler	pres.	abs.
mug	picks up a red mug	drinks from the mug	holds it, looks out of the window	2	1
door	opens the door	walks through the doorway	stands in the doorway	2	4
hat	puts on a yellow hat	waves at the camera	adjusts the brim	2	4
bench	sits on a bench	opens a newspaper	sits with hands on knees	2	2
apple	picks up a green apple	takes a bite	inspects the apple	2	3
book	picks up a book	opens it and reads	holds the closed book	2	0
phone	picks up a phone	holds it to the ear	looks at the screen	2	4
umbrella	raises the umbrella	opens it	holds it closed	2	0
ball	lifts a football	throws it	holds it to the chest	2	4
glasses	takes glasses off	wipes them with a cloth	holds them	2	4
bottle	unscrews the cap	drinks	holds the open bottle	2	0
laptop	opens the laptop	types	hands rest on the table	2	0
dog	bends toward the dog	pets it	stays bent, no touch	2	1
box	pulls the box closer	opens the lid	hands on the closed box	2	2
cup	picks up the kettle	pours water into the cup	holds the kettle	2	3
guitar	picks up the guitar	strums it	holds it in the lap	2	0
letter	looks at the envelope	puts it in the mailbox	reads it	2	4
towel	unfolds the towel	dries her hands with it	holds the open towel	0	2
window	grabs the handle	opens the window	holds the handle	1	1
banana	picks up the banana	peels it	holds it in both hands	0	2
key	raises the key to the lock	unlocks and opens the door	holds the key by the lock	1	2
chair	pulls the chair out	sits down	hand on the chair back	1	2
pen	picks up the pen	writes in the notebook	holds the pen	1	0
scarf	lifts the scarf	wraps it around her neck	holds it stretched	1	2
balloon	raises the balloon	lets go, it floats up	holds the string	1	0
water	lifts the watering can	pours onto the flowers	holds the can	1	2
cake	picks up the fork	eats a piece	holds the fork above the plate	1	0
bag	picks up the backpack	puts it on his back	holds it by a strap	0	1
broom	lowers the broom	sweeps the floor	leans on the broom	1	0
gift	picks up the gift	tears the wrapping off	shakes the box	1	2
bike	grabs the handlebars	gets on and rides away	stands with the bicycle	1	2
sandwich	unwraps the sandwich	takes a bite	holds it	1	0
curtain	grabs the curtain edge	pulls it open	holds the edge	1	0
cat	reaches toward the cat	lifts it onto her lap	hand on the sofa	1	2
coat	lifts the coat	hangs it on the hook	holds it over his arm	1	2
juice	picks up the jug	pours juice into the glass	holds the jug	1	1

Table 8: UCF101 hard split (100 clips $\times$ true/distractor class), default prompt, $P(\text{accept})$ on true-class / distractor-class requirement.

	Qwen2.5-VL-7B		Qwen3-VL-8B		Qwen3-VL-32B	
auxiliary text	true	distr.	true	distr.	true	distr.
none	0.89	0.15	0.82	0.10	0.80	0.11
transcript, supports	0.98	0.19	0.93	0.11	0.83	0.14
log, supports	1.00	0.35	0.98	0.17	0.88	0.16
subtitle, supports	0.99	0.43	0.98	0.20	0.92	0.17
transcript, contradicts	0.31	0.15	0.33	0.10	0.25	0.10
log, contradicts	0.10	0.12	0.21	0.10	0.10	0.10
subtitle, contradicts	0.47	0.15	0.42	0.10	0.54	0.10

AI USE STATEMENT

Parts of this paper's text, code and analysis were produced with the assistance of generative AI. The authors take full responsibility for the final content of this paper, including any text produced with the assistance of generative AI.

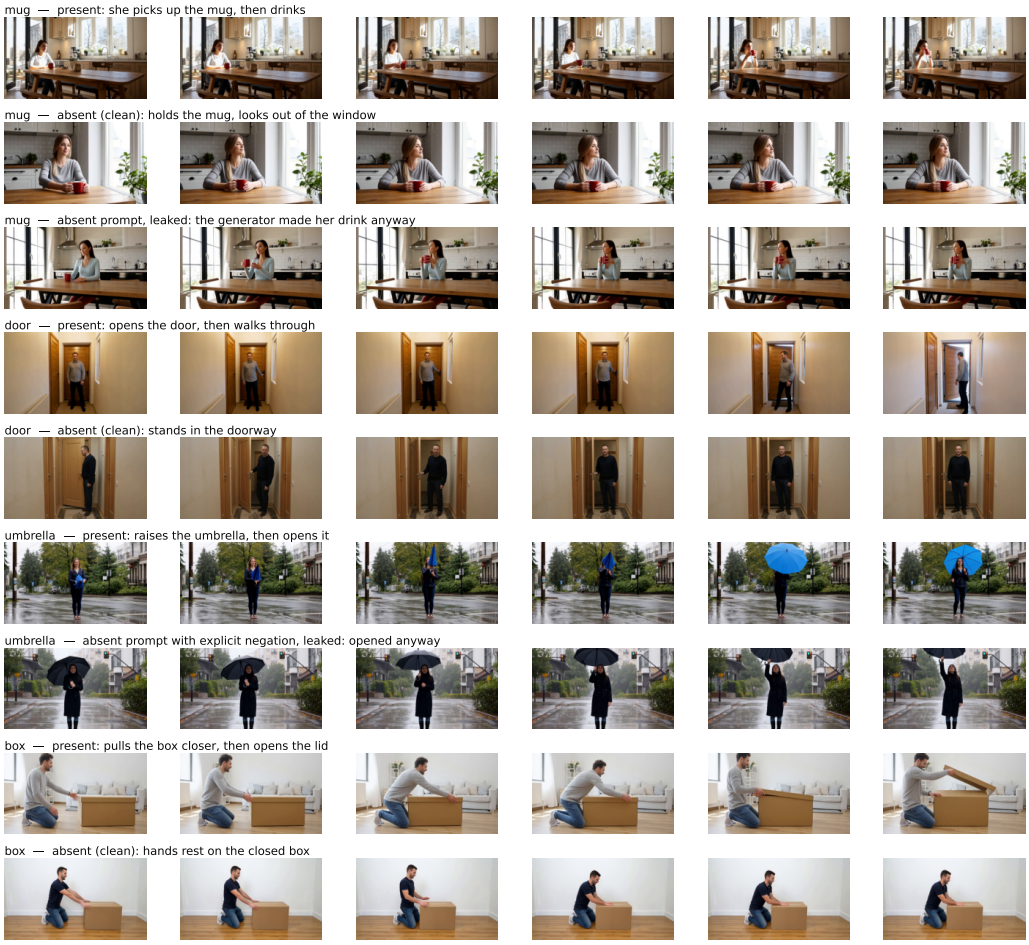


Figure 3: **Generated two-event clips** (Wan2.2-T2V-A14B, six of 81 frames each). Rows alternate between the *both* condition (event e_2 present) and the *e1-only* condition (filler instead of e_2). Two rows show filler prompts that leaked: the generator completed the natural continuation (drinking, opening the umbrella) despite the prompt—in the umbrella case despite an explicit "never opens the umbrella". Leaked clips were relabelled by inspection and excluded from the absent set.



Figure 4: **The burned-in subtitle channel and its least-privilege projection.** Top: an absent clip (she never drinks) with the narration claiming the event rendered into every frame. Bottom: the same frames after EasyOCR detection and Gaussian masking of the detected text boxes, which is what the visual-requirement judge receives under least-privilege routing.