HALLUCLEAR: A DIAGNOSTIC AND EVALUATION SUITE WITH MITIGATION PROBE FOR HALLUCINATIONS IN GUI AGENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

As GUI agents increasingly benefit from industrial-scale training, ungrounded hallucinations can still lead to cascading failures during sequential GUI interaction. Unlike hallucinations in general vision-language tasks, GUI-agent hallucinations span both visual perception and task reasoning and are further complicated by partial observability. Existing benchmarks primarily assess task success, visual grounding, or action fidelity, but lack fine-grained hallucination diagnosis and reliable evaluation. To bridge this gap, we introduce **HalluClear**, a suite for diagnosing and evaluating hallucinations in GUI agents, with a mitigation probe. HalluClear consists of three components: (1) a diagnostic framework that formalizes GUI-agent hallucination as unsupported belief-action commitment under partial observability and derives a fine-grained taxonomy from empirical failure analysis; (2) a three-stage evaluation workflow that combines an expert-annotated benchmark with qualified VLM judges for credibility-calibrated hallucination measurement; and (3) a lightweight mitigation probe that instantiates a structured post-training intervention motivated by patterns identified through our diagnostic analysis. Experiments with representative GUI agents on public benchmarks reveal systematic hallucination patterns across agent families and failure subtypes, providing diagnostic information complementary to conventional visual grounding and action-fidelity metrics. Across three open-source GUI agents, a post-training intervention using 9K samples consistently reduces hallucination rates while improving visual grounding and action fidelity.

1 INTRODUCTION

Recent advances in Vision-Language Models (VLMs) (Qwen Team, 2025a;b), supported by teacher-guided distillation (OpenAI, 2024; Google DeepMind, 2025), large-scale offline corpora (Wu et al., 2025; Wang et al., 2025b), and reinforcement-learning-inspired self-evolving schemes (Guo et al., 2025; Qin et al., 2025; Ye et al., 2025; Schulman et al., 2017; Shao et al., 2024; Yu et al., 2025), have substantially accelerated the development of visual GUI agents. In particular, industrial efforts have largely focused on expanding data and computation to push the performance frontier, as illustrated by the green arrow in fig. 1. However, during sequential GUI interaction, agents may still make unsupported commitments about visual states, task-progress inferences, or action targets, allowing step-level hallucinations to propagate into downstream task failures.

Existing GUI-agent benchmarks (Wu et al., 2025; Li et al., 2024; Lu et al., 2025a; Li et al., 2025; Cheng et al., 2024; Lin et al., 2025) primarily evaluate task completion, UI-element localization, or action validity. While these metrics are essential, they do not directly answer whether an agent hallucinated, where the hallucination occurred, or which failure mechanism caused it. Therefore, hallucination evaluation requires treating GUI-agent hallucination as a distinct evaluation object beyond conventional metrics of task success, visual grounding, and action fidelity.

Although hallucination evaluation has been extensively studied in general VLM settings (Li et al., 2023c; Guan et al., 2024; Bang et al., 2025; Huang et al., 2025), GUI-agent hallucinations pose additional challenges. Unlike static, perception-centric tasks, GUI agents operate in partially observable and dynamically evolving environments, where decisions depend on the current screenshot, task

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.

instructions, interaction history, inferred task states, and executable action constraints. As a result, directly adapting existing hallucination suites does not adequately distinguish genuine hallucinations from uncertainty or non-hallucinatory errors arising from partial observability, motivating evaluation protocols tailored to GUI interaction.

To bridge this gap, we introduce **HalluClear**, a diagnostic and evaluation suite with mitigation probe for hallucinations in GUI agents. As shown in fig. 2, HalluClear consists of three core components.

First, we develop a fine-grained diagnostic framework for GUI agents. Under partially observable GUI interaction, we formalize hallucination as unsupported belief-action commitment. Building on related hallucination taxonomies, we conduct bottom-up inductive clustering and abstraction over empirical failure cases of GUI agents on open-source datasets. This process yields a fine-grained taxonomy that separates hallucinations into perception and reasoning dimensions, with eight subtypes capturing unsupported commitments about visual states, task context, and reasoning.

Second, based on this diagnostic framework, we design a three-stage hallucination evaluation workflow. Human experts first construct an expert-annotated benchmark. The benchmark is then used to qualify and calibrate VLM judges based on their agreement with expert annotations under the fine-grained taxonomy. Only qualified judges participate in large-scale evaluation, where credibility-weighted aggregation is used to estimate hallucination rates.

Finally, motivated by the hallucination patterns identified through diagnostic analysis, we introduce a lightweight mitigation probe. We instantiate a structured reasoning intervention as a controlled post-training probe to examine its effects on hallucination and downstream GUI-agent performance.

Experiments with representative GUI agents on public benchmarks show that HalluClear reveals systematic hallucination patterns across agent families and failure subtypes, providing fine-grained diagnostic information complementary to conventional visual-grounding and action-fidelity metrics. Across three open-source GUI agents, a post-training intervention using 9K training samples consistently reduces hallucination rates and improves visual grounding and action fidelity, as illustrated by the blue arrow in fig. 1.

In summary, our main contributions are as follows:

- We develop a diagnostic framework that formalizes GUI-agent hallucination as unsupported belief-action commitment under partial observability and derives a GUI-specific taxonomy with eight subtypes from systematic failure analysis.
- We design a three-stage credibility-calibrated evaluation workflow that combines an expert-annotated benchmark with qualified VLM judges for reliable subtype-level hallucination measurement, revealing systematic patterns across GUI agents and failure subtypes.
- We introduce a lightweight mitigation probe motivated by these identified patterns, instantiating a structured post-training intervention that reduces hallucination rates and improves visual grounding and action fidelity across three open-source GUI agents.

2 HALLUCINATION DIAGNOSTIC

GUI agents interact with the environment sequentially under partial observability. We model the underlying environment as a discrete-time Partially Observable Markov Decision Process (Kaelbling et al., 1998), where each user task $u \in \mathcal{U}$ induces a POMDP $\mathcal{M}_u = \langle \mathcal{S}, \mathcal{A}, \mathcal{O}, \mathbb{T}, \mathbb{O}, \mathcal{R} \rangle$, with state space $\mathcal{S}$, action space $\mathcal{A}$, observation space $\mathcal{O}$, transition kernel $\mathbb{T}(s_{t+1} \mid s_t, a_t)$, observation kernel $\mathbb{O}(o_t \mid s_t)$, and reward function $\mathcal{R}(\tau)$ defined over a trajectory τ.

Recent advances in VLMs (Qwen Team, 2024; 2025a;b) have driven the adoption of end-to-end GUI agents with explicit System-2-style reasoning (Wei et al., 2022; Yao et al., 2022), now a prominent paradigm for GUI interaction (Xu et al., 2025; Wu et al., 2025; Qin et al., 2025; Wang et al., 2025a; Ye et al., 2025; Wang et al., 2025b). Rather than maintaining explicit belief states, these agents operate on an information state $\tilde{s}_t = (o_t, u, h_t)$, where $h_t = (o_{1:t-1}, a_{1:t-1})$ denotes the interaction history. However, hallucination remains insufficiently characterized in GUI agents.

2.1 HALLUCINATION DEFINITION

Policy optimization aims to maximize expected returns in a POMDP. As illustrated in fig. 1, in GUI interaction, the attainable returns are constrained by partial observability and the limited information exposed through observations. We therefore define hallucination relative to the available *information state* $\tilde{s}_t$, distinguishing agent errors from the unavoidable performance gap induced by partial observability. Conceptually, we define GUI-agent hallucination as an *unsupported belief-action commitment* relative to the available information state $\tilde{s}_t$.

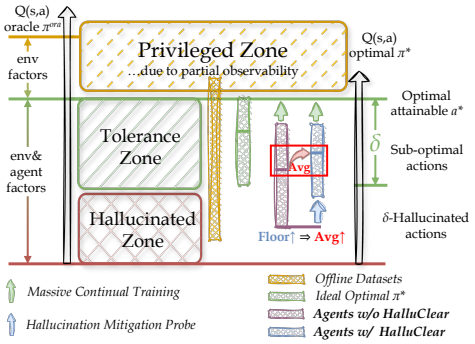

Figure 1: **Hallucination Definition in GUI Agents.** The **Privileged Zone** denotes oracle-level actions requiring unavailable information; the **Tolerance Zone** contains evidence-supported actions near the attainable optimum; and the **Hallucinated Zone** consists of actions driven by unsupported commitments. HalluClear diagnoses and evaluates such failures, while the gain of mitigation probe is visualized by the red arrow.

Definition. *Given a policy π and information state $\tilde{s}_t \in \tilde{\mathcal{S}}$, an action a_t is δ-hallucinated if the policy prefers it over the best attainable action a^* despite a value gap exceeding δ*:

$$\begin{cases} \pi(a_t|\tilde{s}_t) > \pi(a^*|\tilde{s}_t) \\ Q^{\pi^*_{\tilde{s}}}(\tilde{s}_t, a^*) - Q^{\pi^*_{\tilde{s}}}(\tilde{s}_t, a_t) > \delta \end{cases} \tag{1}$$

where $\pi^*_{\tilde{s}} \in \text{argmax}_{\pi \in \Pi_{\tilde{s}}} \mathbb{E}_{\tau(\pi, \mathcal{M}_u)}[\mathcal{R}(\tau)]$ is optimal within $\Pi_{\tilde{s}}$, $a^* \in \text{argmax}_a Q^{\pi^*_{\tilde{s}}}(\tilde{s}_t, a)$ is the best attainable action, and $\delta > 0$ is the tolerance. By construction, $\pi^*_{\tilde{s}}$ serves as the attainable reference policy under eq. (1).

Notice that $\pi^*_{\tilde{s}}$ differs from the "oracle" policy $\pi^{ora}_{\tilde{s}^p} \in \Pi_{\tilde{s}^p}$, implied by ground truth in ideal error-free offline datasets with access to extended **p**rivileged state $\tilde{s}^p \in \tilde{\mathcal{S}}^p$ (e.g., DOM and hindsight replay). Now, under the surjective mapping $\phi : \tilde{\mathcal{S}}^p \to \tilde{\mathcal{S}}$ and its kernel equivalence relation, we have:

$$\Pi_{\tilde{s}} \cong \{\pi \in \Pi_{\tilde{s}^p} \mid \forall \tilde{s}^p_1 \sim_\phi \tilde{s}^p_2,\ \pi(\cdot \mid \tilde{s}^p_1) = \pi(\cdot \mid \tilde{s}^p_2)\} \tag{2}$$

eq. (2) represents the *measurability constraint*. Consequently, we inherently have:

$$\max_{\pi \in \Pi_{\tilde{s}}} \mathbb{E}_{\tau(\pi, \mathcal{M}_u)}[\mathcal{R}(\tau)] \leq \mathbb{E}_{\tau(\pi^{ora}_{\tilde{s}^p}, \mathcal{M}^p_u)}[\mathcal{R}(\tau)] = J_{\text{ora}} \tag{3}$$

A source of strict inequality is *perceptual aliasing*, where states requiring different optimal actions induce the same information state. In such cases, J_{ora} is unattainable under partial observability.

Our definition evaluates **an agent against the best achievable policy $\pi^*_{\tilde{s}}$ given only $\tilde{s}$, rather than penalizing it for lacking the oracle's privileged access**. Since $\pi^*_{\tilde{s}}$ is intractable, we use expert judgments conditioned on the same observable information $\tilde{s}$ as a practical evaluation proxy.

2.2 Hallucination Taxonomy

To operationalize GUI-agent hallucination, we analyze failures of advanced GUI agents and general-purpose VLMs (Qwen Team, 2025a; Qin et al., 2025; Ye et al., 2025) on offline GUI datasets (Li et al., 2024; Lu et al., 2025a). Following our definition of hallucination as unsupported belief-action commitment under partial observability, we derive eight recurring subtypes grouped into two dimensions: perception hallucination, arising from unsupported visual interpretation, and reasoning hallucination, arising from unsupported textual, contextual, or logical inference.

Perception Hallucination (PH). PH refers to hallucinations rooted in visual perception and grounding. In these cases, the agent makes unsupported commitments about the current GUI observation, UI elements, or their visual and semantic relations.

- **PH.1 Screenshot State**: The agent misinterprets the global semantic state of the screenshot, i.e., the observation o_t, beyond individual UI elements.
- **PH.2 Element Existence**: The agent commits to the presence of a UI element that is not supported by the screenshot.
- **PH.3 Element Attribute**: The agent misidentifies intrinsic properties of a UI element, including its appearance, function, state, or interaction affordance.

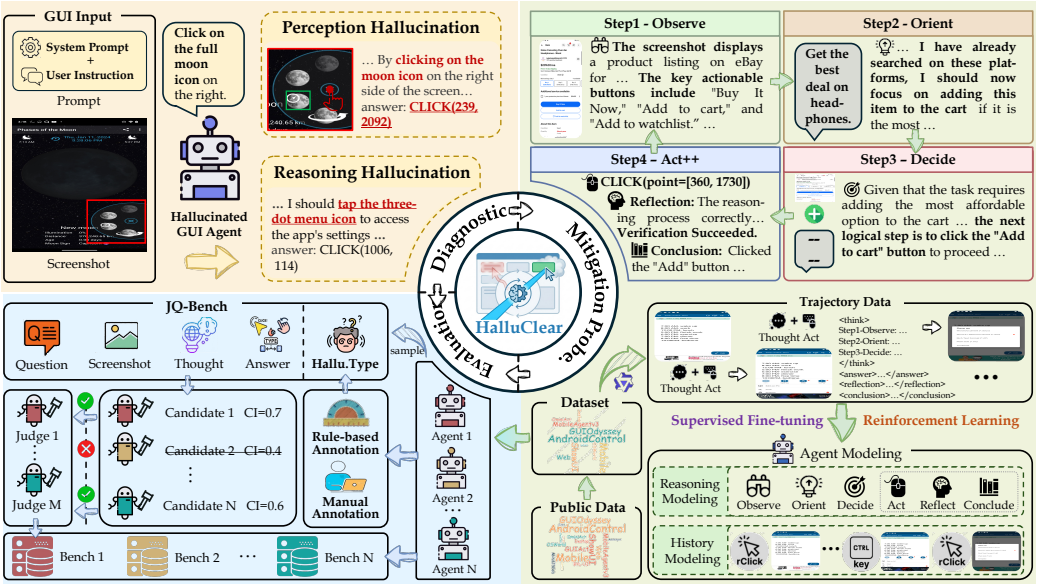

Figure 2: **Overview of the HalluClear Suite.** HalluClear is a diagnostic and evaluation suite with mitigation probes for hallucinations in GUI agents. (1) The fine-grained **diagnostic framework** (top left) formalizes GUI-agent hallucinations and characterizes their subtypes through bottom-up clustering and abstraction of offline failure cases. (2) The three-stage **evaluation workflow** (bottom left) builds the expert-annotated JQ-Bench, qualifies VLM judges via credibility estimation, and reports credibility-calibrated hallucination measurements. (3) The **mitigation probe** (right) instantiates diagnosis-guided structured reasoning as post-training interventions to test whether diagnostic signals can guide targeted intervention and improve visual grounding and action fidelity.

- **PH.4 Element Relation**: The agent misinterprets spatial or semantic relationships between UI elements, or between an element and the global observation o_t.

Reasoning Hallucination (RH). RH refers to hallucinations rooted in textual, contextual, or logical reasoning. In these cases, the agent makes unsupported commitments about the instruction, interaction history, task progress, causal reasoning, or external facts.

- **RH.1 Instruction**: The agent forms an unsupported interpretation of the task instruction or step-specific constraints, leading it to omit, distort, or contradict them.
- **RH.2 Context**: The agent introduces contextual inconsistencies, such as contradictions between the available action space, interaction history, reasoning trace, and final action.
- **RH.3 Logic**: The agent relies on unsupported or invalid inference links when deriving the next action, including unjustified logical or causal transitions.
- **RH.4 Fact**: The agent fabricates or relies on unsupported factual information, including domain-general external knowledge not grounded in the current GUI state or task context.

No Hallucination (NonH). We additionally define a *NonH* category for cases without unsupported commitments relative to the available GUI state and task context. This category includes correct executions, evidence-supported alternatives to the reference action, and benign deviations consistent with the observable information. It prevents ordinary errors or acceptable action variations from being conflated with hallucinations.

3 Hallucination Evaluation

3.1 Philosophy

Prior quantitative hallucination assessments typically rely on three paradigms: (1) **heuristic scoring** (Li et al., 2023c), which provides limited diagnostic coverage; (2) **model-based judging** (e.g., LLM-as-a-judge) (Liu et al., 2023), which may suffer from judge instability and systematic bias; and

(3) **reference-based benchmarks** (Lin et al., 2022; Li et al., 2023b), which are costly to scale and may have limited coverage. Inspired by crowdsourcing qualification tests, we combine model-based judging with reference-based benchmarks in a three-stage hallucination evaluation workflow.

Overall, we conduct **J**udge **Q**ualification (**JQ**) on the expert-annotated **JQ-Bench** to qualify candidate VLM judges and estimate their credibility. Qualified judges then assess GUI-agent responses on **HR-Pool** to estimate **H**allucination **R**ates (**HR**). JQ-Bench contains expert-annotated task-query–agent-response pairs, while HR-Pool is a shared set of challenging task queries across target agents. Although they share the same task queries, hallucination labels are response-conditioned: JQ-Bench annotations are used only for judge qualification and credibility estimation, whereas HR-Pool responses are newly generated by each target agent and judged independently. This design enables scalable automated hallucination evaluation while maintaining a traceable link to expert judgments.

3.2 IMPLEMENTATION

Stage 1: Construction of HR-Pool and JQ-Bench. We first curate **HR-Pool** from public GUI datasets and agent failure cases. We filter out queries with dataset artifacts, uninformative contexts, or irreducible action ambiguity to avoid conflating hallucination with annotation incompleteness or underspecified GUI states.

To construct **JQ-Bench**, we pair HR-Pool queries with responses from advanced GUI agents and general-purpose VLMs, and annotate the resulting query–response pairs with expert hallucination labels. The annotation follows a lookup-style protocol derived from concrete GUI-agent failure cases and detailed in Appendix E. Formally, JQ-Bench is denoted as $\{(q_i, a_i, gt_i)\}_{i=1}^{|\text{JQ-Bench}|}$, where q_i is a task query, a_i is an agent response, and gt_i is the set of acceptable expert labels from the taxonomy in section 2.2. Since hallucinations may overlap or admit ambiguous subtype attribution, each pair may have up to three acceptable labels.

Stage 2: Selection of VLM Judges. Given JQ-Bench, we select qualified VLM judges for large-scale hallucination-rate measurement. We follow a *strict-in, lenient-out* principle: candidate judges must reliably identify fine-grained hallucination subtypes on JQ-Bench before being admitted to the final panel for large-scale HR-Pool evaluation. This admission test encourages taxonomy-grounded understanding rather than shallow error detection. Each candidate judge follows the same taxonomy-grounded labeling protocol as human annotators and is evaluated against the expert annotations in JQ-Bench. For the k-th judge, we define its credibility as:

$$C_k = \frac{1}{|\text{JQ-Bench}|} \sum_{i=1}^{|\text{JQ-Bench}|} \mathbb{1}\{\text{VLM}_k(q_i, a_i) \in gt_i\}, \quad k \in \mathcal{P} \text{ if } C_k \geq \tau_{\text{JQ}}, \tag{4}$$

where $\mathcal{P}$ denotes the admitted judge panel and τ_{JQ} is the qualification threshold. Only judges meeting this threshold are used in Stage 3.

Stage 3: HR Evaluation of Agents. After qualification, each target GUI agent is evaluated on the same HR-Pool queries. For agent m and query $q_j \in$ HR-Pool, the agent generates a response $a_j^{(m)}$, and qualified judges independently classify the pair $(q_j, a_j^{(m)})$ into one of nine categories: eight hallucination subtypes or *NonH*. Thus, Stage-3 labels are response-conditioned and not inherited from JQ-Bench. For judge k, agent m, and category c, the category rate is:

$$R_{k,c}^{(m)} = \frac{1}{|\text{HR-Pool}|} \sum_{j=1}^{|\text{HR-Pool}|} \mathbb{1}\{\text{VLM}_k(q_j, a_j^{(m)}) = c\}. \tag{5}$$

The judge-specific *Overall HR* is obtained by summing the eight hallucination-subtype rates while excluding *NonH*. To reduce judge-specific bias, we compute a credibility-weighted ensemble estimate using normalized credibility scores:

$$\bar{R}_c^{(m)} = \sum_{k \in \mathcal{P}} w_k R_{k,c}^{(m)}, \quad w_k = \frac{C_k}{\sum_{k' \in \mathcal{P}} C_{k'}}. \tag{6}$$

Further details on the evaluation workflow, including its design rationale, implementation, and fairness considerations, are provided in Appendix D; the full labeling protocol is described in Appendix E.

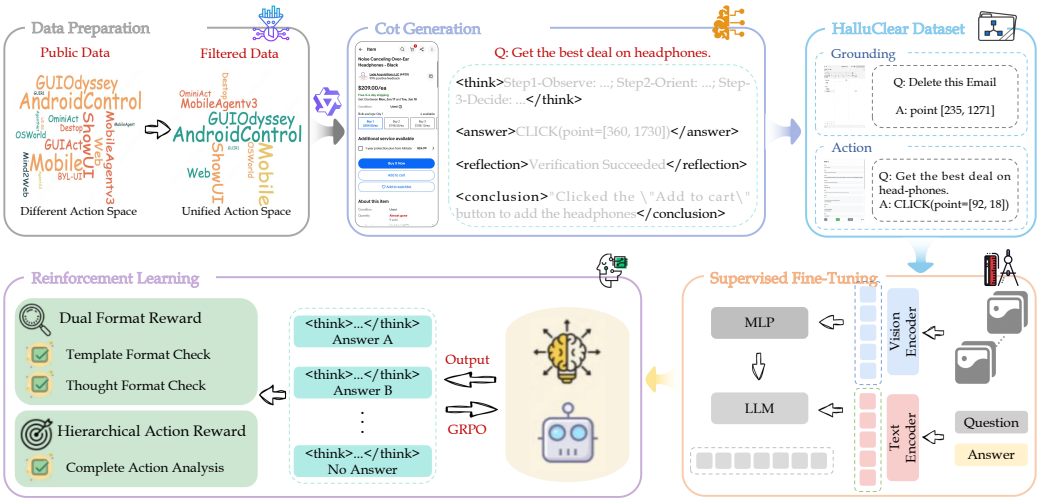

Figure 3: **Overview of the Mitigation-Probe Training Pipeline.** The probe consists of (1) **Cold Start**, involving mitigation dataset construction with structured reasoning traces and subsequent SFT; and (2) **Reinforcement Learning** via GRPO, utilizing Dual Format Reward and Hierarchical Action Reward to optimize perception accuracy and reasoning consistency.

4 Hallucination Mitigation Probe

4.1 OODA-Style Structured Reasoning

Motivated by the hallucination patterns identified through our diagnostic analysis, we instantiate a structured reasoning probe to examine whether lightweight post-training can reduce hallucinations and improve downstream GUI-agent performance. We use Observe-Orient-Decide-Act (OODA) as an interpretable scaffold to separate visual evidence collection, task-state reasoning, action selection, and post-action verification in GUI interaction. Concretely, we augment the base agent with a closed-loop reasoning format consisting of three functional modules:

- **Active Perception (Observe)** encourages the agent to explicitly extract visual evidence from the current GUI observation, helping reduce unsupported perceptual commitments.
- **Joint Reasoning (Orient & Decide)** separates task-environment alignment from action selection, encouraging the agent to ground reasoning in the context and observed state.
- **Introspective Execution (Act++)** adds action verification, reflection, and step-level summarization, allowing the agent to compare intended actions with current evidence and maintain compact trajectory memory.

Formally, at time step t, the agent receives a partially observable information state $\tilde{s}_t = (o_t, u, h_t)$, where $o_t \in \mathcal{O}$ is the current screen observation, $u \in \mathcal{U}$ is the global natural-language instruction, and h_t denotes the interaction history. The OODA-style probe defines a structured policy function:

$$\mathcal{F}(\tilde{s}_t) \rightarrow m_t = (\epsilon_t, a_t, r_t, c_t), \tag{7}$$

where ϵ_t is the reasoning trace produced by the Observe–Orient–Decide steps, a_t is the executable action, and r_t and c_t denote the reflection result and current-step summary produced in Act++. Each action $a_t \in \mathcal{A}$ is represented by an action type α_t (e.g., `CLICK`, `TYPE`) and parameters p_t (e.g., coordinates or text inputs). This structured output provides a controlled probe for evaluating the effects of structured post-training on hallucination and downstream GUI-agent performance.

4.2 Training the Mitigation Probe

To instantiate the mitigation probe, we perform lightweight two-stage post-training, as shown in fig. 3. We first use SFT to teach the structured reasoning format. Seed samples are drawn from ShowUI (Lin et al., 2025), AndroidControl (Li et al., 2024), and GUI-Odyssey (Lu et al., 2025a). We use Qwen3-VL-MAX (Qwen Team, 2025b) to generate structured reasoning trajectories and retain those passing our filtering criteria.

We then optimize the agent with GRPO (Shao et al., 2024) using a rule-based reward:

$$R_{\text{OODA}} = R_{\text{format}} + R_{\text{action}}. \tag{8}$$

The format reward encourages parseable and logically ordered responses. Let $\mathcal{S}_{\text{struct}}$ denote outputs satisfying the required "Thinking-Answer-Reflection-Conclusion" template, and $\mathcal{S}_{\text{logic}}$ outputs following the OODA-style reasoning order:

$$R_{\text{format}}(m_t) = \mathbb{1}(m_t \in \mathcal{S}_{\text{struct}}) + \mathbb{1}(m_t \in \mathcal{S}_{\text{logic}}). \tag{9}$$

The action reward encourages executable and accurate GUI actions. Let the predicted action be $a_t = (\alpha_t, p_t)$ and the ground-truth action be $a^g = (\alpha^g, p^g)$, where α denotes the action type and p its parameters. We use a gated reward:

$$R_{\text{action}}(a_t, a^g) = \mathbb{1}(\alpha_t = \alpha^g) \cdot \left(1 + w_{\alpha^g}\, \mathbb{1}(p_t \approx p^g)\right), \tag{10}$$

where $p_t \approx p^g$ denotes parameter consistency. We set $w_{\alpha^g} = 1.0$ for coordinate-sensitive actions such as `CLICK`, and $w_{\alpha^g} = 0.5$ for semantic-command actions such as `TYPE`.

5 Experiments

5.1 Experimental Setup

Benchmarks and evaluation pools. For hallucination evaluation, we build **JQ-Bench** and **HR-Pool** from public benchmarks (Li et al., 2024; Lu et al., 2025a). Both share 3,410 challenging task queries but serve distinct roles: HR-Pool is a shared query pool for large-scale evaluation, while JQ-Bench augments these queries with agent responses and expert hallucination labels for judge qualification and credibility estimation. For the mitigation probe, we further curate a 9K-sample set.

Agents and judges. We evaluate two closed-source agents, Gemini-3.1 and Claude-4.6, and three open-source agents: the general-purpose Qwen2.5-VL-7B (Qwen Team, 2025a) and two GUI-specialized agents, GUI-Owl-7B (Ye et al., 2025) and UI-TARS-1.5-7B (Qin et al., 2025). The mitigation probe is restricted to the three open-source agents. For judge qualification, we consider seven candidate VLM judges from diverse model families.

Metrics. For hallucination evaluation, we report overall HR and fine-grained subtype rates defined in section 3, using the qualified judges. To assess the downstream effects of the mitigation probe, we evaluate GUI-agent performance on standard benchmarks (Wu et al., 2025; Li et al., 2025; 2024; Lu et al., 2025a), reporting grounding accuracy and action success rate as complementary indicators.

5.2 Measurement Reliability

Reliable hallucination evaluation requires reproducible expert annotations and automated judges that are stable and aligned with human judgments. We assess reliability along three dimensions.

Human annotation reliability. To assess whether experts can consistently apply our taxonomy, we sample 200 outputs from each of three agents and ask three expert annotators to label them independently. As shown in Table 1, Fleiss' Kappa exceeds 0.66 for all agent subsets, with an average of 0.745, indicating substantial agreement and supporting the reproducibility of taxonomy-based annotation.

Judge consistency and screening. For each admitted judge, we run three independent evaluations on JQ-Bench and compute Fleiss' Kappa across runs. Table 2 shows that Gemini-3-Flash, Seed-1.8, and Qwen3-VL-Plus all achieve Kappa scores above 0.80, with an average of 0.855, indicating high self-consistency. We also evaluate all candidate judges under a binary hallucination-detection setting as a sanity check. As shown in Table 4, final admission follows the taxonomy-level JQ protocol in section 3, under which Gemini-3-Flash, Seed-1.8, and Qwen3-VL-Plus form the final judge panel.

Table 1: **Inter-Annotator Agreement on JQ-Bench.**

Model	FK
Qwen2.5-VL-7B	0.667
GUI-Owl-7B	0.777
UI-TARS-1.5-7B	0.791
Avg	**0.745**

Table 2: **Judge Consistency on JQ-Bench.**

Judge	FK
Gemini-3-Flash	0.842
Seed-1.8	0.913
Qwen3-VL-Plus	0.809
Avg	**0.855**

Table 3: **Judge-human Agreement on JQ-Bench among Different Agents.**

Model	Gemini	Seed	Qwen3
Qwen2.5-VL-7B	0.662	0.703	0.601
GUI-Owl-7B	0.768	0.671	0.618
UI-TARS-1.5-7B	0.823	0.751	0.777
Avg	**0.726**	**0.681**	**0.648**

Table 4: **Binary hallucination detection performance of candidate VLM judges on JQ-Bench.** F1 and accuracy averaged over **three** independent runs, formatted as $\mu^{+\Delta_{\max}}_{-\Delta_{\min}}$.

Metric	Gemini-3-Flash	Seed-1.8	Qwen3-VL-Plus	Gemini-2.5-Flash	Seed-1.6-Vision	GLM-4.6V	GPT-4o
F1	$92.6^{+0.4}_{-0.2}$	$80.9^{+0.2}_{-0.1}$	$82.1^{+0.4}_{-0.6}$	$81.1^{+0.1}_{-0.1}$	$81.9^{+0.3}_{-0.3}$	$80.9^{+0.2}_{-0.1}$	$81.2^{+0.4}_{-0.3}$
Acc	$89.8^{+0.5}_{-0.3}$	$78.7^{+0.1}_{-0.1}$	$76.0^{+0.5}_{-0.8}$	$73.7^{+0.2}_{-0.1}$	$75.1^{+0.4}_{-0.5}$	$75.7^{+0.3}_{-0.2}$	$74.5^{+0.6}_{-0.5}$

Table 5: **Agent-level hallucination profiles on HR-Pool.** Results are evaluated using different judges with their credibility scores.

Model	HR	Judge (HR ↓)		
		Gemini-3 (0.814)	Seed-1.8 (0.688)	Qwen3-VL (0.600)
Gemini-3.1-Flash	Overall	21.47	24.42	29.16
	PH	9.80	14.45	16.68
	RH	11.67	9.97	12.48
Claude-4.6-Sonnet	Overall	22.46	26.61	27.21
	PH	11.29	15.54	15.99
	RH	11.17	11.07	11.22
Qwen2.5-VL-7B	Overall	57.54	56.55	56.70
	PH	18.98	23.92	23.54
	RH	38.56	32.63	33.16
GUI-Owl-7B	Overall	54.46	39.02	41.77
	PH	6.33	10.73	11.53
	RH	48.13	28.29	30.24
UI-TARS-1.5-7B	Overall	41.27	38.53	40.17
	PH	8.58	11.05	11.92
	RH	32.69	27.48	28.25

Figure 4: **Fine-grained hallucination distribution on HR-Pool.** (a) Perception and (b) reasoning hallucinations are further decomposed into specific categories. Notably, the y-axis is scaled by $\sqrt{y}$ for better visualization of low-frequency categories.

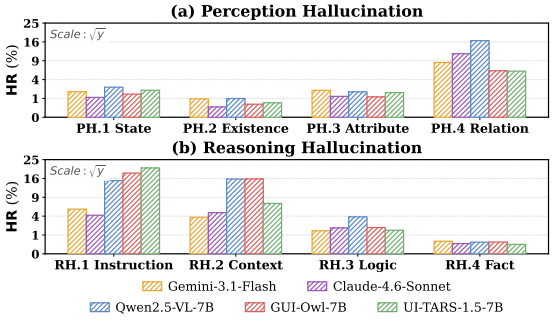

Judge-human agreement. We further compare qualified judges with expert annotations on sampled outputs from the three open-source agents. As shown in Table 3, most Cohen's Kappa scores fall within the 0.6–0.8 range, indicating substantial judge-human agreement. The average agreement trend also aligns with the judges' JQ credibility scores, supporting credibility as a reliability signal. Together, these results support the use of a qualified, credibility-weighted judge ensemble for scalable HR-Pool evaluation while maintaining a traceable link to expert judgments.

5.3 Diagnostic Findings Beyond Conventional Metrics

Conventional GUI-agent metrics indicate whether an action matches the reference, but not whether the decision is supported by the observable GUI state and task context. We therefore use HalluClear to compare hallucination profiles across agent families on the shared HR-Pool.

Agent-level hallucination profiles. Table 5 reports overall HR, PH-HR, and RH-HR under the qualified judges. All evaluated agents exhibit non-negligible hallucination vulnerability, while closed-source agents show substantially lower HR than open-source agents. The PH/RH decomposition further reveals distinct failure profiles: closed-source agents are relatively balanced between perception and reasoning hallucinations, GUI-specialized agents are dominated by reasoning hallucinations, and the general-purpose Qwen2.5-VL-7B shows stronger perception-side hallucinations. These results highlight systematic differences in hallucination profiles across agent families.

Fine-grained subtype distribution. Figure 4 presents fine-grained subtype rates. On the perception side, *PH.4 Element Relation* is the dominant failure type, highlighting persistent difficulty in

Table 6: **Actionability study through the mitigation probe.** We compare the baseline agent, a generic probe without HalluClear diagnosis, and the diagnosis-guided probe with HalluClear.

Model	Setup	Grounding (ACC ↑)		Action (SR ↑)				HR ↓
		SS-V2	SS-Pro	AC-Low	AC-High	GO-Low	GO-High	
Qwen2.5-VL-7B	Baseline	77.18	22.58	60.93	47.22	69.26	47.63	57.54
	Generic Post-training	83.85	18.91	80.50	63.02	71.74	51.36	35.81
	Structured Mitigation Probe	**86.62**	**31.31**	**87.49**	**66.72**	**77.71**	**53.99**	**24.98**
GUI-Owl-7B	Baseline	81.90	38.52	72.65	57.23	62.20	49.00	54.46
	Generic Post-training	84.88	40.87	79.03	66.20	76.11	58.58	30.90
	Structured Mitigation Probe	**88.67**	**43.58**	**87.46**	**68.19**	**79.58**	**58.82**	**21.28**
UI-TARS-1.5-7B	Baseline	86.55	43.83	66.86	53.96	75.88	63.22	41.27
	Generic Post-training	85.62	40.49	81.79	63.04	78.49	60.36	29.98
	Structured Mitigation Probe	**88.36**	**48.77**	**88.81**	**69.44**	**83.09**	**63.34**	**21.73**

modeling spatial and semantic relations among task-relevant UI elements. On the reasoning side, *RH.1 Instruction* and *RH.2 Context* dominate, especially for GUI-specialized agents, indicating persistent vulnerabilities in handling task constraints and interaction history.

5.4 Evaluation of the Mitigation Probe

We further evaluate whether a structured post-training intervention motivated by the hallucination patterns identified in our diagnostic analysis can reduce hallucinations. As shown in Table 6, the mitigation probe consistently reduces overall HR across all three open-source agents compared with their baselines and the generic post-training variant. These reductions are accompanied by improved grounding on ScreenSpot-V2/Pro and action performance on AndroidControl and GUI-Odyssey, showing that lower hallucination rates can coincide with improved downstream GUI-agent performance.

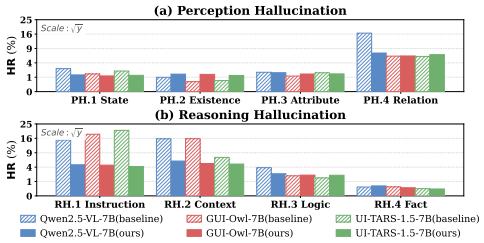

Figure 5: **Fine-grained hallucination reduction under the mitigation probe.**

Figure 5 further shows how subtype distributions change after post-training. The probe substantially reduces high-frequency *PH.4 Element Relation* errors in Qwen2.5-VL-7B and reasoning-side errors such as *RH.1 Instruction* and *RH.2 Context* in GUI-specialized agents. These results show that the mitigation probe reduces both aggregate HR and several high-frequency failure modes revealed by our diagnostic analysis.

5.5 Limitations

HalluClear still relies on qualified VLM judges and may inherit judge-specific biases in ambiguous cases. The current HR-Pool, built from AndroidControl and GUI-Odyssey, does not fully cover enterprise-style CUA workflows, cross-application desktop tasks, or long-horizon online interaction. The mitigation probe is limited to open-source agents and should be viewed as a controlled mitigation study rather than a universal mitigation algorithm. HalluClear also remains challenged by severe visual ambiguity, hidden states, and tasks requiring external factual knowledge beyond the observable context. Additional analyses of successful and failed cases are provided in Appendix B.3.

6 Conclusion

We presented HalluClear, a taxonomy-grounded suite for diagnosing and evaluating hallucinations in GUI agents. HalluClear treats GUI-agent hallucination as an evaluable construct: unsupported belief-action commitment under partial observability, rather than merely a mismatch with a reference action or final task failure. It operationalizes this construct through a fine-grained taxonomy, an expert-annotated JQ-Bench for judge qualification, and an HR-Pool for scalable hallucination-rate measurement. Across representative GUI agents, HalluClear exposes systematic hallucination profiles that conventional grounding and action metrics cannot fully explain. The mitigation probe further shows that these diagnoses can guide lightweight interventions, supporting the actionability of the evaluation signals. We hope HalluClear can serve as a reusable evaluation foundation for understanding, auditing, and improving the reliability of GUI agents.

AI USE STATEMENT

We used generative AI tools for language editing and phrasing assistance during manuscript preparation. All AI-assisted text was reviewed, revised, and verified by the authors, who take full responsibility for the final manuscript and its scientific claims. Generative AI tools were not treated as authoritative sources for experimental results or citations.

Generative models were also used as components of our research methodology, including VLM-based judging and the generation of structured reasoning trajectories. These methodological uses are explicitly described in the corresponding sections of the paper and are distinct from AI-assisted manuscript preparation.

ETHICS STATEMENT

This work studies hallucinations in GUI agents using existing public GUI benchmarks, model-generated responses, and expert annotations. Our experiments focus on offline evaluation and do not involve deployment with end users. GUI agents may have dual-use risks when deployed for automated interaction, particularly when hallucinated decisions lead to unintended actions. In addition, model-based judging may inherit biases from the underlying VLMs. We mitigate these concerns by explicitly evaluating judge reliability, reporting the limitations of automated evaluation, and restricting our conclusions to the evaluated models and benchmarks. Any released artifacts will follow the licenses and usage requirements of the underlying datasets and models.

REPRODUCIBILITY STATEMENT

We provide the formal hallucination definition and taxonomy in sections 2.1 and 2.2, the three-stage evaluation protocol and metrics in section 3, and the mitigation-probe formulation and training procedure in sections 4.1 and 4.2. The experimental section specifies the evaluation pools, agents, judges, and metrics used in our experiments. Additional details on data construction, annotation protocols, judge qualification, and evaluation procedures are provided in Appendix E and Appendix D. We also provide qualitative success and failure analyses in Appendix B.3. These materials are intended to support reproduction of both the hallucination evaluation and mitigation-probe experiments.

REFERENCES

Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zhang Zheng, and Mike Zheng Shou. Hallucination of multimodal large language models: A survey. *arXiv preprint arXiv:2404.18930*, 2024.

Yejin Bang, Ziwei Ji, Alan Schelten, Anthony Hartshorn, Tara Fowler, Cheng Zhang, Nicola Cancedda, and Pascale Fung. Hallulens: Llm hallucination benchmark. *arXiv preprint arXiv:2504.17550*, 2025.

Wentong Chen, Junbo Cui, Jinyi Hu, Yujia Qin, Junjie Fang, Yue Zhao, Chongyi Wang, Jun Liu, Guirong Chen, Yupeng Huo, et al. Guicourse: From general vision language model to versatile gui agent. In *ACL*, pp. 21936–21959, 2025.

Kanzhi Cheng, Qiushi Sun, Yougang Chu, Fangzhi Xu, Li YanTao, Jianbing Zhang, and Zhiyong Wu. Seeclick: Harnessing gui grounding for advanced visual gui agents. In *ACL*, pp. 9313–9332, 2024.

I Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian He, Graham Neubig, Pengfei Liu, et al. Factool: Factuality detection in generative ai–a tool augmented framework for multi-task and multi-domain scenarios. *arXiv preprint arXiv:2307.13528*, 2023.

Lingzhong Dong, Ziqi Zhou, Shuaibo Yang, Haiyue Sheng, Pengzhou Cheng, Zongru Wu, Zheng Wu, Gongshen Liu, and Zhuosheng Zhang. Say one thing, do another? diagnosing reasoning-execution gaps in vlm-powered mobile-use agents. *arXiv preprint arXiv:2510.02204*, 2025.

Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, Rongrong Ji, Caifeng Shan, and Ran He. Mme: A comprehensive evaluation benchmark for multimodal large language models, 2025.

Google DeepMind. Model evaluation – approach, methodology & results: Gemini 3 flash, December 2025. URL https://storage.googleapis.com/deepmind-media/gemini/gemini_3_flash_model_evaluation.pdf. Model id: gemini-3-flash-preview.

Boyu Gou, Ruohan Wang, Boyuan Zheng, Yanan Xie, Cheng Chang, Yiheng Shu, Huan Sun, and Yu Su. Navigating the digital world as humans do: Universal visual grounding for gui agents. In *ICLR*, 2025.

Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, et al. Hallusionbench: an advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *CVPR*, pp. 14375–14385, 2024.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638, 2025.

Wenyi Hong, Weihan Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, et al. Cogagent: A visual language model for gui agents. In *CVPR*, pp. 14281–14290, 2024.

Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55, 2025.

Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.

Raghav Kapoor, Yash Parag Butala, Melisa Russak, Jing Yu Koh, Kiran Kamble, Waseem Alshikh, and Ruslan Salakhutdinov. Omniact: A dataset and benchmark for enabling multimodal generalist autonomous agents for desktop and web, 2024.

Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023a.

Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. Halueval: A large-scale hallucination evaluation benchmark for large language models. In *EMNLP*, 2023b.

Kaixin Li, Ziyang Meng, Hongzhan Lin, Ziyang Luo, Yuchen Tian, Jing Ma, Zhiyong Huang, and Tat-Seng Chua. Screenspot-pro: Gui grounding for professional high-resolution computer use. In *ACM MM*, pp. 8778–8786, 2025.

Wei Li, William E Bishop, Alice Li, Christopher Rawles, Folawiyo Campbell-Ajala, Divya Tyamagundlu, and Oriana Riva. On the effects of data scale on ui control agents. In *NeurIPS*, volume 37, pp. 92130–92154, 2024.

Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. In *EMNLP*, pp. 292–305, 2023c.

Kevin Qinghong Lin, Linjie Li, Difei Gao, Zhengyuan Yang, Shiwei Wu, Zechen Bai, Stan Weixian Lei, Lijuan Wang, and Mike Zheng Shou. Showui: One vision-language-action model for gui visual agent. In *CVPR*, pp. 19498–19508, 2025.

Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. In *ACL*, pp. 3214–3252, 2022.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: NLG evaluation using gpt-4 with better human alignment. In *EMNLP*, pp. 2511–2522, 2023.

Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. In *ICLR*.

Quanfeng Lu, Wenqi Shao, Zitao Liu, Lingxiao Du, Fanqing Meng, Boxuan Li, Botong Chen, Siyuan Huang, Kaipeng Zhang, and Ping Luo. Guiodyssey: A comprehensive dataset for cross-app gui navigation on mobile devices. In *ICCV*, pp. 22404–22414, 2025a.

Zhengxi Lu, Yuxiang Chai, Yaxuan Guo, Xi Yin, Liang Liu, Hao Wang, Han Xiao, Shuai Ren, Guanjing Xiong, and Hongsheng Li. Ui-r1: Enhancing efficient action prediction of gui agents by reinforcement learning. *arXiv preprint arXiv:2503.21620*, 2025b.

Run Luo, Lu Wang, Wanwei He, Longze Chen, Jiaming Li, and Xiaobo Xia. Gui-r1: A generalist r1-style vision-language action model for gui agents. *arXiv preprint arXiv:2504.10458*, 2025.

Potsawee Manakul, Adian Liusie, and Mark Gales. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. In *EMNLP*, pp. 9004–9017, 2023.

Ziyang Meng, Yu Dai, Zezheng Gong, Shaoxiong Guo, Minglong Tang, and Tongquan Wei. Vga: Vision gui assistant – minimizing hallucinations through image-centric fine-tuning, 2024. URL https://arxiv.org/abs/2406.14056.

Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation. *arXiv preprint arXiv:2305.15852*, 2023.

Songqin Nong, Jiali Zhu, Rui Wu, Jiongchao Jin, Shuo Shan, Xiutian Huang, and Wenhao Xu. Mobileflow: A multimodal llm for mobile gui agent. *arXiv preprint arXiv:2407.04346*, 2024.

OpenAI. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.

Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao Li, Yunxin Li, Shijue Huang, et al. Ui-tars: Pioneering automated gui interaction with native agents. *arXiv preprint arXiv:2501.12326*, 2025.

Qwen Team. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.

Qwen Team. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025a.

Qwen Team. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025b.

Vipula Rawte, Swagata Chakraborty, Agnibh Pathak, Anubhav Sarkar, SM Towhidul Islam Tonmoy, Aman Chadha, Amit Sheth, and Amitava Das. The troubling emergence of hallucination in large language models-an extensive definition, quantification, and prescriptive remediations. In *EMNLP*, pp. 2541–2573, 2023.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. In *EuroSys*, 2024.

Yunpeng Song, Yiheng Bian, Yongtao Tang, Guiyu Ma, and Zhongmin Cai. Visiontasker: Mobile task automation using vision based ui understanding and llm task planning. In *ACM UIST*, pp. 1–17, 2024.

Qiushi Sun, Kanzhi Cheng, Zichen Ding, Chuanyang Jin, Yian Wang, Fangzhi Xu, Zhenyu Wu, Chengyou Jia, Liheng Chen, Zhoumianze Liu, et al. Os-genesis: Automating gui agent trajectory construction via reverse task synthesis. In *ACL*, pp. 5555–5579, 2025a.

Yuchen Sun, Shanhui Zhao, Tao Yu, Hao Wen, Samith Va, Mengwei Xu, Yuanchun Li, and Chongyang Zhang. Gui-xplore: Empowering generalizable gui agents with one exploration. In *CVPR*, pp. 19477–19486, 2025b.

Derek Tam, Anisha Mascarenhas, Shiyue Zhang, Sarah Kwan, Mohit Bansal, and Colin Raffel. Evaluating the factual consistency of large language models through news summarization. In *ACL*, 2023.

Xingjian Tao, Yiwei Wang, Yujun Cai, Zhicheng Yang, and Jing Tang. Understanding gui agent localization biases through logit sharpness. *arXiv preprint arXiv:2506.15425*, 2025.

Haoming Wang, Haoyang Zou, Huatong Song, Jiazhan Feng, Junjie Fang, Junting Lu, Longxiang Liu, Qinyu Luo, Shihao Liang, Shijue Huang, et al. Ui-tars-2 technical report: Advancing gui agent with multi-turn reinforcement learning. *arXiv preprint arXiv:2509.02544*, 2025a.

Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. Measuring multimodal mathematical reasoning with math-vision dataset. In *NeurIPS*, volume 37, pp. 95095–95169, 2024a.

Xinyuan Wang, Bowen Wang, Dunjie Lu, Junlin Yang, Tianbao Xie, Junli Wang, Jiaqi Deng, Xiaole Guo, Yiheng Xu, Chen Henry Wu, et al. Opencua: Open foundations for computer-use agents. In *NeurIPS*, 2025b.

Yuxia Wang, Minghan Wang, Muhammad Arslan Manzoor, Fei Liu, Georgi Georgiev, Rocktim Jyoti Das, and Preslav Nakov. Factuality of large language models in the year 2024. *CoRR*, 2024b.

Ziwei Wang, Leyang Yang, Xiaoxuan Tang, Sheng Zhou, Dajun Chen, Wei Jiang, and Yong Li. History-aware reasoning for gui agents. *arXiv preprint arXiv:2511.09127*, 2025c.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*, volume 35, pp. 24824–24837, 2022.

Zhiyong Wu, Zhenyu Wu, Fangzhi Xu, Yian Wang, Qiushi Sun, Chengyou Jia, Kanzhi Cheng, Zichen Ding, Liheng Chen, Paul Pu Liang, et al. Os-atlas: Foundation action model for generalist gui agents. In *ICLR*, 2025.

Yiheng Xu, Zekun Wang, Junli Wang, Dunjie Lu, Tianbao Xie, Amrita Saha, Doyen Sahoo, Tao Yu, and Caiming Xiong. Aguvis: Unified pure vision agents for autonomous gui interaction. *ICML*, 2025.

Yan Yang, Dongxu Li, Yutong Dai, Yuhao Yang, Ziyang Luo, Zirui Zhao, Zhiyuan Hu, Junzhe Huang, Amrita Saha, Zeyuan Chen, et al. Gta1: Gui test-time scaling agent. *arXiv preprint arXiv:2507.05791*, 2025.

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

Jiabo Ye, Xi Zhang, Haiyang Xu, Haowei Liu, Junyang Wang, Zhaoqing Zhu, Ziwei Zheng, Feiyu Gao, Junjie Cao, Zhengxi Lu, et al. Mobile-agent-v3: Fundamental agents for gui automation. *arXiv preprint arXiv:2508.15144*, 2025.

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025.

Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. Alignscore: Evaluating factual consistency with a unified alignment function. In *ACL*, pp. 11328–11348, 2023.

Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Yu Qiao, et al. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? In *ECCV*, pp. 169–186. Springer, 2024.

Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang, Zhikai Wu, Baole Ai, Ang Wang, Wenmeng Zhou, and Yingda Chen. Swift:a scalable lightweight infrastructure for fine-tuning. In *AAAI*, 2025.

A Related Work

A.1 GUI Agents

The integration of Large Language Models (LLMs) and Multimodal Large Language Models (MLLMs) has revolutionized GUI automation. Early GUI agents (Hong et al., 2024; Nong et al., 2024; Song et al., 2024) primarily relied on text-only LLMs to interpret structured representations, such as HTML and accessibility trees. However, these approaches often struggle with the modality gap, as structured text fails to capture critical visual semantics (e.g., icons, layout relations) inherent in graphical interfaces, leading to suboptimal performance in visually complex scenarios.

Consequently, recent advancements have shifted toward MLLMs, especially Vision-Language Models (VLMs), which leverage visual-language alignment to ground actions directly in screen pixels, often adopting the ReAct paradigm (Yao et al., 2022) to interleave perception with stepwise reasoning. Notable works (Gou et al., 2025; Cheng et al., 2024; Chen et al., 2025) significantly enhance element localization by training on large-scale screenshot collections. By unifying action spaces and conducting large-scale post-training, a series of foundational GUI agent models (Wu et al., 2025; Wang et al., 2025b) have achieved robust cross-platform task execution capabilities. Furthermore, data scaling has garnered increasing attention, with numerous studies (Sun et al., 2025a;b) utilizing synthetic data or video trajectories to further unleash agent potential. To transcend the limitations of supervised fine-tuning, recent efforts (Lu et al., 2025b; Luo et al., 2025) have progressively incorporated reinforcement learning, while industrial applications (Qin et al., 2025; Wang et al., 2025a; Ye et al., 2025) have pushed online multi-turn RL to its limits to handle complex interaction dynamics. Beyond training paradigms, researchers are also exploring efficient history organization (Wang et al., 2025c) and test-time scaling strategies (Yang et al., 2025).

While these methods have significantly advanced GUI automation by optimizing task success rates, the explicit diagnosis of the intermediate reasoning process has received less attention. Exploring this dimension offers a promising opportunity to further mitigate hallucinations, which persist as a challenge to agent reliability in dynamic or partially observable environments.

A.2 Hallucination in (M)LLMs

Hallucination in (M)LLMs has emerged as a critical research focal point. Extensive studies have comprehensively explored hallucinations in LLMs, establishing fine-grained taxonomies (Huang et al., 2025; Wang et al., 2024b; Rawte et al., 2023), measurement benchmarks (Lin et al., 2022; Tam et al., 2023; Li et al., 2023b) and automatic evaluation metrics (Zha et al., 2023; Chern et al., 2023; Mündler et al., 2023; Manakul et al., 2023), which have provided valuable insights into the behavioral patterns of hallucination.

The challenge becomes significantly more intricate in MLLMs, where generated content may contradict visual inputs or deviate from sound logical reasoning. To assess the impact of these phenomena, recent works have quantified performance degradation across diverse domains, including object perception (Li et al., 2023c), mathematical reasoning (Lu et al.; Wang et al., 2024a; Zhang et al., 2024) and general multimodal capabilities (Li et al., 2023a; Fu et al., 2025). However, while these methods offer advanced multimodal evaluations, directly applying them to the GUI domain remains challenging. This difficulty primarily stems from the unique characteristics of GUI environments, such as their dynamic, interactive nature and the requirement for precise action execution.

In the context of GUI agents, recent preliminary studies (Tao et al., 2025; Dong et al., 2025; Meng et al., 2024) have observed phenomena such as spatial localization biases and inconsistencies between reasoning and execution. Yet, a systematic analysis of the underlying behavioral patterns and effective mitigation strategies for these domain-specific hallucinations remains underexplored. To bridge this gap, we introduce HalluClear, a comprehensive suite designed to diagnose, evaluate, and mitigate hallucinations in GUI agents. Through three core components, HalluClear precisely detects fine-grained hallucination categories, profiles distinct behavioral patterns and effectively suppresses hallucinations to enhance agent performance.

Table 7: **Performance Comparison on Web and Desktop Tasks (OOD).** Evaluation is performed on GUI-Act-Web, OmniAct-Web and OmniAct-Desktop. The best results are highlighted in **bold**.

Model	Method	GUI-Act-Web			OmniAct-Web			OmniAct-Desktop		
		Type	SR	HR	Type	SR	HR	Type	SR	HR
Qwen2.5-VL-7B	base	99.69	86.91	10.21	99.73	80.11	14.59	98.82	83.21	13.48
	ours	**100.00**	**88.25**	**7.32**	**100.00**	**82.23**	**12.47**	**100.00**	**88.88**	**8.56**
GUI-Owl-7B	base	88.87	75.88	19.59	94.96	68.70	23.61	93.90	75.08	20.86
	ours	**99.79**	**87.94**	**8.87**	**99.73**	**79.57**	**12.47**	**100.00**	**90.37**	**7.59**
UI-TARS-1.5-7B	base	91.24	76.80	22.06	96.29	70.29	28.91	95.08	76.04	22.35
	ours	**100.00**	**87.84**	**8.25**	**100.00**	**79.58**	**14.32**	**100.00**	**86.63**	**10.80**

Table 8: **Complete ablation study on the proposed components.** Performance is evaluated via Grounding ACC on ScreenSpot-V2 (SS-V2) and ScreenSpot-Pro (SS-Pro), Action SR on Android-Control (AC) and GUI-Odyssey (GO) (categorized into Low and High levels), and the overall HR. The best results are highlighted in **bold**.

Model	Setup			Gruonding (ACC ↑)		Action (SR ↑)				HR ↓
	SFT	RL	OODA	SS-V2	SS-Pro	AC-Low	AC-High	GO-Low	GO-High	
Qwen2.5-VL-7B Qwen Team (2025a)	-	-	-	77.18	22.58	60.93	47.22	69.26	47.63	57.54
	✓			81.77	17.29	76.82	58.77	73.84	51.16	39.17
	✓	✓		83.85	18.91	80.50	63.02	71.74	51.36	35.81
	✓		✓	84.74	18.28	79.88	62.61	76.80	53.22	29.33
	✓	✓	✓	**86.62**	**31.31**	**87.49**	**66.72**	**77.71**	**53.99**	**24.98**
GUI-Owl-7B Ye et al. (2025)	-	-	-	81.90	38.52	72.65	57.23	62.20	49.00	54.46
	✓			82.92	39.14	77.85	62.19	74.59	56.02	34.01
	✓	✓		84.88	40.87	79.03	66.20	76.11	58.58	30.90
	✓		✓	86.31	43.26	80.53	65.39	77.11	58.10	26.91
	✓	✓	✓	**88.67**	**43.58**	**87.46**	**68.19**	**79.58**	**58.82**	**21.28**
UI-TARS-1.5-7B Qin et al. (2025)	-	-	-	86.55	43.83	66.86	53.96	75.88	63.22	41.27
	✓			82.53	39.60	72.60	62.54	76.71	56.96	35.80
	✓	✓		85.62	40.49	81.79	63.04	78.49	60.36	29.98
	✓		✓	85.76	41.56	79.46	63.53	79.43	57.88	29.13
	✓	✓	✓	**88.36**	**48.77**	**88.81**	**69.44**	**83.09**	**63.34**	**21.73**

B Extended Experiment Results

B.1 Generalization Across Environments

Although the mitigation probe is trained on limited open-source GUI data, its intended effect is to transfer a structured reasoning and verification paradigm rather than memorize domain-specific interface patterns. We therefore provide additional evidence for cross-environment generalization.

First, the grounding results on ScreenSpot-V2 and ScreenSpot-Pro in the main paper cover Mobile, Desktop, and Web environments. Across these environments and model families, the HalluClear-guided probe consistently improves visual grounding performance, suggesting that the learned reasoning structure is not restricted to a single GUI domain.

Second, we evaluate out-of-distribution action generalization. Although the action-training data are primarily derived from open-source mobile GUI datasets, we follow the OS-ATLAS evaluation setting and test the trained agents on Desktop and Web benchmarks, including GUIcourse (Chen et al., 2025) and OmniAct (Kapoor et al., 2024). As shown in table 7, the probe improves action performance on seen settings while maintaining robust performance in unseen environments. These results indicate that lightweight post-training does not merely memorize mobile-specific action patterns, but can transfer the structured observation–reasoning–action paradigm to broader GUI environments.

Overall, this OOD study provides additional support for the actionability of HalluClear's diagnostic signals: the mitigation probe can improve reliability beyond the environments from which its training data are primarily drawn, while still requiring further validation in more complex enterprise-style and long-horizon CUA workflows.

B.2 Complete Ablation Study

We provide a complete ablation study in table 8 to analyze the components of the mitigation probe. This study is intended to support the actionability analysis in the main text, rather than to claim a standalone training algorithm.

Comparing vanilla SFT with SFT+OODA shows that structured reasoning trajectories improve both downstream capabilities and hallucination-related metrics. This suggests that explicitly separating observation, orientation, decision, and action helps expose intermediate commitments and reduce unsupported grounding or reasoning errors. Adding RL refinement further improves generalization and mitigates SFT-induced degradation. For general-purpose models, RL helps recover performance drops under data-constrained SFT, as seen on ScreenSpot-Pro for Qwen2.5-VL. For GUI-specialized agents, it also alleviates post-training degradation, with similar trends observed for UI-TARS-1.5-7B on ScreenSpot-V2.

Overall, the ablation shows that structured reasoning supervision and lightweight RL refinement play complementary roles in the mitigation probe. This supports HalluClear's actionability claim: diagnostic signals can guide controlled interventions whose effects appear in both hallucination reduction and downstream GUI-agent behavior.

B.3 Qualitative Case Analysis

To complement the quantitative results in the main text, we provide qualitative case studies that illustrate how the HalluClear-guided probe affects different hallucination subtypes. We focus on representative cases from the taxonomy, including *PH.2 Element Existence*, *PH.4 Element Relation*, and *RH.1 Instruction*. These examples are not intended to claim a universal mitigation algorithm, but to show how HalluClear's diagnostic signals can make failure modes more visible and, in some cases, actionable.

B.3.1 Success Case Studies

Addressing *PH.2 Element Existence*. Figure 6 shows a case where the task requires returning to the previous screen. The baseline Qwen2.5-VL agent exhibits a typical *PH.2 Element Existence* hallucination: it incorrectly assumes the presence of a "back arrow" at the bottom-left corner, although no such visual element is present in the screenshot, and therefore attempts to click a nonexistent target. In contrast, the HalluClear-guided probe explicitly checks the observable GUI state, recognizes that no visual back button is available, and selects the system-level action `PRESS_BACK`. This case illustrates how the probe can suppress unsupported commitments to nonexistent GUI elements.

Correcting *PH.4 Element Relation*. Figure 7 illustrates a spatial grounding failure corresponding to *PH.4 Element Relation*. Although the GUI-Owl baseline correctly identifies the semantic target, i.e., the "Continue listening" button, it fails to map this element to the correct spatial location and clicks a non-interactive area. The HalluClear-guided probe mitigates this failure by explicitly separating observation from action selection: it first identifies the target element and its bounding region, then verifies the alignment between the intended element and the predicted action coordinates. This leads to a correctly grounded action and shows how fine-grained diagnosis can help distinguish semantic recognition from spatial relation grounding.

Enforcing *RH.1 Instruction* adherence. Figure 8 presents an *RH.1 Instruction* hallucination. The current-step instruction explicitly asks the agent to "click on the first slide," but the UI-TARS baseline ignores this immediate constraint and instead attempts to open the settings panel, violating the required procedural order. With the HalluClear-guided probe, the agent uses the orientation stage to align its decision with the current-step goal, prioritizing the immediate instruction over a plausible but unsupported longer-term plan. This case shows that structured diagnosis can reduce unsupported commitments arising from instruction misalignment.

Overall, these success cases show that HalluClear-guided reasoning can expose and mitigate different failure mechanisms: nonexistent-element hallucination, incorrect element-relation grounding, and unsupported instruction-level reasoning.

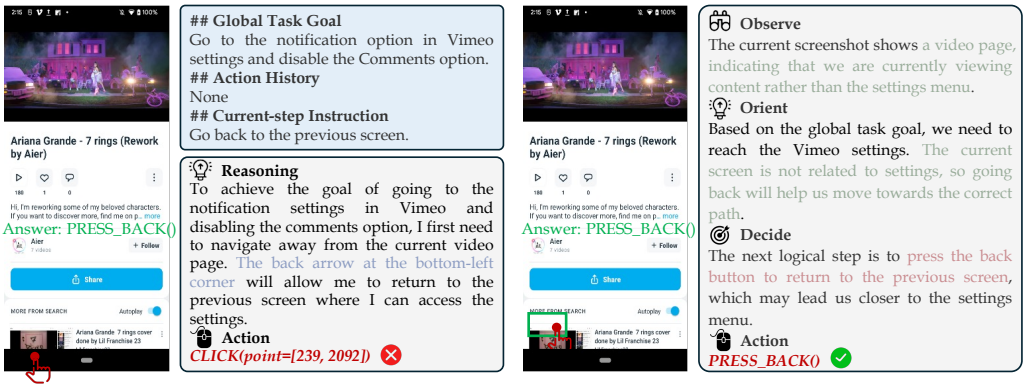

Figure 6: **Qualitative analysis about *PH.2 Element Existence* in GUI tasks.** Qwen2.5-VL (Left) misinterpreting visual cues (hallucinating a "back arrow"), while ours (Right) successfully decides to execute PRESS_BACK through the structured introspection.

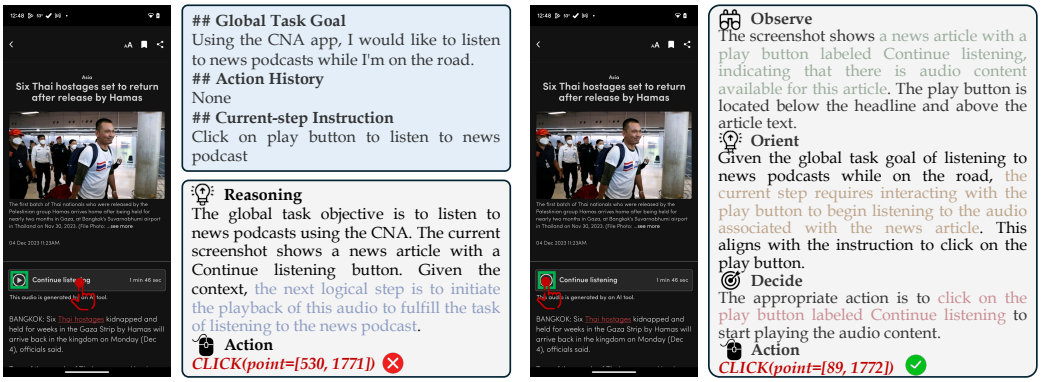

Figure 7: **Qualitative analysis about *PH.4 Element Relation* in GUI tasks.** GUI-Owl (Left) generates coordinates that do not correspond to the target button, while ours (Right) accurately identifies the "Continue listening" button via the "Observe" phase and executes the correct click.

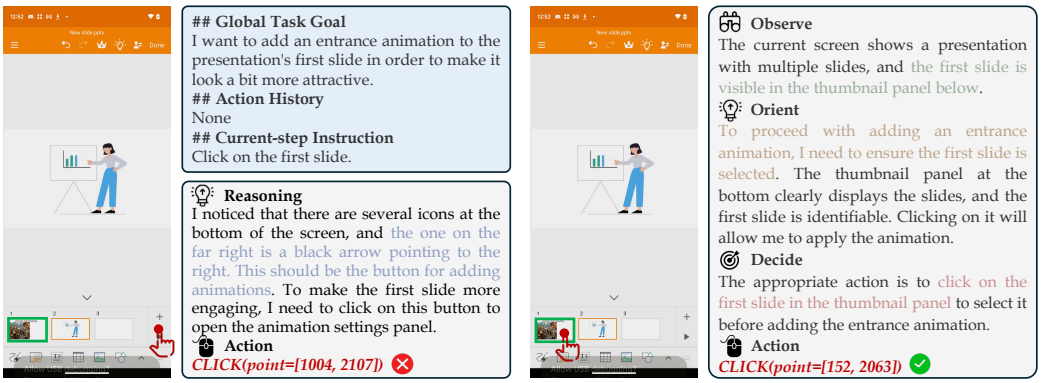

Figure 8: **Qualitative analysis about *RH.1 Instruction* in GUI tasks.** UI-TARS-1.5 (Left) fails to adhere to the instruction in the current-step, while ours (Right) correctly prioritizes the immediate instruction via the "Orient" phase.

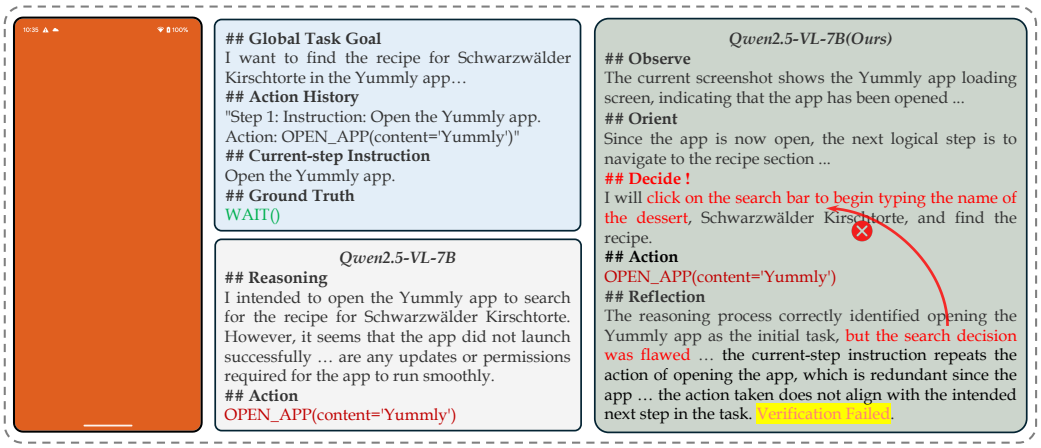

Figure 9: **Qualitative analysis of failure scenarios for Qwen2.5-VL-7B.** Qwen2.5-VL-7B (left) outputs unverified actions after reasoning, while ours (right) identifies a decision deviation in the ”Decide” stage through self-reflection and issues a ”Verification Failed” warning.

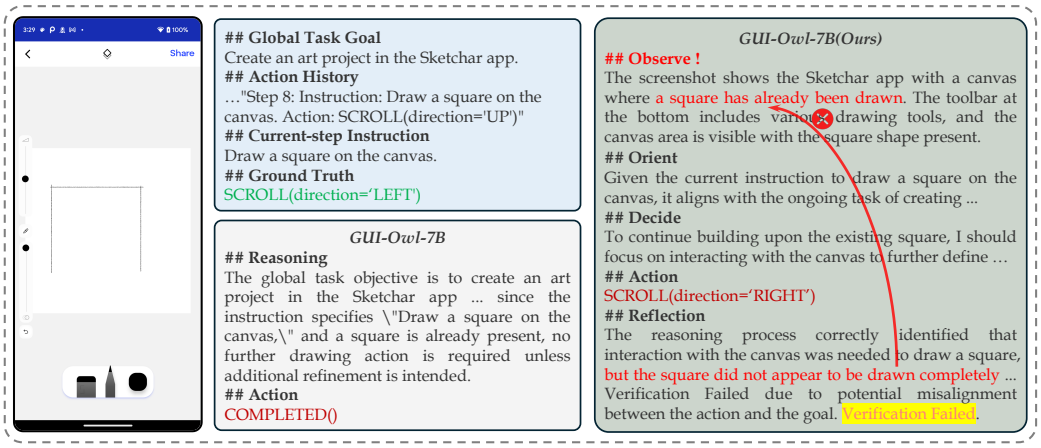

Figure 10: **Qualitative analysis of failure scenarios for GUI-Owl-7B.** GUI-Owl-7B (left) outputs unverified actions after reasoning, while ours (right) localizes the visual hallucination in the ”Observe” stage through self-reflection and issues a ”Verification Failed” warning.

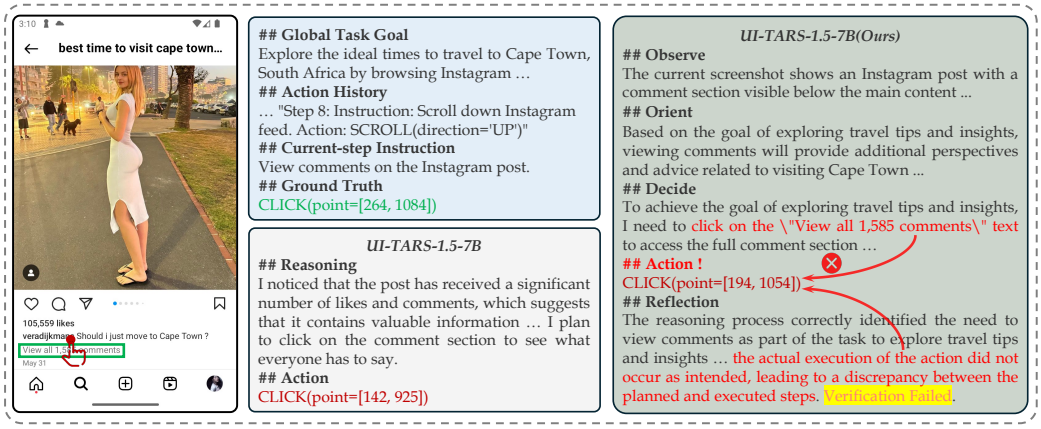

Figure 11: **Qualitative analysis of failure scenarios for UI-TARS-1.5-7B.** UI-TARS-1.5-7B (left) outputs unverified actions after reasoning, while ours (right) detects inconsistencies between the ”Decide” and ”Action” stage through the self-reflection and promptly issues the ”Verification Failed” warning to stop untrustworthy operations.

B.3.2 Failure Case Studies

HalluClear does not eliminate all GUI-agent failures. Agents may still fail in highly dynamic scenarios, such as real-time video streams, or in extremely dense interfaces, such as professional creative software with many small or overlapping controls. Such failures may arise from limited data coverage, insufficient model capacity, severe visual ambiguity, or hidden environment states that cannot be resolved from the current observation alone.

Nevertheless, the failure mode changes qualitatively after applying the HalluClear-guided probe. Native VLM-based GUI agents often fail silently: they commit to an incorrect action path without exposing the unsupported belief or reasoning step that caused the error, and the mistake is directly propagated to the environment. In contrast, HalluClear-guided agents produce structured OODA-style traces, which make it easier to localize the initial failure to a specific stage, such as observation, orientation, decision, or action verification. Moreover, the introspective verification step can sometimes issue a *Verification Failed* signal before executing an unreliable action, converting some black-box execution failures into inspectable and potentially interruptible failures.

As shown in Figures 9, 10, and 11, both general-purpose and GUI-specialized agents can identify the stage at which the reasoning loop first breaks down after HalluClear-guided training. Although this self-diagnosis does not guarantee task success, it provides a more transparent failure trace and helps prevent unverified actions from being blindly executed. These cases support the role of the mitigation probe as an actionability test for HalluClear's diagnoses, while also highlighting the remaining limitations of current GUI-agent reliability.

C More Implementation Details

C.1 Probe Execution Protocol

We details the operational protocol of the OODA loop, using the cross-platform shopping task in Figure 12 to illustrate how the cognitive pipeline physically grounds the agent in the environment:

- **Step 1: Observe.** To mitigate perception hallucinations, we enforce an explicit semantic extraction protocol. As shown in Figure 12(a), instead of implicitly processing embeddings, the agent must output a textual description of critical UI elements. In this instance, by explicitly parsing attributes such as the price ("$209.00") and condition ("Used"), the agent creates a verified evidence base, preventing downstream reasoning from relying on distorted visual features.
- **Step 2: Orient.** This phase addresses the disconnection between the current state and the global goal. The agent synthesizes the **Observe** output with the interaction history. In Figure 12(b), the agent does not evaluate the product in isolation; it cross-references the current price against previous findings (e.g., "Thread AliExpress") stored in memory. This alignment confirms that the current item satisfies the "best deal" constraint, effectively reducing task ambiguity.
- **Step 3: Decide.** The agent formulates a high-level plan derived strictly from the oriented context. As seen in Figure 12(c), the decision to "add the item to the cart" is generated as a logical consequence of the price comparison. This step separates strategic planning from low-level execution, ensuring that the intent is established before any action code is synthesized.
- **Step 4: Act++.** To prevent execution errors, the final output adheres to a $\langle \textit{Answer}, \textit{Reflection}, \textit{Conclusion} \rangle$ structure. The *Answer* specifies the precise coordinate-based command. Crucially, the *Reflection* triggers a consistency check, verifying whether the generated action aligns with the **Decide** plan—for instance, confirming that clicking "Add to Cart" is consistent with the intent to purchase. Finally, the *Conclusion* updates the trajectory memory, closing the loop for the next iteration.

C.2 Unified Action Space Mapping

Existing GUI datasets (Lin et al., 2025; Li et al., 2024; Lu et al., 2025a) often employ heterogeneous action definitions—ranging from raw Python code snippets to JSON-formatted command sequences.

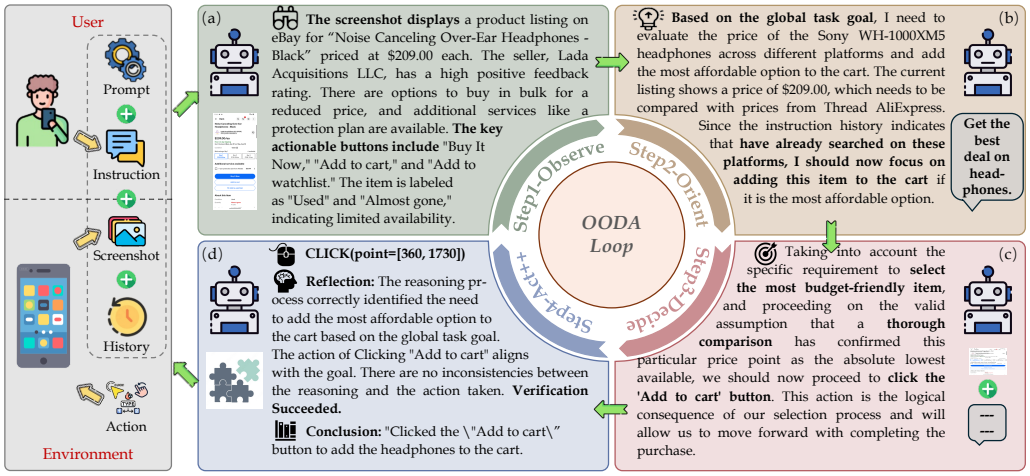

Figure 12: **Execution instance of the OODA loop on a shopping task.** The cognitive trajectory proceeds as follows: (a) Observe extracts specific UI attributes (e.g., price, condition) for visual grounding; (b) Orient aligns the current state with interaction history to confirm the goal; (c) Decide formulates the logical plan; and (d) Act++ generates the operation and verify consistency between reasoning and action, ensuring a self-correcting closed loop.

Table 9: **Definition of the Unified Action Space.** We standardize diverse GUI interactions into a cohesive set of atomic agent actions to simulate human actions.

@] X]@

Human Action	Action Description	Agent Action
Click	Click at a specific point using the coordinates (x, y) in the 'point' field	CLICK(point=[x, y])
Long Press	Long press at a specific point using the coordinates (x, y) in the 'point' field	LONG_PRESS(point=[x, y])
Type	Type the text in the 'content' field	TYPE(content='xxx')
Scroll	Scroll in a specific direction set ('DOWN' or 'UP' or 'RIGHT' or 'LEFT')	SCROLL(direction='xxx')
Open App	Open an application specified by the text in the 'content' field	OPEN_APP(content='xxx')
Press Home	Go to the home screen	PRESS_HOME()
Press Back	Go back to the previous screen	PRESS_BACK()
Press Appselect	Go to the recently used apps screen	PRESS_APPSELECT()
Wait	Wait 5 seconds for the screen to load	WAIT()
Completed	Indicate that the task is completed, with reason in 'content' field	COMPLETED(content='xx')
Incomplete	Indicate that the task is incomplete, with reason in 'content' field	INCOMPLETE(content='xxx')

This heterogeneity hinders the training of a generalized agent and complicates the evaluation of hallucinations.

To address this challenge, we standardize these interactions into a cohesive set of atomic actions that mimic human behavior, thereby establishing a **Unified Action Space**. The detailed definitions and mapping rules are presented in Table 9. Specifically, we categorize diverse dataset-specific labels into three distinct types of operations:

- **Element Interaction:** Precise operations on specific UI elements, including `CLICK`, `LONG_PRESS`, `TYPE`, and `SCROLL`.
- **System Navigation:** Global control commands such as `OPEN_APP`, `PRESS_HOME`, `PRESS_BACK`, and `PRESS_APPSELECT`.
- **Task Flow Control:** Mechanisms to manage task progression, including `WAIT` for loading states, and `COMPLETED` or `INCOMPLETE` to terminate the trajectory with explicit reasoning.

Furthermore, we uniformly categorize all samples from the grounding datasets under the `CLICK` action.

C.3 Training Configurations

Our training pipeline comprised two stages: Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL). For SFT, we utilized the SWIFT (Zhao et al., 2025) framework. To ensure parameter efficiency and mitigate overfitting, we adopted LoRA training for a single epoch with the visual encoder frozen. We employed a cosine learning rate scheduler with a peak learning rate of 2×10^{-5}, a warm-up ratio of 0.05, and a global batch size of 16. For RL, we employ the VeRL (Sheng et al., 2024) framework for one epoch, maintaining most hyperparameters at their default configurations. All experiments are conducted using 8×NVIDIA A100-80G GPUs.

C.4 Prompt Templates

To ensure full reproducibility, we provide the specific prompts used for the baselines—Qwen2.5-VL-7B (Qwen Team, 2025a) in fig. 15, GUI-Owl-7B (Ye et al., 2025) in fig. 14, and UI-TARS-1.5-7B (Qin et al., 2025) in fig. 13—as well as the prompts for our method in fig. 16 and fig. 17. In our experiments, we adhere to the standard prompt formats recommended by the official repositories of the baseline models to ensure a fair comparison. In contrast, the prompt for our method is meticulously designed to enforce the OODA reasoning loop, guiding the agent to explicitly generate observation, orientation, and decision steps before executing an action.

Figure 13: **Prompt of UI-TARS-1.5-7B**

You are a GUI agent. You are given a task and your action history, with screenshots. You need to perform the next action to complete the task.
Output Format
Thought: ...
Action: ...
Action Space

click(start_box='<|box_start|>(x1,y1)<|box_end|>')
long_press(start_box='<|box_start|>(x1,y1)<|box_end|>')
type(content='xxx')
scroll(start_box='<|box_start|>(x1,y1)<|box_end|>', end_box='<|box_start|>(x2,y2)<|box_end|>')
open_app(content='xxx')
press_home()
press_back()
press_appselect() # go to app selection page to show all the opened apps
wait() # wait 5s for the screen to load.
finished()
call_user() # submit the task as failed.

Note
- Use English in ‘Thought‘ part.
- Write a small plan and finally summarize your next action (with its target element) in one sentence in ‘Thought‘ part.

Figure 14: **Prompt of GUI-Owl-7B**

You are a helpful assistant.
Tools
You may call one or more functions to assist with the user query. You are provided with function signatures within <tools></tools> XML tags:

<tools>
{
”type”: ”function”,
”function”: {
”name”: ”mobile_use”,
”description”: ”Use a touchscreen to interact with a mobile device, and take screenshots. ∗ This is an interface to a mobile device with touchscreen. You can perform actions like clicking, typing, swiping, etc. ∗ Some applications may take time to start or process actions, so you may need to wait and take successive screenshots to see the results of your actions. ∗ The screen’s resolution is ${height}x${width}. ∗ Make sure to click any buttons, links, icons, etc with the cursor tip in the center of the element. Don’t click boxes on their edges unless asked.”,
”parameters”: {
”type”: ”object”, ”required”: [”action”],
”properties”: {
”action”: {”type”: ”string”, ”description”: ”The action to perform. The available actions are: ∗ ‘key‘: Perform a key event on the mobile device. − This supports adb’s ‘ keyevent‘ syntax. − Examples: \”volume_up\”, \”volume_down\”, \”power\”, \”camera\”, \”clear\”. ∗ ‘click‘: Click the point on the screen with coordinate (x, y). ∗ ‘long_press‘: Press the point on the screen with coordinate (x, y) for specified seconds. ∗ ‘swipe‘: Swipe from the starting point with coordinate (x, y) to the end point with coordinates2 (x2, y2). ∗ ‘type‘: Input the specified text into the activated input box. ∗ ‘system_button‘: Press the system button. ∗ ‘open‘: Open an app on the device. ∗ ‘wait‘: Wait specified seconds for the change to happen. ∗ ‘terminate‘: Terminate the current task and report its completion status .”, ”enum”: [”key”, ”click”, ”long_press”, ”swipe”, ”type”, ”system_button”, ” open”, ”wait”, ”terminate”]},
”coordinate”: {”type”: ”array”, ”description”: ”(x, y): The x (pixels from the left edge) and y (pixels from the top edge) coordinates to move the mouse to. Required only by ‘action=click‘, ‘action=long_press‘, and ‘action=swipe‘.”},
”coordinate2”: {”type”: ”array”, ”description”: ”(x, y): The x (pixels from the left edge) and y (pixels from the top edge) coordinates to move the mouse to. Required only by ‘action=swipe‘.”},
”text”: {”type”: ”string”, ”description”: ”Required only by ‘action=key‘, ‘action=type‘, and ‘action=open‘.”},
”time”: {”type”: ”number”, ”description”: ”The seconds to wait. Required only by ‘ action=long_press‘ and ‘action=wait‘.”},
”button”: {”type”: ”string”, ”description”: ”Back means returning to the previous interface, Home means returning to the desktop, Menu means opening the application background menu, and Enter means pressing the enter. Required only by ‘action=system_button‘”, ”enum”: [”Back”, ”Home”, ”Menu”, ”Enter”]},
”status”: {”type”: ”string”, ”description”: ”The status of the task. Required only by ‘ action=terminate‘.”, ”enum”: [”success”, ”failure”]}
}
}
}
}
</tools>

For each function call, return a json object with function name and arguments within <tool_call></tool_call> XML tags:

<tool_call>
{ ”name”: <function−name>, ”arguments”: <args−json−object> }
</tool_call>

Figure 15: **Prompt of Qwen2.5-VL-7B**

You are a helpful assistant.
Tools
You may call one or more functions to assist with the user query. You are provided with function signatures within <tools></tools> XML tags:

<tools>
{
"type": "function",
"function": {
"name_for_human": "mobile_use",
"name": "mobile_use",
"description": "Use a touchscreen to interact with a mobile device, and take screenshots. * This is an interface to a mobile device with touchscreen. You can perform actions like clicking, typing, swiping, etc. * Some applications may take time to start or process actions, so you may need to wait and take successive screenshots to see the results of your actions. * The screen's resolution is ${height}x${width}. * Make sure to click any buttons, links, icons, etc with the cursor tip in the center of the element. Don't click boxes on their edges unless asked.",
"parameters": {
"type": "object", "required": ["action"],
"properties": {
"action": {
"type": "string",
"description": "The action to perform. Available actions: * key: Perform a key event (adb keyevent syntax, e.g. \"volume_up\", \"home\"). * click: Click point (x,y). * long_press: Press (x,y) for seconds. * swipe: Swipe from (x,y) to (x2,y2). * type: Input text. * system_button: Press system button. * open: Open app. * wait : Wait seconds. * terminate: Terminate task.",
"enum": ["key", "click", "long_press", "swipe", "type", "system_button", "open", " wait", "terminate"]},
"coordinate": { "type": "array", "description": "(x, y): Coordinates (pixels from left/top). Required by click, long_press, swipe." },
"coordinate2": { "type": "array", "description": "(x, y): End coordinates. Required only by action=swipe."},
"text": { "type": "string", "description": "Required only by action=key, action=type, and action=open."},
"time": { "type": "number", "description": "The seconds to wait. Required only by action=long_press and action=wait."},
"button": {
"type": "string",
"description": "System button type. Required only by action=system_button",
"enum": ["Back", "Home", "Menu", "Enter"]},
"status": {
"type": "string", "description": "Task status. Required only by action=terminate.",
"enum": ["success", "failure"]}
}
},
"args_format": "Format the arguments as a JSON object."
}
}
</tools>

For each function call, return a json object with function name and arguments within <tool_call></tool_call> XML tags:

<tool_call>
{ "name": <function−name>, "arguments": <args−json−object> }
</tool_call>

Figure 16: **Low-level Prompt of Ours**

As a Reasoning GUI Agent, your responsibility is to provide the correct solution that specifies the action to be executed, based on the global task goal, the action history, the current-step instruction and the screenshot.
The action space is as follows:

CLICK(point=[x, y]) ## Click at a specific point on the screen using the coordinates (x , y) in the 'point' field.
LONG_PRESS(point=[x, y]) ## Long press at a specific point on the screen using the coordinates (x, y) in the 'point' field.
TYPE(content='xxx') ## Type the text in the 'content' field.
SCROLL(direction='DOWN' or 'UP' or 'RIGHT' or 'LEFT') ## Scroll in a specific direction set in the 'direction' field.
OPEN_APP(content='xxx') ## Open an application specified by the text in the ' content' field.
PRESS_HOME() ## Go to the home screen.
PRESS_BACK() ## Go back to the previous screen.
WAIT() ## Wait 5 seconds for the screen to load.
PRESS_APPSELECT() ## Go to the recently used apps screen.
COMPLETED(content='xxx') ## Indicate that the task is completed, and the information in the 'content' field explains why the task is completed.
INCOMPLETE(content='xxx') ## Indicate that the task is incomplete, and the information in the 'content' field explains why the task is incomplete.

The solution includes the following four parts: thinking, answer, reflection, and conclusion. Each part is to be enclosed within specific tags:

1. <thinking>thinking</thinking>: Present your complete logical chain of problem−solving. It follows a clear and concise three−step logical reasoning process, i.e., Step 1: Observe; Step 2: Orient; Step 3: Decide.
− Step 1: Observe: Describe in detail the layout, state, and key elements of the current −step screenshot.
− Step 2: Orient: Infer what you should do in the current step based on the global task goal, the action history, the current−step instruction and the screenshot.
− Step 3: Decide: Decide which action to execute and focus the corresponding region in the screenshot, based on the analysis from "Step 1: Observe" and "Step 2: Orient".
2. <answer>answer</answer>: Provide the action to be executed in the specified format of the Action Space defined above. If you conclude that the task cannot be completed, output exactly: "Task Failed".
3. <reflection>reflection</reflection>: Review the accuracy of the reasoning process within <thinking> and then verify the consistency between the reasoning process within <thinking> and the result within <answer>. If any error or inconsistency exists, end with: "Verification Failed"; otherwise, end with: "Verification Succeeded".
4. <conclusion>conclusion</conclusion>: Summarize the action taken in the current step. \newline

Respond according to the user's input, supplying the requested sections of the problem −solving process, i.e., <thinking>thinking</thinking> <answer>answer</answer ><reflection>reflection</reflection><conclusion>conclusion</conclusion>.

Solve the problem in accordance with these guidelines.

Figure 17: **High-level Prompt of Ours**

As a Reasoning GUI Agent, your responsibility is to provide the correct solution that specifies the action to be executed, based on the global task goal, the action history, and the screenshot.
The action space is as follows:

CLICK(point=[x, y]) ## Click at a specific point on the screen using the coordinates (x , y) in the 'point' field.
LONG_PRESS(point=[x, y]) ## Long press at a specific point on the screen using the coordinates (x, y) in the 'point' field.
TYPE(content='xxx') ## Type the text in the 'content' field.
SCROLL(direction='DOWN' or 'UP' or 'RIGHT' or 'LEFT') ## Scroll in a specific direction set in the 'direction' field.
OPEN_APP(content='xxx') ## Open an application specified by the text in the ' content' field.
PRESS_HOME() ## Go to the home screen.
PRESS_BACK() ## Go back to the previous screen.
WAIT() ## Wait 5 seconds for the screen to load.
PRESS_APPSELECT() ## Go to the recently used apps screen.
COMPLETED(content='xxx') ## Indicate that the task is completed, and the information in the 'content' field explains why the task is completed.
INCOMPLETE(content='xxx') ## Indicate that the task is incomplete, and the information in the 'content' field explains why the task is incomplete.

The solution includes the following four parts: thinking, answer, reflection, and conclusion. Each part is to be enclosed within specific tags:

1. <thinking>thinking</thinking>: Present your complete logical chain of problem−solving. It follows a clear and concise three−step logical reasoning process, i.e., Step 1: Observe; Step 2: Orient; Step 3: Decide.
− Step 1: Observe: Describe in detail the layout, state, and key elements of the current −step screenshot.
− Step 2: Orient: Infer what you should do in the current step based on the global task goal, the action history, and the screenshot.
− Step 3: Decide: Decide which action to execute and focus the corresponding region in the screenshot, based on the analysis from "Step 1: Observe" and "Step 2: Orient".
2. <answer>answer</answer>: Provide the action to be executed in the specified format of the Action Space defined above. If you conclude that the task cannot be completed, output exactly: "Task Failed".
3. <reflection>reflection</reflection>: Review the accuracy of the reasoning process within <thinking> and then verify the consistency between the reasoning process within <thinking> and the result within <answer>. If any error or inconsistency exists, end with: "Verification Failed"; otherwise, end with: "Verification Succeeded".
4. <conclusion>conclusion</conclusion>: Summarize the action taken in the current step.

Respond according to the user's input, supplying the requested sections of the problem −solving process, i.e., <thinking>thinking</thinking><answer>answer</answer ><reflection>reflection</reflection><conclusion>conclusion</conclusion>.

Solve the problem in accordance with these guidelines.

D HALLUCLEAR EVALUATION WORKFLOW

D.1 PHILOSOPHY OF EVALUATION WORKFLOW DESIGN

We aim for our evaluation workflow to exhibit three core features: (1) **quantification**, (2) **automation** and (3) **scalability**. However, the non-mutually exclusive label space and subjective human judgments makes the problem akin to reference-free evaluation. As a result, achieving quantitative, automated, and scalable assessment remains nontrivial in practice.

Our ultimate goal is to evaluate the Hallucination Rate (HR) of M agents on an HR-Pool consisting of N datasets. However, it is clearly impractical to repeat $M \times N$ manual evaluation (similar to the JQ-Bench construction) on $M \times N$ $\{q_j, \text{agent}_m(q_j)\}$ output sets. Directly resorting to VLM-as-a-judge, on the other hand, makes it difficult to conduct an objectively, systematically, and unbiased quantitatively evaluation of HR. This is why we therefore introduce a novel three-stage paradigm. In summary, to achieve scalable and reliable automated quantitative evaluation, we integrate the classic benchmark with golden reference paradigm and the VLM-as-a-judge paradigm to form a three-stage hallucination evaluation workflow:

- **Stage 1.** Manual construction of JQ-Bench;
- **Stage 2.** Selection and member credibility assessment of the VLM judge panel based on JQ-Bench;
- **Stage 3.** HR evaluation based on the selected VLM judges.

D.2 STAGE 1: CONSTRUCTION OF JQ-BENCH

Filtering (q_i, a_i). In practice, JQ-Bench, as an admission test for judge candidates, requires high-quality manual annotation. We first need to collect a batch of erroneous responses from advanced open-source agents on public datasets and then filter out desired $\{(q_i, a_i)\}_i^{|\text{JQ-Bench}|}$ that exhibit hallucinations. Specifically, during the annotation process, we filter out these data points:

- low-quality queries caused by the dataset;
- low-quality responses generated by the agent; and
- responses that meet the quality standards, mismatch the currently annotated unique golden action in the offline dataset, but are manually identified as reasonable immediate decisions given the current context.

Thus, the remaining (q_i, a_i) pairs can be treated as cases of hallucinations for expert labeling.

Annotators. Given a pair of (q_i, a_i), the annotator must determine the type of hallucination present in this pair. Due to the subjectivity and non-mutually exclusive nature of hallucination classification, and the requirement that annotators have a sufficient understanding of the agent's characteristics and the current GUI environment to identify reasonable immediate decisions, building JQ-Bench is actually much more difficult than many other types of related benchmarks, and also places higher demands on the professionalism of the annotators. Therefore, in practice, the initial JQ-Bench was built **entirely by the authors of this paper, without using any additional crowdsourcing outside the author list or any automated VLM-based evaluation**.

Introduction to the labeling protocol. The labeling protocol is a lookup-style set of rules that is expanded from specific cases and constructed bottom-up. Annotators are required to follow it when performing the manual mapping $\mathcal{Q} \times \mathcal{A} \to 2^{\mathcal{HS}}$, where $\mathcal{HS}$ denotes the **H**allucination **S**pace. In practice, we cap the number of potential hallucination labels in gt_i associated with each (q_i, a_i) pair at 3 (i.e., $\forall (q_i, a_i), |gt_i| < 3$). When multiple hallucination subtypes are assigned, their relationship can be:

- **"and"**: hallucination A occurs in one part, while hallucination B occurs elsewhere; or
- **"or"**: a hallucination is present at a specific point, but it may be attributed to either subtype A or subtype B.

For more internal details of the labeling protocol, please refer to Appendix E.

Table 10: **JQ-Bench leaderboard for candidate VLM judges across annotators.** For each table cell, we report two credibility scores: the **top row** is the **binary** accuracy for hallucination detection (hallucinated vs. non-hallucinated), while the **bottom row** is the **fine-grained** accuracy (8 hallucination types + 1 correct). Each entry is computed over **three** independent runs and shown as $\mu^{+\Delta_{\max}}_{-\Delta_{\min}}$.

Annotator	Gemini-3 flash-preview	Seed-1.8	Qwen3-VL Plus	Gemini-2.5 flash	Seed-1.6 Vision	GLM-4.6V	GPT-4o
Annotator A	$89.8^{+0.7}_{-0.6}$	$84.0^{+0.1}_{-0.2}$	$79.6^{+1.1}_{-0.6}$	$78.0^{+0.6}_{-0.4}$	$81.8^{+0.5}_{-0.5}$	$83.4^{+0.4}_{-0.6}$	$78.5^{+0.3}_{-0.5}$
	$80.6^{+0.6}_{-0.9}$	$73.6^{+0.2}_{-0.2}$	$62.1^{+0.9}_{-0.8}$	$59.7^{+1.3}_{-0.8}$	$72.1^{+0.4}_{-0.2}$	$56.0^{+1.2}_{-0.6}$	$63.1^{+0.5}_{-0.5}$
Annotator B	$89.0^{+0.4}_{-0.4}$	$87.0^{+0.4}_{-0.5}$	$82.7^{+0.9}_{-1.2}$	$79.8^{+0.9}_{-0.6}$	$78.7^{+0.6}_{-0.5}$	$83.5^{+0.4}_{-0.3}$	$81.4^{+0.8}_{-0.9}$
	$81.1^{+0.9}_{-1.2}$	$77.4^{+0.2}_{-0.3}$	$68.9^{+0.5}_{-0.7}$	$64.5^{+1.0}_{-0.6}$	$69.3^{+1.0}_{-0.5}$	$63.7^{+0.3}_{-0.4}$	$69.7^{+0.2}_{-0.3}$
Annotator C	$90.7^{+0.3}_{-0.3}$	$62.4^{+0.0}_{-0.1}$	$63.3^{+1.1}_{-0.6}$	$61.1^{+0.4}_{-0.3}$	$63.1^{+0.6}_{-0.4}$	$57.1^{+0.3}_{-0.4}$	$61.2^{+0.8}_{-0.7}$
	$82.5^{+0.4}_{-0.6}$	$52.5^{+0.1}_{-0.2}$	$46.3^{+0.8}_{-1.3}$	$44.9^{+0.6}_{-0.7}$	$44.4^{+0.7}_{-0.7}$	$36.8^{+1.0}_{-0.8}$	$45.4^{+0.5}_{-0.7}$
Average	$89.8^{+0.5}_{-0.3}$	$78.7^{+0.1}_{-0.1}$	$76.0^{+0.5}_{-0.8}$	$73.7^{+0.2}_{-0.1}$	$75.1^{+0.4}_{-0.5}$	$75.7^{+0.3}_{-0.2}$	$74.5^{+0.6}_{-0.5}$
	$81.4^{+0.6}_{-0.6}$	$68.8^{+0.1}_{-0.2}$	$60.0^{+0.6}_{-0.9}$	$57.2^{+0.3}_{-0.2}$	$62.8^{+0.5}_{-0.5}$	$53.2^{+0.2}_{-0.2}$	$60.4^{+0.2}_{-0.2}$

Necessity of the labeling protocol. As discussed above, considering the inherent subjectivity of the definition of hallucination, ensuring the high quality of JQ-Bench requires a strict and explicit labeling protocol. Spatially, the protocol must systematically disambiguate potential divergences in semantic interpretation across different annotators to improve inter-annotator consistency. Temporally, given the long annotation cycle, it must also ensure that the same annotator assigns consistent labels to similar cases under the same decision rules, thereby maintaining intra-annotator consistency over time. Finally, in the subsequent Stage 2/3 adoption of the VLM-as-a-judge paradigm, we likewise abstract this protocol into the system prompt for VLM judges, and the clarity and low ambiguity of this prompt is crucial to the credibility of VLM-based judgments. We provide the specific prompt of VLM judges in fig. 18.

D.3 Stage 2: Selection of VLM-Judges

"Strict-in, Lenient-out." So far, we have constructed the JQ-Bench, denoted as $\{(q_i, a_i, gt_i)\}_i^{|\text{JQ-Bench}|}$. Next, we must select, from the pool of candidate VLM judges, those individuals that perform well on JQ-Bench to serve on the final judge panel in Stage 3. We assume that, judges that can correctly identify not only whether a (q_i, a_i) pair contains hallucination, but also its fine-grained subtype (i.e., meeting eq. (4)), have a more reliable understanding of the hallucination phenomena that may arise in GUI agents. This naturally yields a **strict-in, lenient-out** scheme: in Stage 2, we impose a more stringent admission criterion (8 subtypes + no-fatal-hallucination), so that the VLM judges that pass Stage 2 are more reliable when we later conduct binary HR evaluation (hallucination vs. no-fatal-hallucination) in Stage 3.

Credibility assessment. We assess the credibility C_k of the kth VLM judge based on eq. (4), and this credibility is directly related to the labeling quality of JQ-Bench. If, in Stage 2, all judge candidates perform poorly on the current version of JQ-Bench, this could indicate that existing general-purpose VLMs are not yet capable of a fine-grained hallucination judgment task. However, this interpretation is less plausible than issues in the evaluation setup. Here, a more likely explanation can be:

- **(1)** the prompt provided to the judges is semantically ambiguous; or
- **(2)** the annotators introduce random errors when adhering to the labeling protocol.

To address (1), we strengthen the judges' understanding of hallucination subtypes by providing additional in-context learning examples in the system prompt and encouraging deliberate System-2-style reasoning. To address (2), we adopt a strategy that we call "Back to Stage 1".

Table 11: **Binary hallucination detection performance of candidate VLM judges on JQ-Bench.** We report precision, recall, F1, and accuracy under the **binary** setting. Numbers are averaged over **three** independent runs and formatted as $\mu^{+\Delta_{\max}}_{-\Delta_{\min}}$.

Metric	Gemini-3 flash-preview	Seed-1.8	Qwen3-VL Plus	Gemini-2.5 flash	Seed-1.6 Vision	GLM-4.6V	GPT-4o
Precision	$90.9^{+0.2}_{-0.1}$	$86.7^{+0.3}_{-0.3}$	$82.8^{+0.5}_{-0.5}$	$78.8^{+0.2}_{-0.2}$	$80.4^{+0.4}_{-0.4}$	$86.7^{+0.3}_{-0.3}$	$81.0^{+0.5}_{-0.6}$
Recall	$94.3^{+0.5}_{-0.3}$	$75.8^{+0.1}_{-0.1}$	$81.5^{+0.4}_{-0.7}$	$83.5^{+0.2}_{-0.2}$	$83.5^{+0.2}_{-0.1}$	$75.8^{+0.1}_{-0.1}$	$81.5^{+0.4}_{-0.4}$
F1	$92.6^{+0.4}_{-0.2}$	$80.9^{+0.2}_{-0.1}$	$82.1^{+0.4}_{-0.6}$	$81.1^{+0.1}_{-0.1}$	$81.9^{+0.3}_{-0.3}$	$80.9^{+0.2}_{-0.1}$	$81.2^{+0.4}_{-0.3}$
Acc	$89.8^{+0.5}_{-0.3}$	$78.7^{+0.1}_{-0.1}$	$76.0^{+0.5}_{-0.8}$	$73.7^{+0.2}_{-0.1}$	$75.1^{+0.4}_{-0.5}$	$75.7^{+0.3}_{-0.2}$	$74.5^{+0.6}_{-0.5}$

Back to Stage 1: Human-Judge Alignment. Concretely, for cases on which multiple judges consistently fail in a given iteration between Stage 1 and Stage 2, the annotators conduct a detailed failure case study grounded in the judges' reasoning traces, and explicitly discuss whether the hallucination subtype predicted by the judges is reasonable for the corresponding case. If the annotators unanimously agree that it is reasonable, we add this subtype to the gt_i set. If the existing gt_i is already full (up to 3), we compare the judges' predicted subtype against each existing ground-truth subtype in gt_i and decide whether to update gt_i. We refer to this iterative procedure as "**Human–Judge Alignment**." It's worth noting that, this is not an improper practice such as benchmark contamination or data leakage based on judges' outputs; rather, it is a conservative mechanism for improving JQ-Bench labeling quality, triggered only when the annotators unanimously deem it necessary.

Back to Stage 1: Human-Human Alignment. Another aspect of the aforementioned "Back to Stage 1" strategy is to diagnose annotator alignment via the VLM judges' output credibility. We divide the JQ-Bench according to the annotators of the data and record the credibility of different VLM judges on subsets constructed by different annotators. Ideally, the credibility of different judges should all maintain consistency across corresponding subsets of annotators. We also focus on this in the iterative optimization of JQ-Bench.

JQ-Bench Leaderboard. To date, we have obtained the final, optimized version of JQ-Bench and the results of various advanced VLM judge candidates on JQ-Bench, as shown in tables 10 and 11. As one would expect, the newly released models (e.g., Gemini-3 and Seed-1.8) outperform earlier-generation models (e.g., GPT-4o) on JQ-Bench. This, in turn, provides evidence that our labeling protocol is is clear and reasonable for the judges to follow, and that the construction of JQ-Bench exhibits high annotation quality due to its strict adherence to the labeling protocol.

D.4 Stage 3: HR Evaluation

From Stage 2 to Stage 3: Generalizability. We primarily rely on two strategies to ensure the generalizability of the evaluation transfer from JQ-Bench to HR-Pool:

- To avoid excessively large domain gaps, the queries in JQ-Bench maintain a distributional consistent with those in HR-Pool; and
- To prevent the judges from overfitting on JQ-Bench, we use more robust commercial/open-source general-purpose VLMs as the judge candidates, rather than any fine-tuned hallucination detectors.

After constructing JQ-Bench, we determine the HR-Pool and the VLM judge panel based on the above principles.

Selecting HR-Pool. The upper bound of the evaluation accuracy in our evaluation workflow is determined by the credibility of the VLM judges, and the evaluation of their credibility (Stage 2) is directly related to the quantity and quality of JQ-Bench, not the quantity of HR-Pool. In other words, from a quantitative perspective, the upper bound of the evaluation accuracy is limited by JQ-

Bench rather than HR-Pool. Therefore, we argue that, selecting an HR-Pool with a larger number of queries than JQ-Bench is counterintuitive.

Metrics and Calibration. To strictly quantify the hallucination metrics, we introduce a control category designated as *No Fatal Hallucination*, aiming to decouple genuine errors from faithful executions and mitigate the confounding effects of True Negatives and False Positives. We employ three VLM judges from **Stage 2** to independently classify the samples into a taxonomy of nine distinct categories. For each judge, the rate for a specific category is determined by the ratio of samples assigned to that category to the total number of samples in the HR-Pool. Subsequently, by explicitly excluding the *No Fatal Hallucination* class, the judge-specific *Overall HR* is derived by accumulating the individual rates across the remaining eight fine-grained categories, representing the total prevalence of actual hallucinations identified by that judge.

Furthermore, to mitigate the variance and potential bias of individual judges, we compute the *Final Calibrated Distribution*. This process involves a weighted ensemble approach: we first normalize the raw credibility scores of the three judges so that they sum to unity. Subsequently, the calibrated rate for each fine-grained hallucination category is calculated as the weighted sum of the rates reported by the individual judges, using their normalized credibility scores as weights. This results in a debiased distribution that reflects the consensus probability for each specific type of hallucination.

E LABELING PROTOCOL

Here, we revisit section 2.2 to provide a more detailed description of the protocol.

NONH.1 NO FATAL HALLUCINATION - "CORRECT"

Definition. NonH.1 designates the normative baseline wherein the agent successfully executes the task devoid of any significant hallucinations classified under the Perception (PH) or Reasoning (RH) taxonomies.

Details. NonH.1 also accommodates instances where the reasoning trace exhibits minor, inconsequential aberrations. Provided that these deviations do not propagate downstream, or are autonomously rectified via self-correction, and thus exert no adverse impact on the final decision, such cases are rigorously classified as "No Fatal Hallucination".

Note that the broader concept of "No Fatal Hallucination" theoretically encompasses both True Negatives (NonH.1) and False Positives (NonH.2) during offline evaluation.

NONH.2 "FALSE POSITIVE" - "CONFIRM"

Definition. NonH.2 refers to cases flagged as errors by automated ground truth matching on offline datasets, but subsequently validated as correct by expert annotators with limited access to $\tilde{s}_t$.

Details. These cases are NOT hallucinations; rather, they represent legitimate behavioral variations **under the current $\tilde{s}_t$, given no additional privileged states**, while highlighting the inherent limitations of rigid string-matching or coordinate-matching evaluations in offline datasets.

We deliberate retained these cases within JQ-Bench to serve as adversarial controls during Stage 2, where the ground truth actions is also provided to the judge candidates as a reference. **These "False Positive" cases are critical to mitigating shortcut judging**: they prevent the judge from achieving inflated classification accuracy through superficial pattern matching or rigid alignment between the GT and the agent's prediction. Therefore, we ensure that the evaluation metric reflects a robust capability for deep semantic reasoning, rather than a mere proficiency in text comparison.

We categorize these "False Positives" into 4 primary dimensions of equivalence:

- **Action Isomorphism (System vs. Element)**: The agent substitutes a **legal** system-level API call with a functionally equivalent GUI interaction (or vice versa, notice that API call MUST be one of the legal actions with in the action space). Prevalent examples involve "Go Back" (`PRESS_BACK` vs. Back-Arrow Icon, see fig. 19) and invoking an application (`OPEN_APP` vs. clicking the app's icon, see fig. 20).

Figure 18: **Prompt of VLM Judges**

You are an expert in hallucination evaluation for GUI Agents.

Task Requirement
You are given an input consisting of a question, an accompanying screenshot (represented by <image>) and the reference answer. You are also provided with the output of a Reasoning GUI Agent, including the reasoning steps within <thinking> and predicted answer within <answer>. Your task is to analyze the output and determine whether it contains *any* hallucinations.
Types of Hallucinations
1. **Screenshot State Hallucination**: The agent misinterprets the holistic state of the current screenshot. This reflects a failure in global scene understanding that transcends the perception of specific UI elements.
2. **Element Existence Hallucination**: The agent erroneously identifies or fabricates non-existent elements that do not appear in the screenshot.
3. **Element Attribute Hallucination**: The agent misinterprets the intrinsic attributes of UI elements, specifically regarding their visual appearance, intended function, and operational affordances.
4. **Element Relation Hallucination**: The agent misunderstands the relationships between UI elements, or between elements and the overall screen, primarily concerning spatial arrangements. Instances where grounding coordinates fall outside the screenshot boundaries are categorized here.
5. **Instruction Hallucination**: The agent fails to adhere to or explicitly disregards low-level step instructions provided within the query. The primary focus is whether the reasoning steps within <thinking> demonstrate a clear intent to follow current-step instruction.
6. **Context Hallucination**: The agent demonstrates inconsistencies within the context, specifically regarding discrepancies between the query and reasoning steps, or between the reasoning steps and the predicted answer. Typical examples of the former include invoking illegal actions (the predicted answer within <answer> does not follow the action space defined in the question) and failing to utilize required historical information, while the latter involves generating answers decoupled from the agent's own reasoning steps.
7. **Logical Hallucination**: The agent exhibits manifest logical errors or discontinuities within its reasoning steps. This focuses on flawed causal transitions rather than simple correspondence errors.
8. **Factuality Hallucination**: The agent lacks relevant external knowledge, leading to overconfident fabrication of information or incorrect factual assertions.
9. **No Fatal Hallucination**: The agent exhibits no obvious hallucinations, or its hallucinations are subsequently self-corrected or simply so mild as not to affect the correctness of its decision.
Few-shot Examples for Reference
——- Examples are omitted to fit one-page style —-
Evaluation Process
1. **Perception & Grounding Verification**:
- Cross-reference every UI element mentioned within <thinking> against the actual <image>. Verify if the element actually exists and if its appearance, function and affordance are recognized correctly.
- Validate the grounding coordinates within <answer>. Visually estimate if the point falls within the visual bounding box of the target element. Flag any coordinates that are obviously off-target or fall outside the screenshot boundaries.
2. **Reasoning & Consistency Check**:
- Assess whether the reasoning steps within <thinking> logically derive from the observed GUI state and strictly adhere to the context and user's instructions.
- Verify that the final predicted answer within <answer> is a consistent and necessary conclusion of the reasoning steps. Ensure the action aligns with the user's intent (reference answer as a guide) and follows the action space defined in the question.
Output Format
<reason>- Step 1 (Perception & Grounding Verification): [Describe findings] - Step 2 (Reasoning & Consistency Check): [Describe findings]</reason>
<result>Based on the analysis within <reason>, determine whether the agent's output contains hallucinations. If *any* hallucination is detected, output the exact type name from the "Types of Hallucinations" list defined above; otherwise, output "CONFIRM".</result>
Note
- For "CLICK" action, the parameters include click coordinate and bounding box.
- The reference answer is a guide, not the absolute truth. Verify against the question and <image> yourself.
- If the agent's output contains multiple hallucinations, output the first one encountered.

Now, for the following input, analyze the output of agent and provide your final evaluation according to the schema above.
Question: {question}
Screenshot: <image>
Reference Answer: {ref_answer}
Output of Agent: {output}

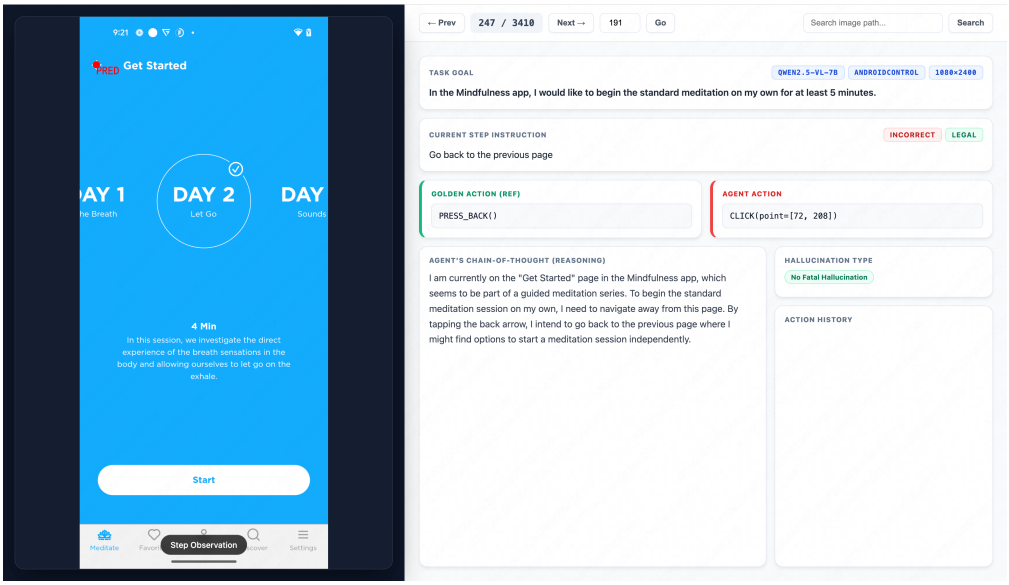

Figure 19: **Case of NonH.2 "False Positive"**: "Go Back."

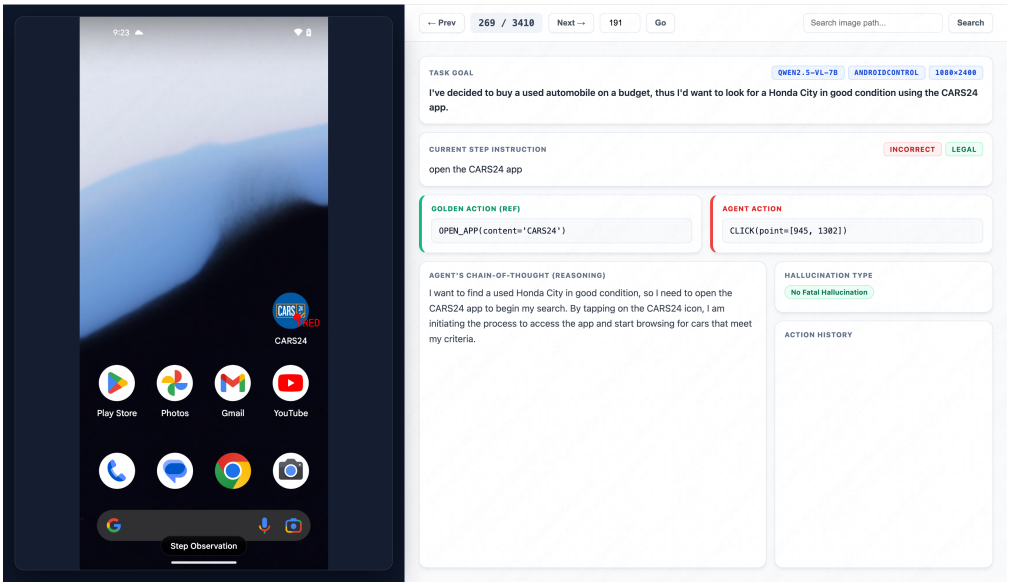

Figure 20: **Case of NonH.2 "False Positive"**: "Open APP."

- **Interaction Pathway Equivalence**: The agent achieves the target state transition through a different but equally effective interaction modality. For instance, the GT typically dictates typing the full query (e.g., `TYPE("<query>")`), while the search engine displays the correct target as the suggestion/history, or this query is already displayed in the clipboard area (see fig. 21). Although this shortcut is valid and often more efficient, sometimes it may attribute to RH.1, and we will elaborate later.

- **Grounding Redundancy**: GUI elements often possess redundant clickable regions that trigger the same effect. For example, a list item may consist of a textual label, an icon, and a chevron arrow, all of which are interactive. If the GT annotation is anchored to the text, but the agent clicks the adjacent icon, strict coordinate matching produces a false negative. However, functionally, the outcome is identical (see fig. 22).

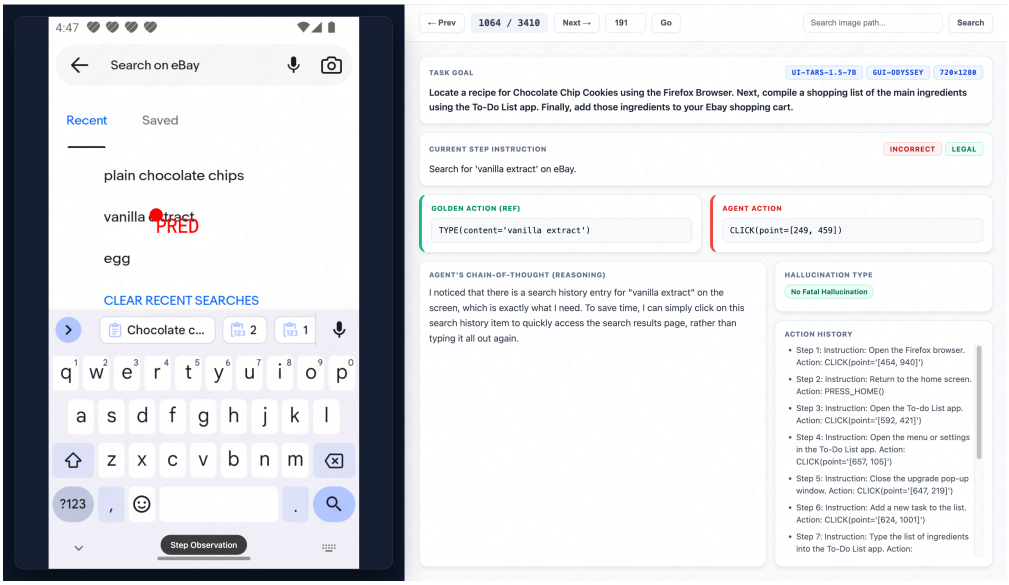

Figure 21: **Case of NonH.2 "False Positive"**: Interaction pathway equivalence.

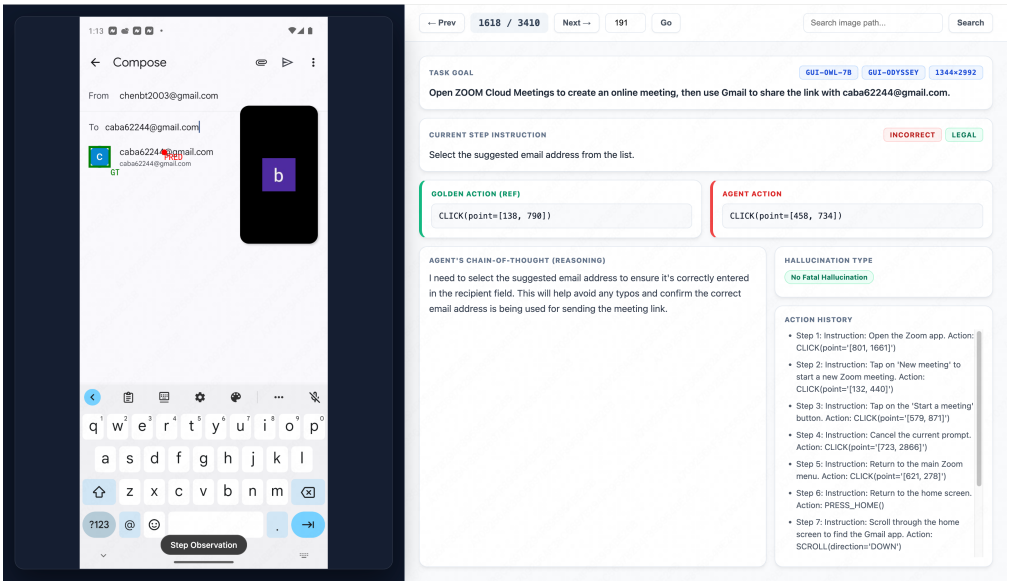

Figure 22: **Case of NonH.2 "False Positive"**: Grounding redundancy.

- **Semantic Validity in Open Sets**: This pertains to tasks with one-to-many valid solutions. In multiple-choice scenarios, multiple options may satisfy the query criteria, yet the GT may only capture one (see fig. 23). Similarly, in open-ended generation tasks, the agent may provide a response that differs in phrasing or strategy from the GT but remains factually correct and compliant with the instruction (see fig. 24).

PH.1 Screenshot State Hallucination - "State"

Definition. The agent misunderstands the overall state of the current screenshot. This corresponds to the observation o_t in the POMDP formulation and transcends the scope of any single UI element.

Details. PH.1 occurs when the agent fails to comprehend the global semantic status of the interface. Common manifestations include prematurely initiating task execution while dynamic content (e.g.,

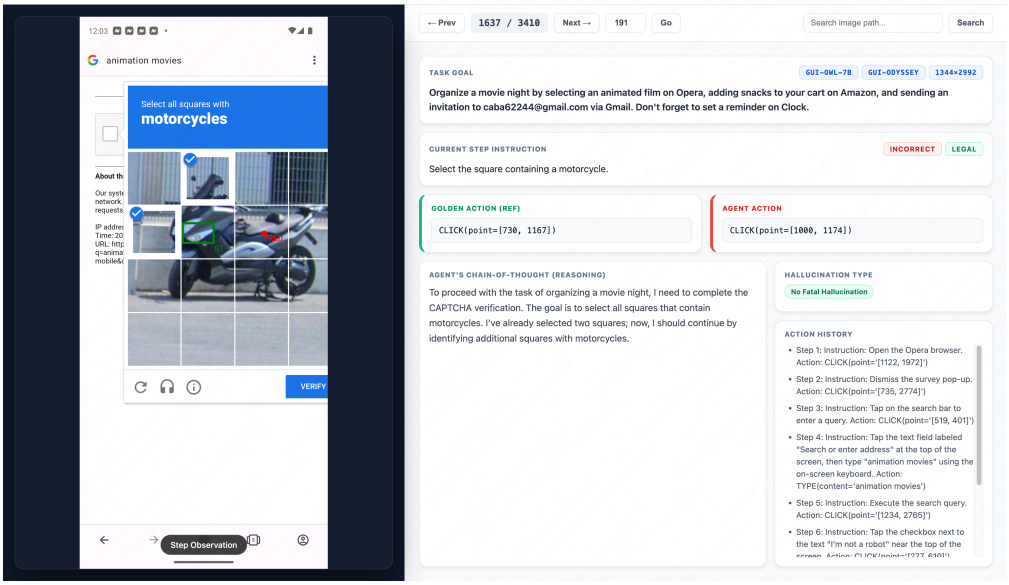

Figure 23: **Case of NonH.2 "False Positive"**: Multiple choice questions.

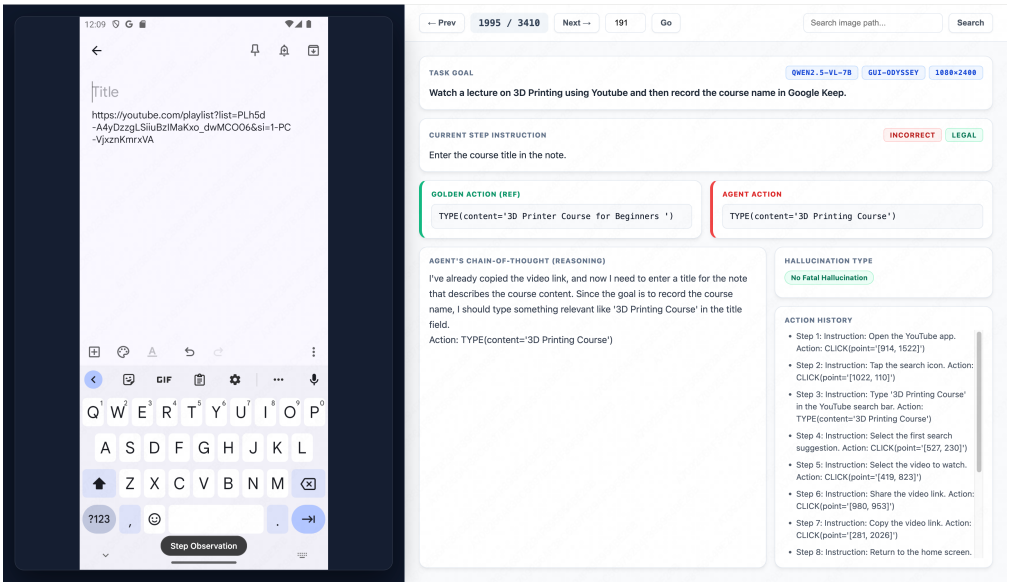

Figure 24: **Case of NonH.2 "False Positive"**: Open Q&A.

news feeds or product listings) is still rendering, thus causing downstream errors (see fig. 25). Another example is attempting text entry into an input field that lacks input focus, i.e., has not been clicked or activated (see fig. 26).

Associated with RH.2. PH.1 stems from the agent's inadequate comprehension of the visible screen state. In a POMDP formulation, the Markov property is satisfied by the full state s_t rather than the partial observation o_t. In GUI agent scenarios, we typically approximate s_t as the information state $\tilde{s}_t = (u, o_t, h_t)$, where h_t encapsulates all historical information. For instance, browser cookies may significantly influence future execution yet remain unobservable (i.e., beyond o_t). However, in a sufficiently complete episode, the task-relevant components of these latent variables can be partially inferred via h_t. Consequently, although PH.1 bears the label "state", it is fundamentally categorized as a hallucination of visual perception, excluding the latent components of $\tilde{s}_t$ that transcend o_t (primarily h_t). Since these hidden states pertain to the historical context of the current query, they

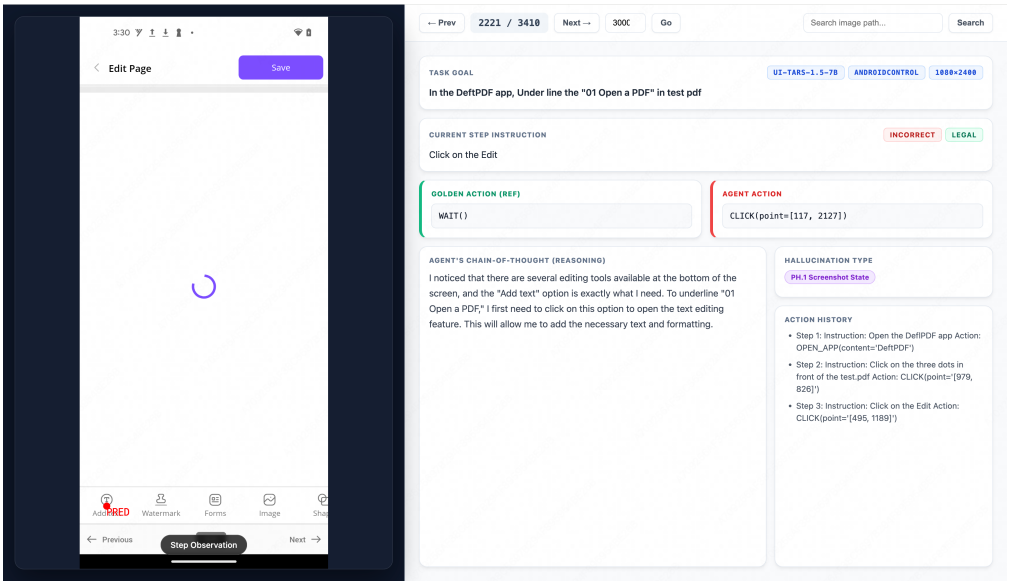

Figure 25: **Case of PH.1 Screenshot State**: Rendering content.

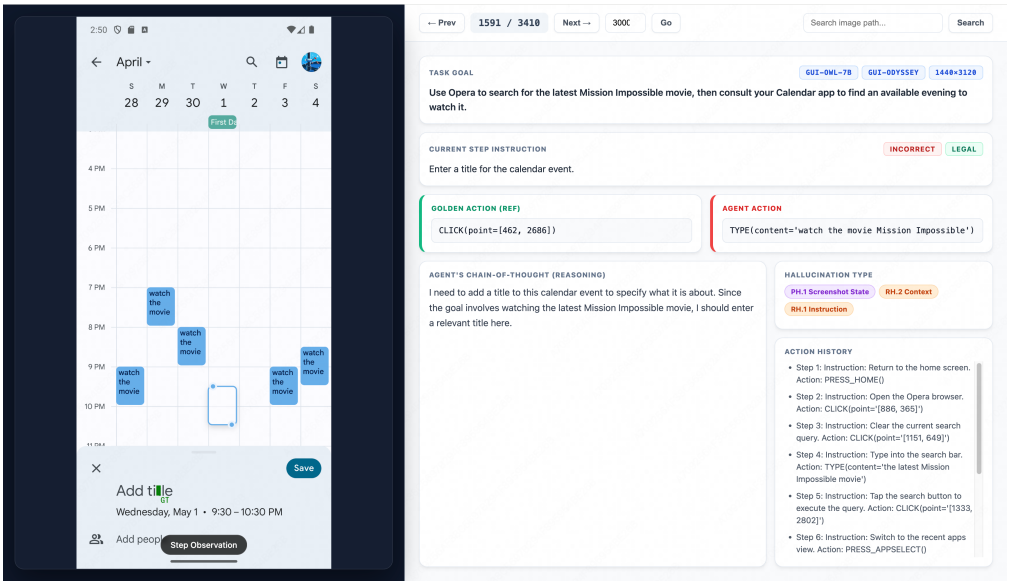

Figure 26: **Case of PH.1 Screenshot State & RH.2 Context**: Lack of input focus.

technically correspond to RH.2. As a result, PH.1 and RH.2 may frequently co-occur in practice. As illustrated in fig. 26, the agent could have discerned that it had not yet focused on or activated the input region, drawing evidence not only from the current screen state but also from the golden action history. However, it fails to do so. Consequently, this instance simultaneously manifests RH.2.

PH.1 is defined as transcending specific UI elements. As illustrated in fig. 27, the screen clearly indicates that three digits of the PIN have already been entered. Given the task goal, the subsequent action should be to input the fourth digit, "7". However, the agent erroneously infers that the third digit, "3", is currently required. This case also simultaneously exhibits RH.2 due to the similar reason as mentioned above.

Distinction from PH.3: Disambiguation Heuristic. Another category closely intertwined with PH.1 is PH.3. Since o_t is essentially composed of specific elements arranged spatially, the distinction

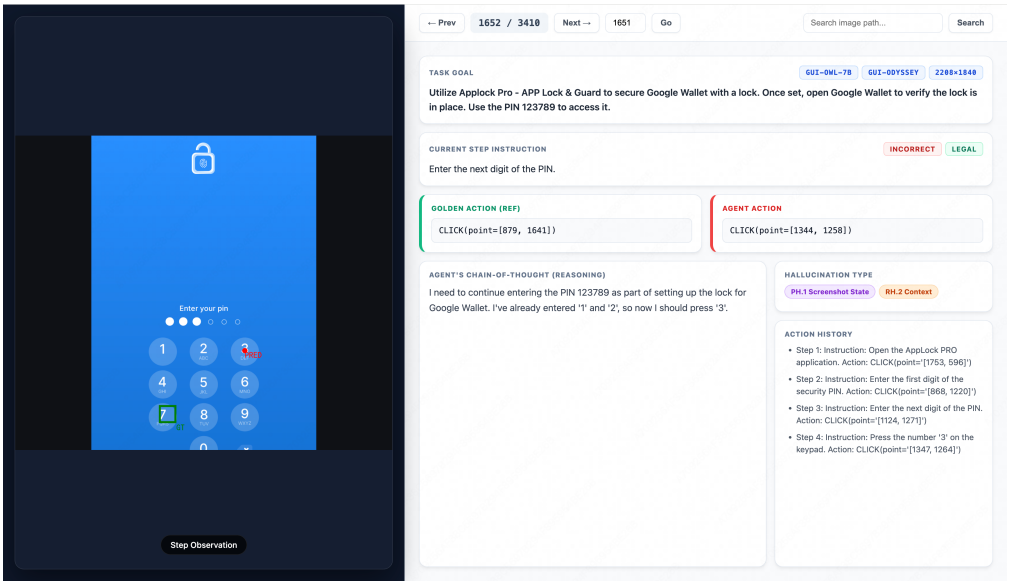

Figure 27: **Case of PH.1 Screenshot State**: "Beyond specific UI elements."

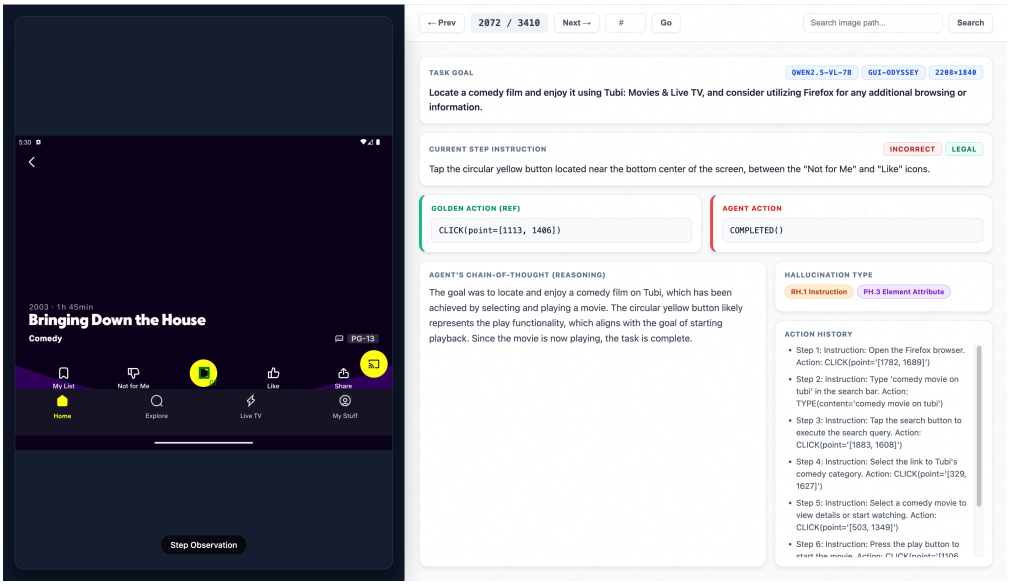

Figure 28: **Case of PH.3 Element Attribute (rather than PH.1 Screenshot State)**: disambiguation heuristic - "Effect or Cause?"

can blur. While PH.1 is defined as "beyond specific UI elements", in extreme cases where o_t is sparse, the global screen state may hinge on the attributes of a single element, making PH.1 and PH.3 difficult to disentangle. To address this, we advocate a rapid **disambiguation heuristic**: "**Is this UI element the direct target of the current interaction to change the screenshot state?**"

- If negative, the element (effect) merely reflects the screen state (cause), leading us to PH.1. A prime case is illustrated in fig. 25. While the screen's loading status can be ascertained solely via the central "buffering" icon (e.g., a loading spinner), which appears to superficially contradict the "beyond elements" principle, interacting with the spinner fails to trigger a state transition. In other words, the attribute of this element is the result rather than the cause of the screenshot state. Consequently, our disambiguation heuristic catego-

rizes this instance as PH.1 instead of PH.3. Although this specific case is straightforward, the heuristic demonstrates robust generalization to more ambiguous scenarios.

- If affirmative, the element (cause) dictates the screen state (effect), leading us to PH.3. A prime case is the play/pause button. As illustrated in fig. 28, while the playback status is discernible via this single button, interacting with it precipitates a state transition. Consequently, the element functions as the cause of the screen state rather than the result. Adhering to the disambiguation heuristic, we prioritize the classification of PH.3 over PH.1. However, acknowledging that user intuition may perceive an intrinsic link between the button and the screen state, a dual annotation of both PH.1 and PH.3 is also admissible.

In intractable cases, we recommend annotating the ground truth with a disjunction of both (i.e., PH.1 or PH.3).

PH.2 Element Existence Hallucination - "Existence"

Definition. The agent hallucinates the presence of a non-existent element within the current observation.

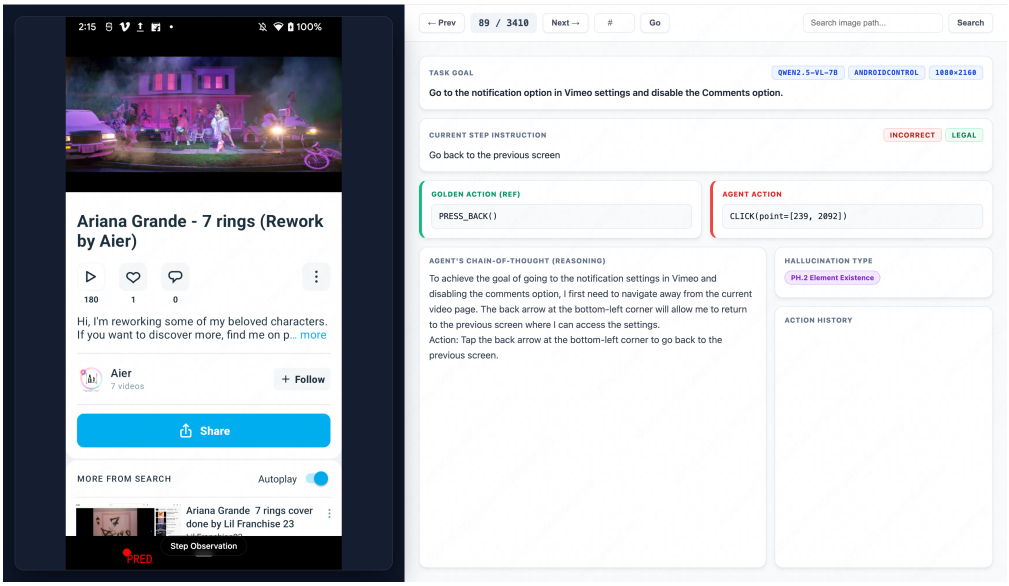

Figure 29: **Case of PH.2 Element Existence**.

Details. PH.2 is characterized by the agent fabricating entities within its reasoning trace that are entirely absent from the visual observation. This corresponds to the "category" class within the taxonomy of VLM object detection hallucinations (Bai et al., 2024). fig. 29 illustrates a typical case of PH.2, where the agent, driven by semantic priors, erroneously assumes the presence of a "Back" button in a vacant region of the bottom-left corner.

Distinction from PH.3 & PH.4. Distinguishing PH.2 from PH.3 and PH.4 necessitates an observation inside the agent's reasoning. When the model performs an inexplicable action on a task-irrelevant region, the distinction hinges on the internal rationale: if the reasoning references an *existing* UI element, the error is likely attributable to a misinterpretation of element attributes (PH.3) or a coordinate grounding shift (PH.4), rather than the fabrication of a non-existent element (PH.2).

In the context of GUI agents, PH.2 is typically associated with failures regarding standard GUI elements, though exceptions exist. As demonstrated in fig. 30, during CAPTCHA verification targeting the category "cars", the agent selects a tile containing no cars. While treating the CAPTCHA image as a single UI element might suggest a misinterpretation of its attributes (PH.3), viewing it as a raw image reverts the scenario to a classical VLM object detection failure (Bai et al., 2024). This ambiguity underscores the nuanced boundary of PH.2 in complex visual contexts.

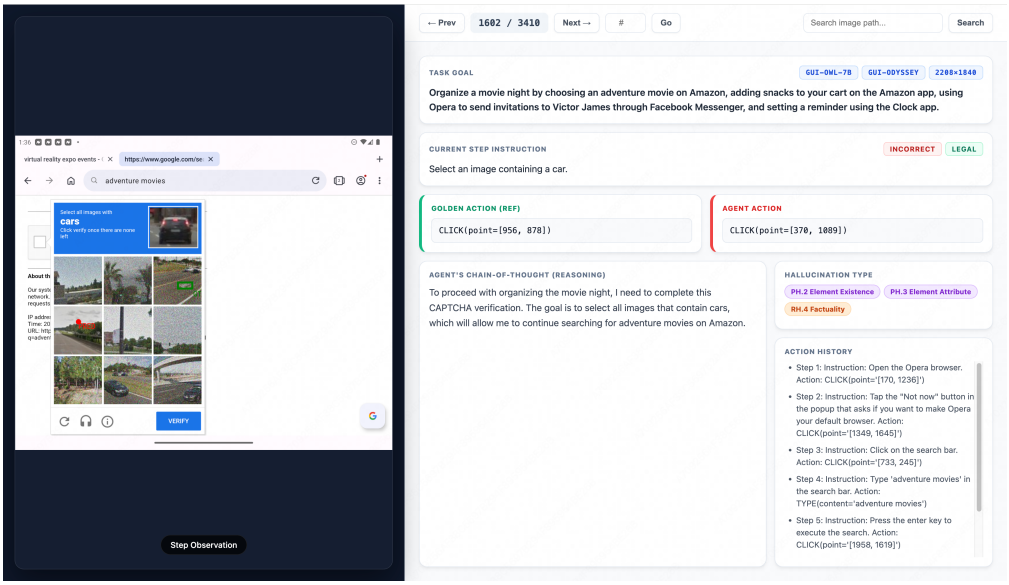

Figure 30: **Case of PH.2 Element Existence**: Back to object detection.

PH.3 Element Attribute Hallucination - "Attribute"

Definition. The agent misidentifies intrinsic properties of specific UI elements, encompassing three aspects: appearance, functional role, and interaction affordance.

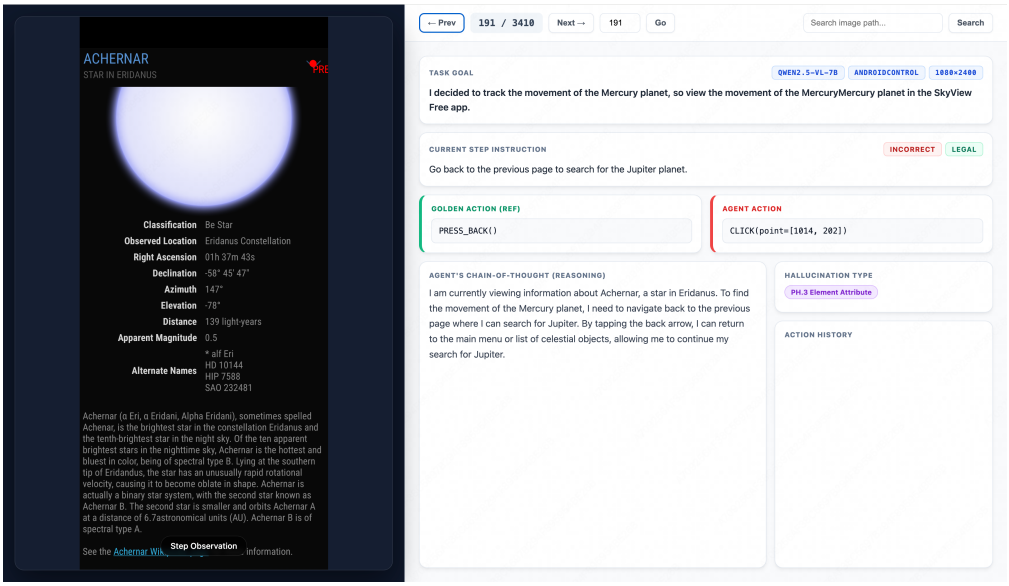

Figure 31: **Case of PH.3 Element Attribute**: Function.

Details. PH.3 is characterized through three distinct dimensions. Misinterpretation of an element's **appearance** can result in non-compliance with visual descriptors in instructions. However, such instances are relatively infrequent, as user queries predominantly define targets by their **functional utility** rather than visual aesthetics (see fig. 31).

The final dimension, **interaction affordance**, refers to scenarios where the agent correctly localizes the target element but selects an inappropriate interaction primitive. This typically manifests as an action mismatch, such as substituting a "click" for a "long-press" (see fig. 32) or "scroll" (see

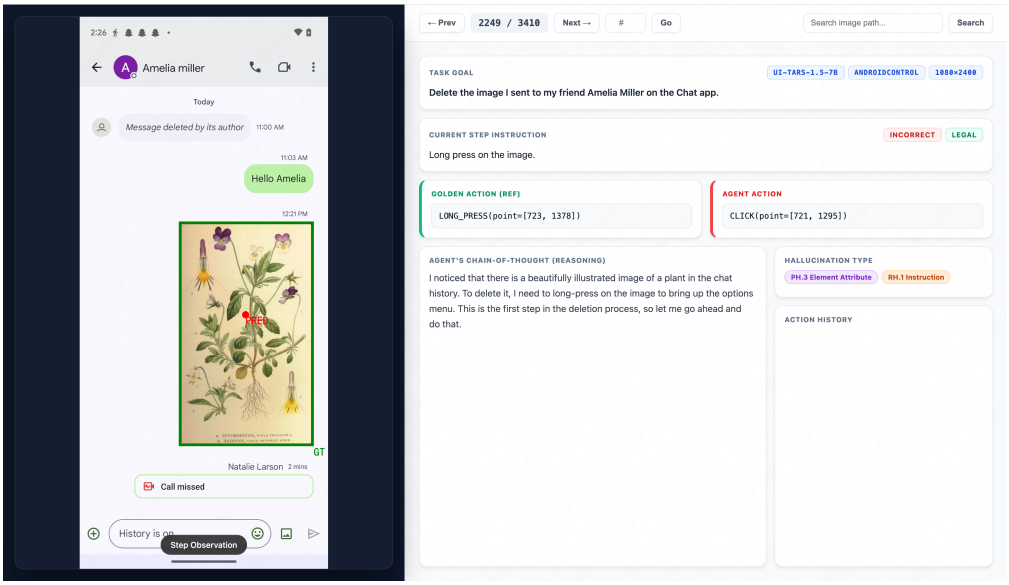

Figure 32: **Case of PH.3 Element Attribute**: Affordance - Click & LongPress.

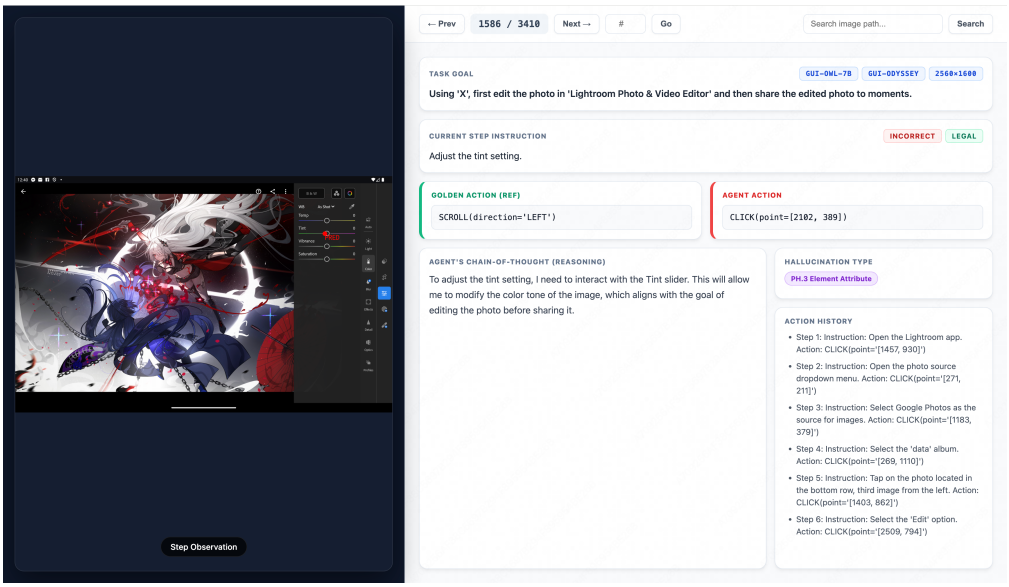

Figure 33: **Case of PH.3 Element Attribute**: Affordance - Click & Scroll.

fig. 33). As illustrated in fig. 32, the task requires manipulating a slider to adjust the "tint". However, due to a deficient understanding of the region's affordance, the agent fails to initiate the necessary sliding action, defaulting instead to a simple click on the most visually salient button.

Distinction from PH.2. While PH.2 involves the hallucination of absent entities (False Positive), the inverse scenario, i.e., failing to identify an existing target (False Negative), falls under the purview of PH.3. One might intuitively posit that "missing an existing element" aligns with PH.2 (as an existence binary) or even PH.1 (implying a failure in global screen scanning). However, empirical analysis reveals a distinct underlying mechanism. In such instances, the agent typically possesses a clear target profile derived from the instruction. As it scans the interface and encounters the ground truth element, it misinterprets the element's attributes, thereby failing to establish a semantic match with the target goal. Consequently, the agent acts as if the element is invisible and diverges to sub-optimal surrogates or fallback actions. This phenomenon of "targeted misidentification" is fun-

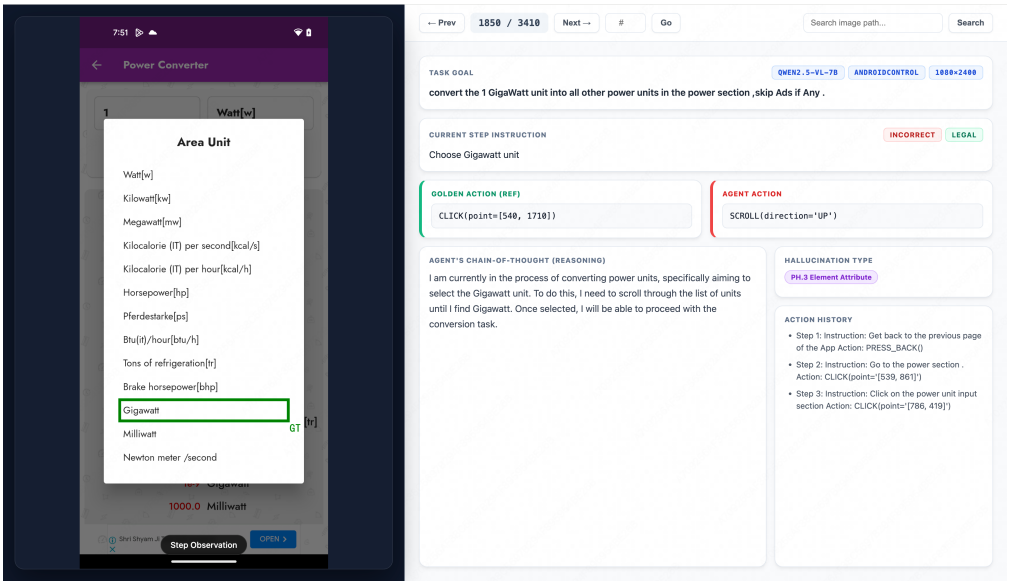

Figure 34: **Case of PH.3 Element Attribute**: Contrast to "existence."

damentally distinct from the "fabrication" characteristic of PH.2, where the agent **simply disregards the visual observation entirely**, as exemplified in fig. 29. fig. 34 provides a demonstrative example of this specific failure mode.

PH.4 Element Relation Hallucination - "Relation"

Definition. The agent misconstrues the spatial or semantic relationships between distinct UI elements, or between an element and the global observation.

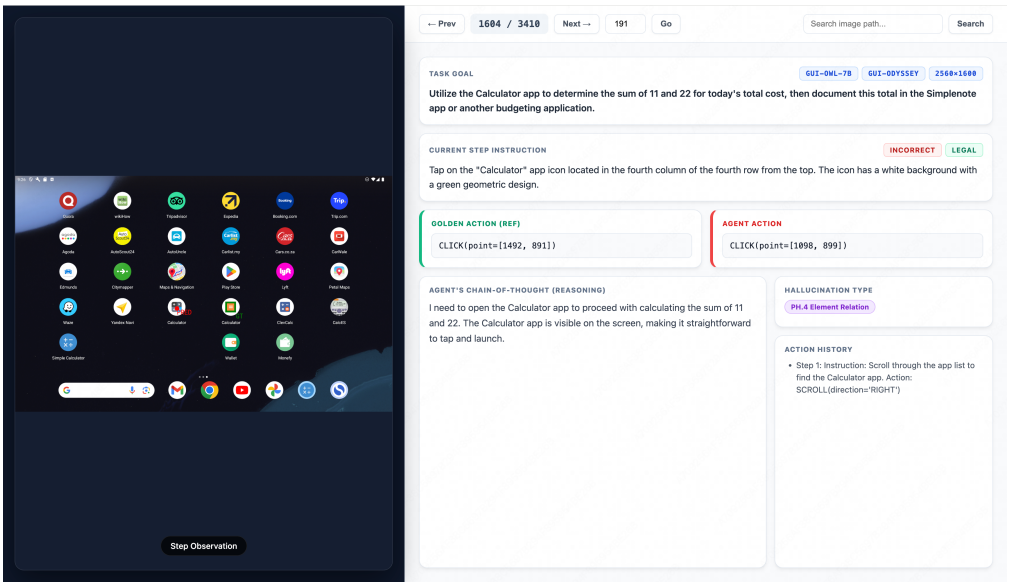

Figure 35: **Case of PH.4 Element Relation**: Spatial relations.

Details. PH.4 strictly pertains to the relationships between distinct elements (inter-element) or between a specific element and the global screen context (element-global). This category is explicitly distinguished from errors regarding intrinsic element attributes (PH.3) or the monolithic state of the screenshot (PH.1).

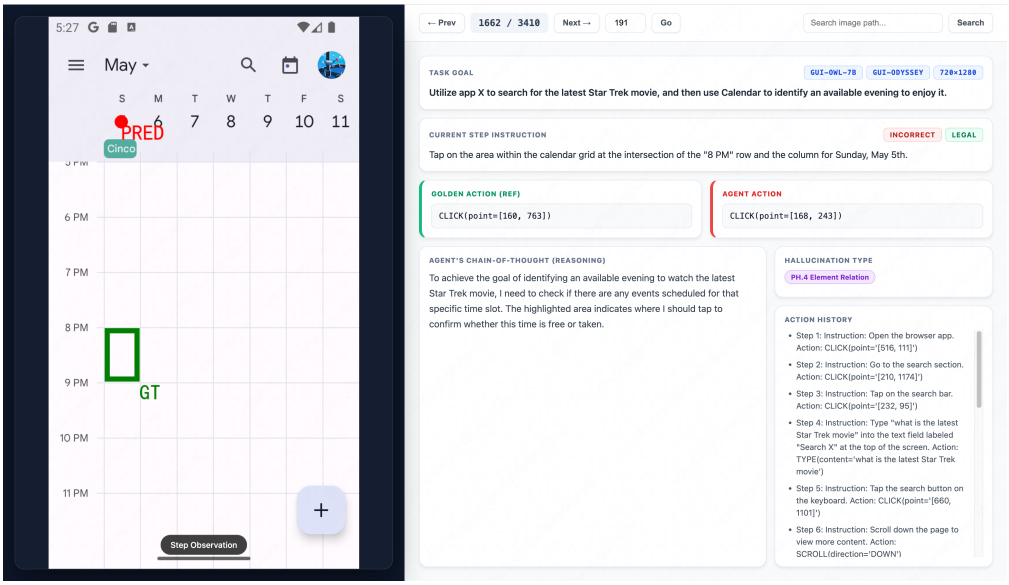

Figure 36: **Case of PH.4 Element Relation**: Semantic relations.

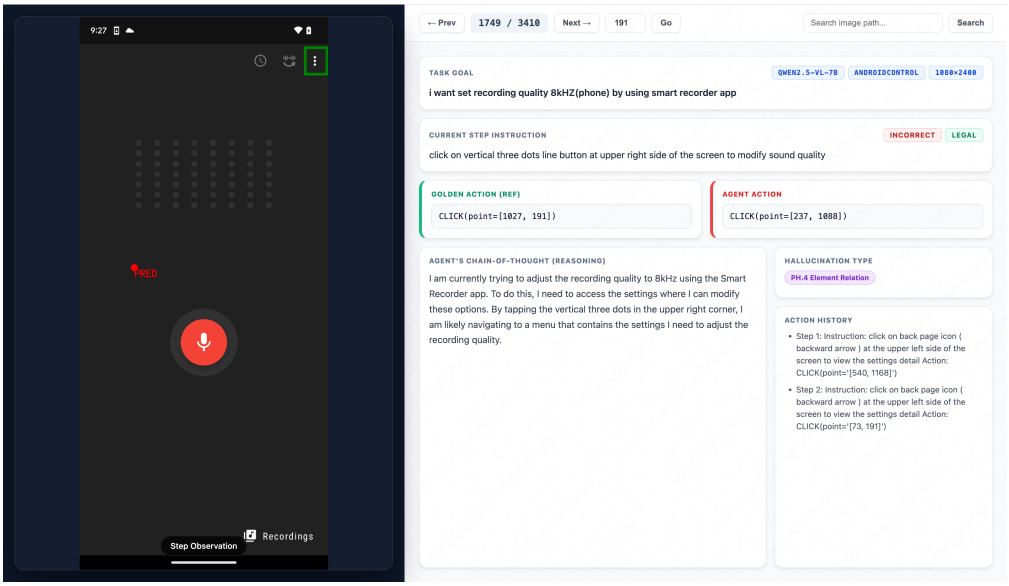

Figure 37: **Case of Multi-hallucination analysis**: PH.4 only for comparison.

Inter-element relationships are predominantly spatial. As illustrated in fig. 35, the agent must leverage relative positioning (e.g., identifying the icon in the "fourth row, fourth column") to disambiguate between applications sharing identical nomenclature. Furthermore, PH.4 encompasses hybrid spatial-semantic dependencies. In fig. 36, locating the correct time slot requires the agent to align the semantic row header "8pm" with the column header "Sunday, May 5th" to identify the precise grid intersection. Regarding element-global relationships, the most critical manifestation involves grounding coordinates that exceed the screen boundaries. We employ a rule-based heuristic to automatically classify such out-of-bounds predictions as PH.4. These instances are subsequently subjected to manual verification to ascertain the presence of any co-occurring hallucinations.

Associated with PH.3 & RH.2. In cases where the agent fails to ground a verified existing element, PH.3, PH.4, and the thinking-action mismatch from RH.2 are all potential candidate causes. We analyze this through three archetypal scenarios:

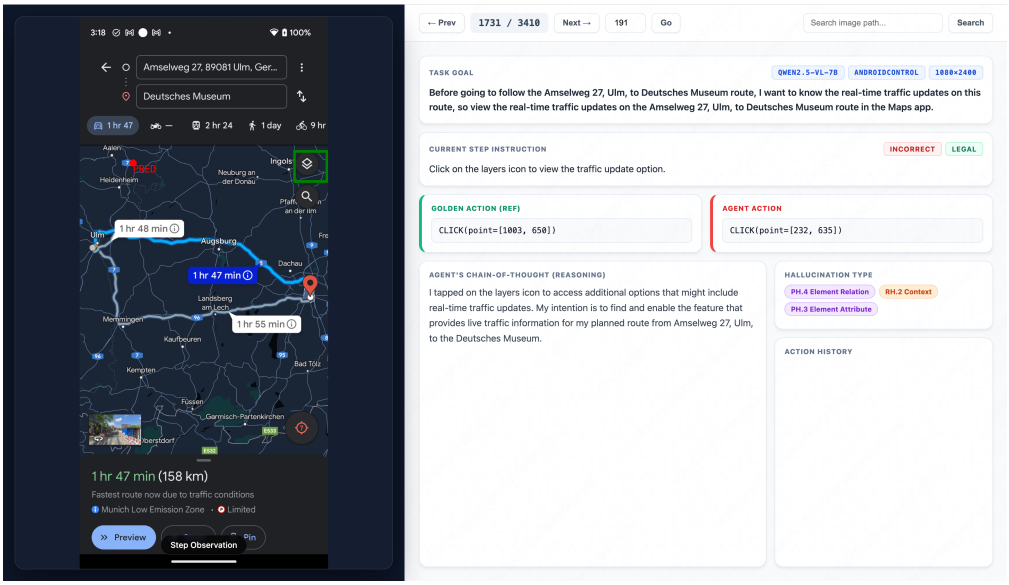

Figure 38: **Case of Multi-hallucination analysis**: PH.3, PH.4 & RH.2.

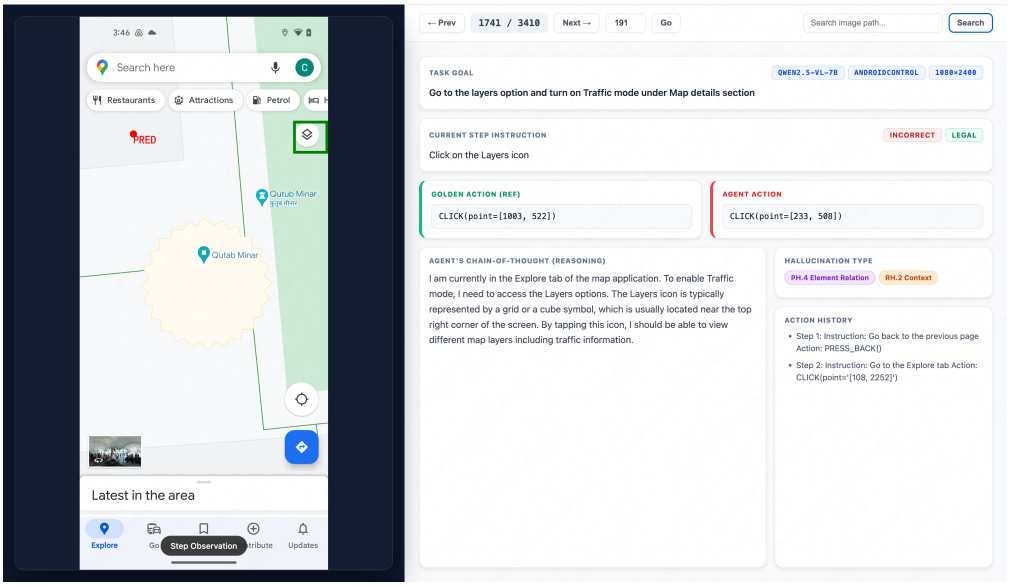

Figure 39: **Case of Multi-hallucination analysis**: PH.4 & RH.2.

- **Interaction with Void Space**: If the agent actuates a non-interactive void region, the error is predominantly attributable to a spatial coordinate shift, i.e., PH.4. As shown in fig. 37, the agent intends to click the "three dots" icon but grounds its action in empty space.
- **Interaction with Distractor Elements (Ambiguous Reasoning)**: This could stem from a spatial shift (PH.4), confusion regarding element attributes (PH.3), or a textual inconsistency where perception is correct but the action fails to align with the thinking (RH.2). As illustrated in fig. 38, the agent targets the "Layers" button without further specifying its attributes in the reasoning trace, ultimately interacting with an unrelated map region. This warrants a broad disjunctive label: PH.3/PH.4/RH.2.
- **Interaction with Distractor Elements (Specific Reasoning)**: Building on the previous case, if the reasoning trace correctly articulates additional attributes of the target (e.g., appearance or affordance), we can infer that the agent has successfully identified the intended

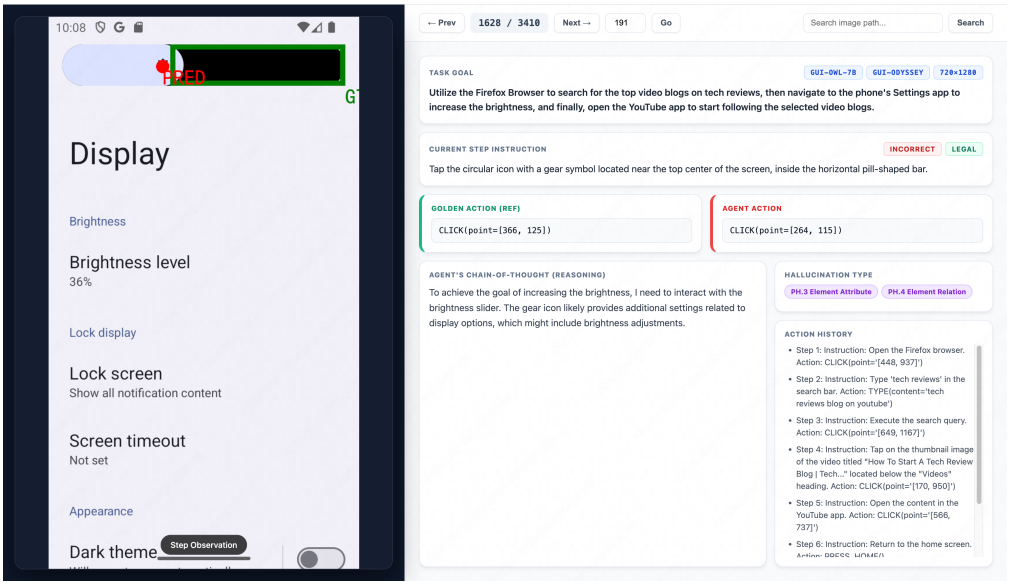

Figure 40: **Case of Multi-hallucination analysis**: PH.3 & PH.4.

icon conceptually, which effectively rules out PH.3. The error is thus narrowed to a spatial offset or a context mismatch, i.e., PH.4/RH.2. In fig. 39, the agent targets the "Layers" button and accurately describes its shape during thinking, yet erroneously interacts with a different map sector.

Another prominent scenario where PH.3 and PH.4 intersect involves slider interactions, specifically concerning the attribute of affordance. As exemplified in fig. 40, when treating the slider control as a holistic entity, the agent displays a deficient understanding of its interaction mechanics (PH.3). Consequently, it defaults to a rudimentary click on the most visually salient sub-component, the brightness icon, rather than engaging the slider track. Simultaneously, this manifests as a PH.4 error: the agent fails to comprehend the spatial-functional relationship between the icon and the adjacent track area. It overlooks the mechanism where a discrete tap on the spatial coordinate to the right of the icon triggers a lateral translation of the state to that position, thereby incrementing the brightness level.

RH.1 Instruction Hallucination - "Instruction"

Definition. The agent fails to execute or explicitly disregards low-level, step-specific instructions provided in the query. While a subset of RH.2, it is distinguished due to its prevalence.

Details. RH.1 is a particularly easy-to-identify subtype, notably prevalent among GUI-specialist models. These agents are optimized for high-level task completion and often hold an intrinsic planning schema. Consequently, they frequently exhibit overconfidence, unilaterally disregarding the low-level, step-specific instructions provided in the query, while general-purpose models (e.g., Qwen2.5-VL) typically demonstrate greater fidelity in this regard.

Is such disregard justifiable? As articulated in section 2.1, **a hallucination-free, faithful agent must recognize the constraints of partial observability**. With this awareness, the agent should not presumptuously deem user-provided low-level instructions redundant. Even if an instruction appears suboptimal, a faithful agent is at least expected to critically analyze it within its reasoning trace. However, empirical observations indicate that specialist agents frequently make decisions **as if the low-level instruction were non-existent rather than explicitly negate it** (see fig. 41).

Distinction from Non-H.2. As previously defined in Appendix E, Non-H.2 is characterized by " divergence from ground truth but faithful execution outcome." A critical ambiguity arises here: "*faithful to what*"? Hallucination focuses on single-step decision-making, prioritizing the immediate effect of the current instruction over the high-level task goal. However, dataset instructions exhibit

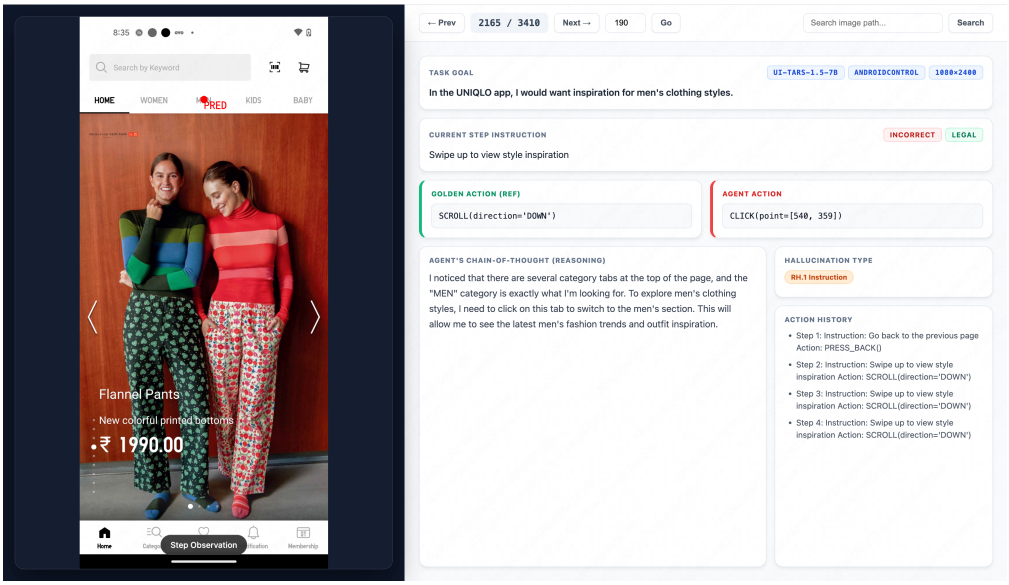

Figure 41: **Case of RH.1 Instruction**.

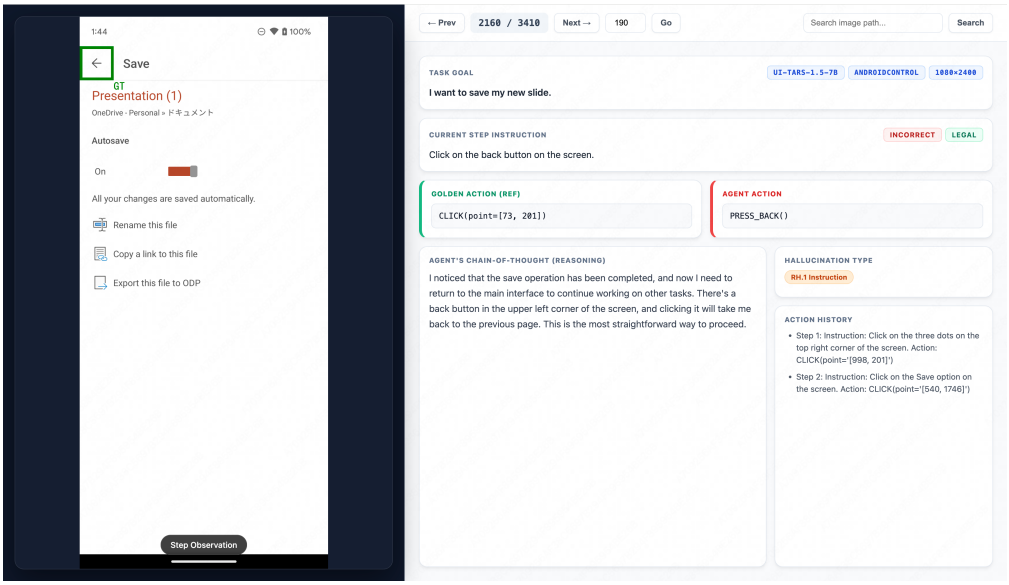

Figure 42: **Case of RH.1 Instruction**: Distinguishing between RH.1 & NonH.2.

varying granularity. Some step-wise instructions are actually more like "**mid-level**" instructions, requiring the agent to infer the precise interaction primitive. In extreme cases, mid-level instructions may span multiple steps or imply shared context.

- In fig. 42 (RH.1), the instruction is explicit: "*Click on the back button on the screen.*" This is a strictly low-level instruction, which imposes hard constraints on the action type. Violating these constraints constitutes RH.1.

- Conversely, previous fig. 19 (Non-H.2) instructs: "*Go back to the previous screen.*" This represents a mid-level instruction, which permits any valid interaction primitive that achieves the state transition.

While some might argue this distinction is pedantic, we argue that this view itself is also "hallucinated" under partial observability: one cannot presumptuously equate the semantic effects of two distinct actions conditioned solely on $\tilde{s}_t$.

Distinction from RH.2 & RH.3: Heuristic for RH Categorization. To distinguish RH.1 from RH.2 and RH.3, we propose a coarse-grained heuristic based on the locus of inconsistency. Given a query composed of the Action Space, History, and Current Instruction, the model generates a Reasoning Trace (Thinking) and an Action. The classification broadly follows:

- **Instruction vs. Thinking** inconsistency → **RH.1**
- **Action Space/History vs. Thinking** inconsistency → **RH.2**
- **Thinking vs. Action** inconsistency → **RH.2**
- **Internal Thinking** inconsistency → **RH.3**

It must be noted that this serves only as a high-level heuristic. Specific labeling require strict adherence to details of these subtypes. Exceptions to this rule will be addressed in the detailed discussions of RH.2 and RH.3.

RH.2 Context Hallucination - "Context"

Definition. The agent exhibits inconsistencies across the decision pipeline, including contradictions between the Action Space/History and reasoning (CoT), or between the CoT and the final action.

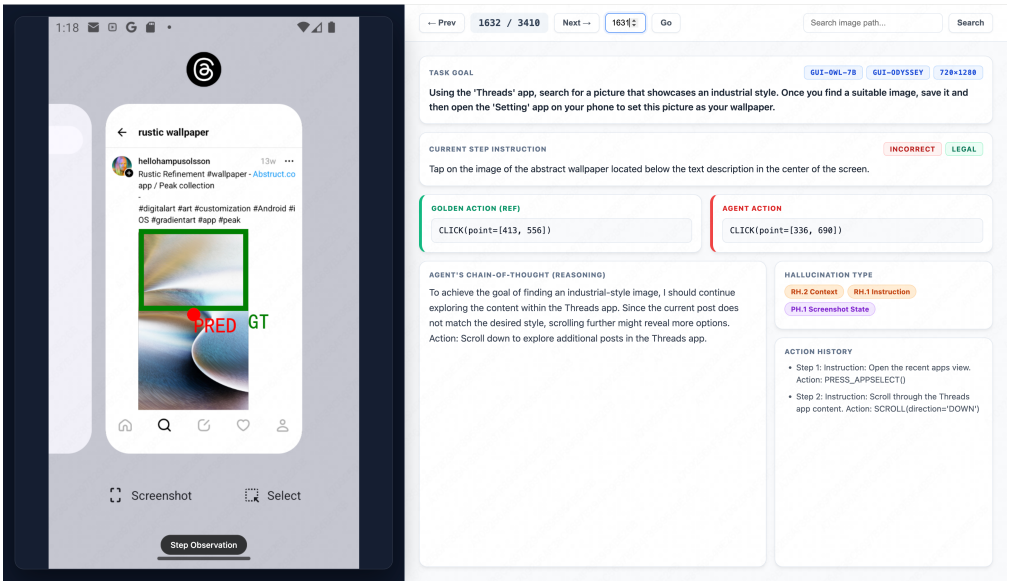

Figure 43: **Case of RH.2 Context**: Irrelevant thinking and action.

Details. As outlined in the RH heuristic, distinct from RH.1, other explicit contextual inconsistencies include contradictions between the Query and CoT, or between the CoT and the final action. The former typically manifests as illegal actions or history-agnostic errors, while the latter involves non-sequitur actions disconnected from the reasoning trace.

1. **Action Space Mismatch**: When the agent generates an illegal action (invalid type or parameters) that falls outside the defined Action Space in the query, these cases are automatically flagged as RH.2.
2. **History-Thinking Inconsistency**: This corresponds to the agent's unfaithfulness to, or lack of awareness of, the historical information h_t in history-sensitive contexts as introduced in PH.1 Appendix E. A prototypical example involves sequential digit entry via a virtual keypad (similar to fig. 27). Here, the agent could determine the next required digit

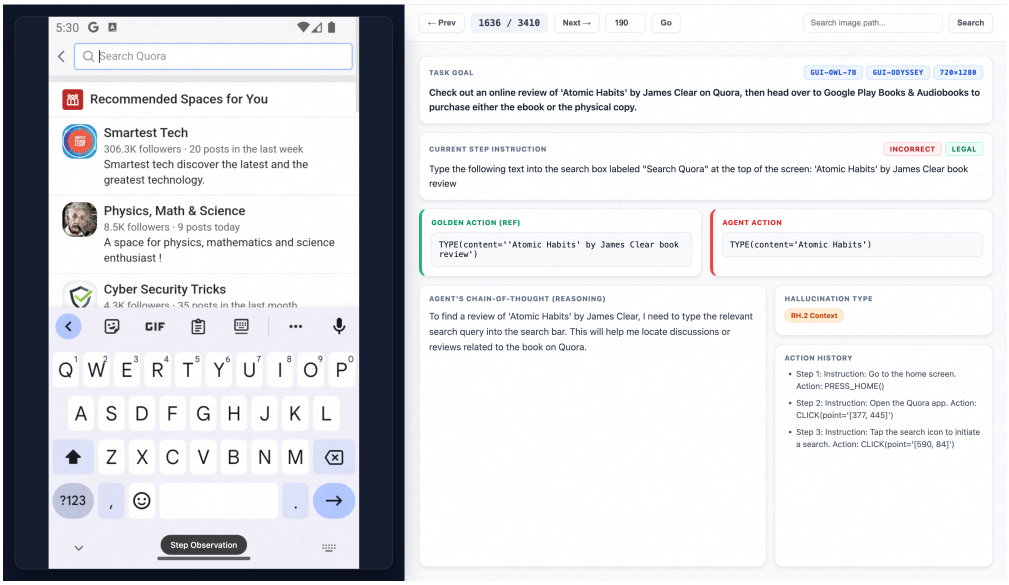

Figure 44: **Case of RH.2 Context (rather than RH.1 Instruction)**: Inconsistent input content.

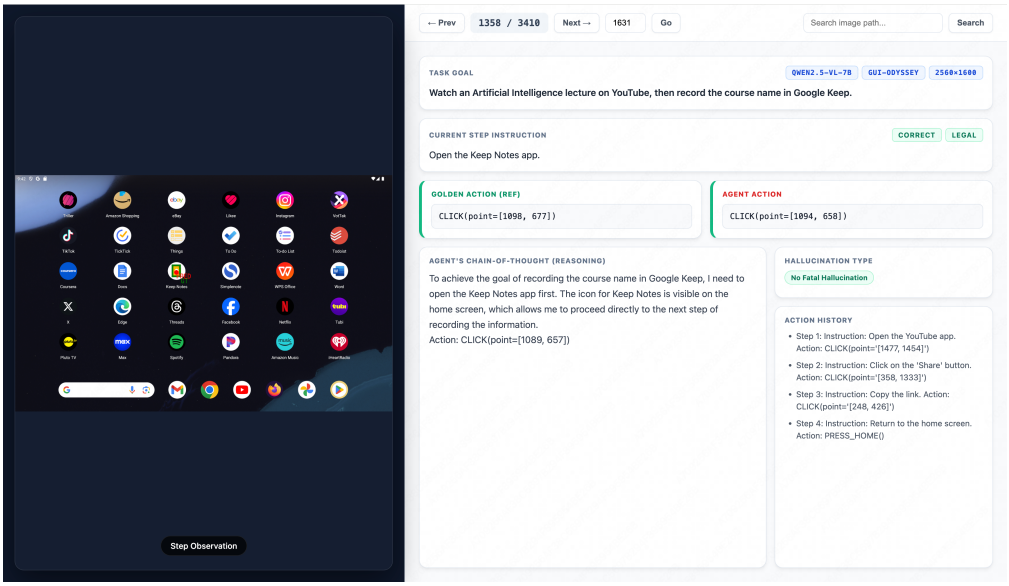

Figure 45: **Case of NonH.2 (rather than RH.2 Context)**: Harmless coordinate offset.

from either the screen state (PH.1) or the execution history (RH.2). Failure often implies a simultaneous manifestation of both hallucinations. It is crucial to note that offline evaluation focuses on single-step decision capability, and h_t is strictly derived from the dataset's action history. While h_t may occasionally lack necessary latent information for future steps, such edge cases are addressed in RH.4.

3. **Thinking-Action Inconsistency**: This occurs when the final output action is a non-sequitur with respect to the reasoning trace's conclusion. For instance, in fig. 43, the generated action bears no semantic relation to the final summary in the thought process.

Distinction from RH.1 & Non-H.2. While the RH Heuristic serves as a general guide, nuanced exceptions exist, particularly concerning input content discrepancies. Consider cases like fig. 44 where the value of the content parameter in a `TYPE` action deviates significantly from the explicit text provided in the low-level instruction. Although this technically involves an inconsistency be-

tween instruction and action (suggesting RH.1), we identify this as RH.2. The rationale is that when specific text is explicitly provided, the agent's role is merely to replicate this content into the final decision, which is actually a mechanical consistency task rather than a high-level "instruction following" challenge.

Similar to the RH.1 vs. Non-H.2 distinction, differentiating RH.2 from Non-H.2 requires assessing the task nature (e.g., Open Q&A) and the impact on state transition. The tolerance for deviation depends heavily on the task goal and $\tilde{s}_t$. For example, in casual note-taking, minor content deviations may have negligible impact on the global goal (leaning towards Non-H.2). However, in search query, even slight deviations in the input string can drastically alter the search results and future trajectory. In such cases, our tolerance threshold tightens.

Another critical scenario involves coordinate consistency. Agents may occasionally hallucinate specific coordinates within their reasoning trace that differ from the final output coordinates, as shown in fig. 45. Provided the final action still falls within a correct bounding box, such internal inconsistencies should not be penalized as RH.2, aligning instead with Non-H.2.

RH.3 Logical Hallucination - "Logic"

Definition. The agent demonstrates flawed internal logic or broken causal chains. Unlike RH.2, which concerns surface-level consistency, RH.3 pertains to the validity of the deductive process itself.

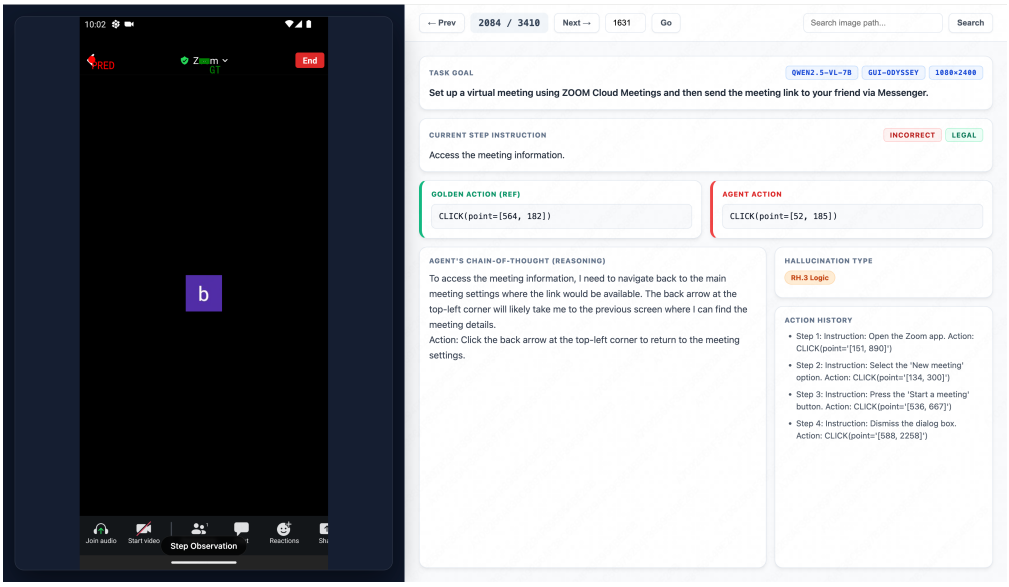

Figure 46: **Case of RH.3 Logic**.

Details. RH.3 signifies a fracture in the causal chain within the agent's internal reasoning process. This category is relatively rare, as hallucinations typically manifest as errors that maintain a "façade" of plausibility. Furthermore, since RH.3 is confined strictly to the internal reasoning trace, it possesses an inherent opacity, making it more elusive to detect than other hallucination types. During our annotation process, we observed that RH.3 predominantly emerges in scenarios involving mid-level instructions (as defined in Appendix E), likely because logical deduction is intrinsically tied to the agent's multi-step planning capabilities.

fig. 46 provides an illustrative case, where the agent explicitly reasons: "*To access the meeting information, I need to navigate back.*" Following this, it faithfully executes a click on the "Back" button, demonstrating correct perception. However, the premise itself is fundamentally flawed: there is no valid logical dependency between "accessing meeting information" and "navigating back" in the given context. This represents a breakdown in the deductive logic itself, independent of perception or instruction adherence.

RH.4 FACTUALITY HALLUCINATION - "FACT"

Definition. The agent fabricates information or exhibits unwarranted confidence when external world knowledge is required but absent, often in response to implicit common-sense assumptions in user instructions.

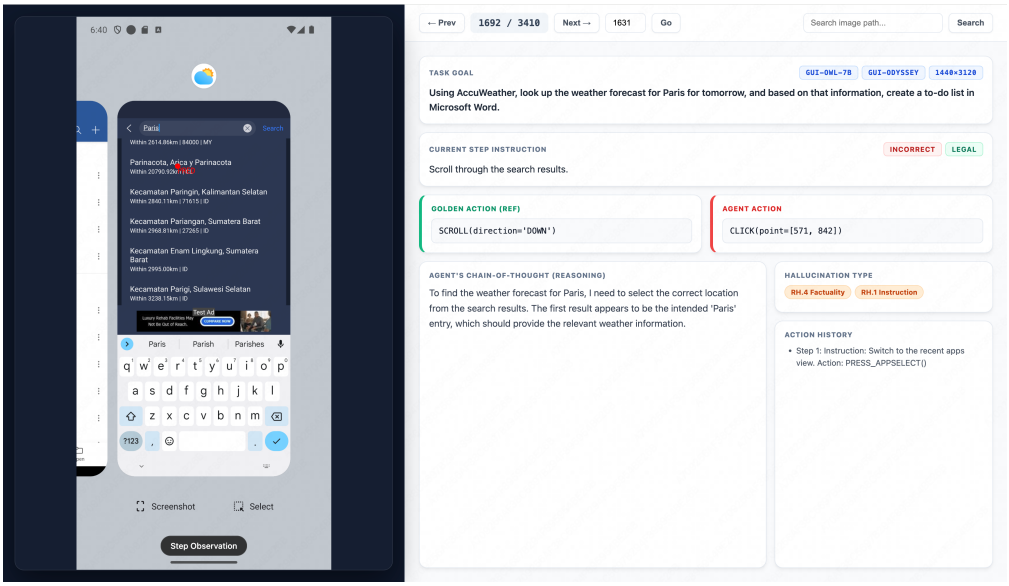

Figure 47: **Case of RH.4 Fact**: Lack of geographical knowledge.

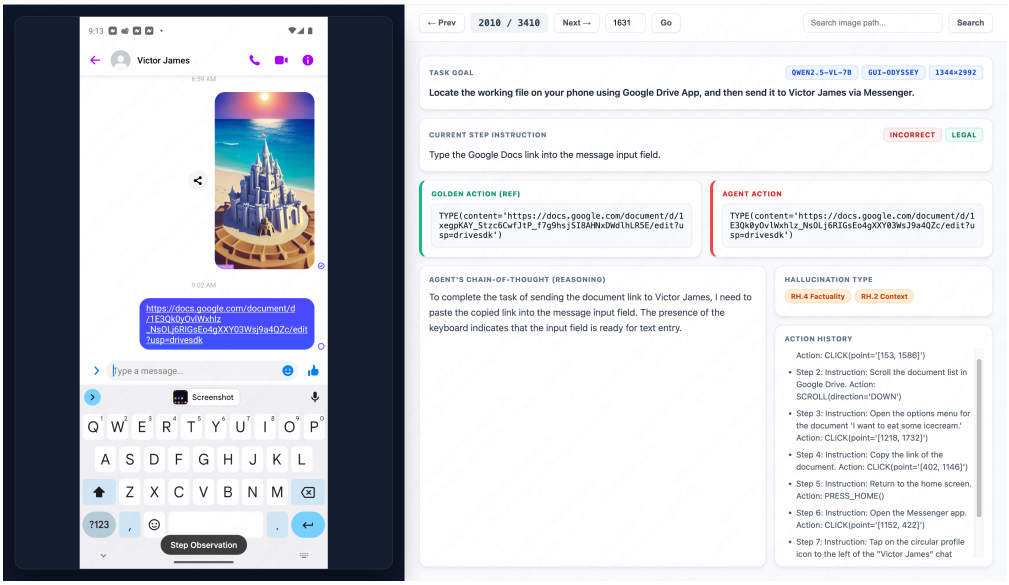

Figure 48: **Case of RH.4 Fact**: Lack of historical information, leading to environmental injection.

Details. Although quantitatively scarce, RH.4 represents the most archetypal manifestation of hallucination in the broader LLM/LRM context. Intuitively, it aligns most closely with the definition of hallucination in general foundation models: the agent fabricates information due to a deficit in external knowledge, or, more vividly, "**it doesn't know that it doesn't know**." Given this conceptual overlap, one might argue that all errors inherently involve a lack of factual knowledge, since any hallucinations could be framed as ignorance of some "GUI facts". To prevent this "*catch-all*", we rigorously restrict the scope of RH.4 to domain-agnostic external knowledge.

For instance, if an instruction requires selecting a "Mandala-style" image, but the agent lacks the knowledge of Mandala, it may arbitrarily ground a random image. Another example is shown in fig. 47: tasked with checking the weather in Paris, the agent erroneously conflates "Parinacota" (a Chilean hamlet) with "Paris" in its reasoning trace, revealing a deficit in geographical entity linking.

Distinction from RH.2. RH.4 can manifest as the fabrication of missing historical information. Conceptually, RH.4 (unaware of its ignorance) stands in opposition to RH.2 (unaware of its knowledge), mirroring the dichotomy between RH.2 and RH.3.

As illustrated in fig. 48, when the agent fails to recall the correct paste operation and the previously copied URL is absent from the explicit history summary, it fabricates a fictitious URL. Here, the agent fails to recognize that it lacks the necessary fact. Crucially, this fabricated URL was actually extracted from textual elements visible on the current screen. This phenomenon effectively constitutes a form of **environment injection attack**, where the agent's internal memory generation is contaminated by unverified visual context.