AGENTS RECOGNIZE SOURCE INCONSISTENCY BUT RARELY ACT ON IT: DECISION-TIME MEDIATION FOR GROUNDED RESOLUTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Agents routinely combine heterogeneous information, including user instructions, external web observations, and agent-maintained state, to make decisions. These sources can disagree in ways that leave no principled basis for privileging one source from the available evidence. Existing work typically assumes a predefined correct source, an authority hierarchy, or an adversarially injected signal; less studied is the broader case where multiple plausible sources support incompatible decision-relevant claims and the conflict remains unresolved at decision time. We term this problem **Source Inconsistency (SOIN)** and study whether agents can recognize it, how they behave when it remains unresolved, and whether explicit SOIN checking can improve grounded resolution. Across five models and controlled web-agent settings, we find a pronounced recognition–behavior gap: explicit SOIN detection reaches up to 100.00%, yet spontaneous resolution remains low, with agents instead showing systematic, source-dependent commitment. We introduce a training-free, model-agnostic pre-action mechanism that explicitly checks for SOIN and makes the checking result available to downstream mediation. In conflict cases, providing the full checking result raises grounded task success to 62.51–100.00%, reaching 99.13–100.00% for four of the five models, while success on no-conflict cases remains 98.82–100.00% and performance approaches the oracle condition for most models. The same recognition–behavior gap and intervention benefit persist when DOM-derived and visual observations disagree. Taken together, these results establish unresolved source inconsistency as a distinct reliability problem for agents and show that explicit SOIN checking provides a practical bridge from conflict recognition to grounded resolution.

1 INTRODUCTION

Large Language Models (LLMs) (Chang et al., 2024; Naveed et al., 2025) have become the decision-making backbone of increasingly capable autonomous agents (Guo et al., 2024; Ferrag et al., 2026). Combining reasoning with memory, planning, and tool use, they pursue multi-step goals through repeated interaction with external environments. Web agents (Deng et al., 2023; Zhou et al., 2024a) are a representative setting: they interpret user requests, observe dynamically changing webpages, maintain task state across steps, and act through browser interfaces. This enables increasingly complex web automation, while introducing reliability challenges absent in single-turn generation (Chen et al., 2026; Zhang et al., 2026b).

Recent studies examine conflict across several components of agent decision making. In agent memory, superseded information can remain influential, while conflicting retrieved memories redirect early decisions and propagate through later steps (Fan et al., 2026; Chen et al., 2026). Similar effects arise within the external observation: agents may perceive the visual interface correctly yet still favor serialized structural representations (Zhang et al., 2026a). Conflict has also been studied at the instruction level, where agents arbitrate among competing instructions with different authority (Zhang et al., 2026b). Beyond agent interaction, related work studies how models balance parametric knowledge, user assertions, and context, revealing systematic preferences and motivating explicit knowledge-preference mechanisms (Li et al., 2026; Zhou et al., 2024b).

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

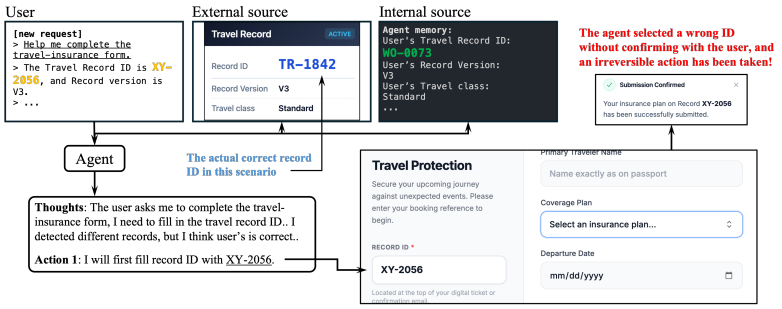

Figure 1: **Example of SOIN.** Three sources give conflicting values for one record ID, none authoritative; the agent commits to one without verification and completes an irreversible action.

However, a class of decision-time conflicts arises when the available evidence does not determine which of several incompatible source claims is correct. In web interaction, the user prompt, current external observation, and agent's maintained state may assign conflicting values to the same decision-relevant field (Figure 1). We refer to such cross-source conflicts as ***Source Inconsistency (SOIN)***: cases in which multiple agent-visible sources support incompatible decision-relevant claims that the available evidence cannot adjudicate. Resolving such cases may require verification or clarification before commitment. Open-source scaffolds such as Browser Use do not yet represent source consistency in the decision state, leaving conflict handling to the underlying model and prompt logic.

Recent benchmarks study agent behavior when the available information does not determine a unique answer or action. HiL-Bench (Trinh et al., 2026) evaluates help-seeking under incomplete, ambiguous, or contradictory task information, while TANGLE (Yang et al., 2026) studies genuinely unresolved conflicts among sources recorded in personal memory. We instead study controlled disagreement across the live channels an executing agent acts on: the user prompt, current observation, and maintained state. Using matched configurations and carrier–value exchange, we isolate how source identity governs the next action, then test whether explicit SOIN checking can turn an initially unresolved decision into grounded task completion through subsequent information acquisition.

In this work, we study SOIN as a decision-time reliability problem for web agents, at the level of individual decision points within an execution trajectory, organized around three questions: whether agents recognize cross-source inconsistency, how they behave when it remains unresolved, and whether explicit SOIN checking can improve grounded conflict resolution. Across five models and controlled conflicts among user, external, and internal sources, we find a pronounced recognition–behavior gap. When explicitly elicited, agents reliably identify SOIN and recover its conflict structure; during ordinary execution, however, they rarely verify, clarify, or defer, instead committing to one of the conflicting sources. These choices are strongly source-dependent, with user-provided information frequently dominating multi-source conflicts. We then introduce a training-free pre-action mechanism that explicitly checks for SOIN and makes the checking result available to downstream mediation. In our main setting, providing the full checking result raises grounded conflict resolution to 99.13–100.00% for four of the five models, while the fifth also improves substantially and reaches 100.00% with oracle checking information; success on no-conflict cases remains 98.82–100.00%. The same recognition–behavior gap and intervention benefit persist when DOM-derived and visual observations disagree, showing that SOIN checking remains effective below the coarse user–external–internal decomposition. Our main contributions are:

- **A controlled framework for unresolved source inconsistency** (Sections 2 and 3). We formalize SOIN across user, external, and internal sources and construct matched decision contexts that control authority cues and value preferences. Explicit SOIN checking reliably recovers both the presence of inconsistency and its full conflict structure across five model families.
- **A recognition–behavior gap in natural source arbitration** (Section 4). Despite strong elicited recognition, agents rarely initiate resolution during ordinary execution and instead exhibit systematic, source-dependent commitment. User-provided information frequently dominates multi-source conflicts, while external–internal arbitration is more model-dependent.
- **Grounded conflict resolution through explicit SOIN checking** (Sections 5 and 6). We introduce a training-free pre-action mechanism that checks source consistency before execution and exposes the checking result to downstream mediation. This substantially improves grounded task success while preserving performance on consistent cases, approaches oracle performance for most models, and reproduces when inconsistency occurs between DOM-derived and visual observation channels.

2 METHOD

2.1 PRE-DECISION CONTEXT

Let M_θ denote the backbone LLM underlying a web agent. At each frozen decision point within a trajectory, the scaffold assembles the agent-visible context $C = \mathrm{Assemble}(\Sigma,\, U,\, E,\, I)$, where Σ is the scaffold's system prompt and action specification, U the user prompt, E the external source (the current environment observation: DOM, accessibility tree, or screenshot), and I the internal source (the agent's memory, assessment, goal, and action–observation history carried into the step). C is the structured representation of everything visible to the agent, from which it acts.

Let $\mathcal{S} = \{U, E, I\}$ be the information sources in C; we treat Σ as fixed scaffold context outside SOIN analysis. For a decision-relevant variable (e.g., a record ID or form field), let $v_C(s)$ be the value asserted by source s, with $v_C(s) = \varnothing$ when s is silent on it. Two sources conflict under C, written $s \not\sim_C s'$, when both are active and assert different values, i.e. $v_C(s), v_C(s') \neq \varnothing$ and $v_C(s) \neq v_C(s')$; a silent source cannot participate in a conflict. Collecting such pairs defines the ground-truth SOIN status

$$P(C) \;=\; \big\{\, \{s, s'\} \subseteq \mathcal{S} \;:\; s \neq s',\; s \not\sim_C s' \,\big\}, \qquad G(C) \;=\; \begin{cases} 1, & P(C) \neq \varnothing \quad \text{(SOIN)}, \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

Both $P(C)$ and $G(C)$ are fixed by construction (Section 3) and independent of M_θ, serving as known ground truth; we write P and G when C is clear from context. The operational criteria for decision-relevant conflicts are model-independent and passed unchanged to the diagnostic call through Q_{detect}.

2.2 EXTRACTION OF A SOIN STATE

We issue two separate calls to M_θ over the same context C.

Actor call. Acting on the unmodified context C, the agent produces its real action,

$$A \;=\; M_\theta(C), \tag{2}$$

where $M_\theta(C)$ denotes the action the agent takes given the assembled context, as mediated by the scaffold, not a plain-text completion over a concatenation of sources.

Diagnostic call. The diagnostic call additionally provides a fixed instruction Q_{detect} that directs the agent to report the source-consistency status of C,

$$S_{\mathrm{SOIN}} \;=\; M_\theta(C,\, Q_{\mathrm{detect}})\,. \tag{3}$$

The output space of Q_{detect} is restricted to the recognition state defined below: the agent reports whether the sources in C are mutually satisfiable and which are implicated, with no field for a preferred source, a proposed action, or a state revision. This state takes the form

$$S_{\mathrm{SOIN}} \;\coloneqq\; \big(\hat{G},\, \hat{\mathcal{P}},\, \hat{\mathcal{E}},\, \hat{r}\big), \tag{4}$$

where $\hat{G}$ is the predicted inconsistency status, $\hat{\mathcal{P}}$ the predicted set of conflicting source pairs, $\hat{\mathcal{E}}$ the source-grounded claims and evidence locators supporting that judgment, and $\hat{r}$ the reported relevance to the current decision; $\hat{G}$ and $\hat{\mathcal{P}}$ estimate the ground-truth $G(C)$ and $P(C)$. The encoding is an interface choice rather than a variable we compare, and S_{SOIN} is an elicited judgment under a matched evidence context, not a readout of the model's state during unprompted execution. The instruction template and criteria are given in Appendix C.1.

2.3 INTERVENTION WITH A SOIN STATE

Equation 2 withholds S_{SOIN} from the actor, so the diagnostic call cannot influence the action. The intervention makes it available, providing the agent with the extracted state and a fixed set of handling policies π alongside the context,

$$A^{\mathrm{int}} \;=\; M_\theta(C,\, S_{\mathrm{SOIN}},\, \pi)\,. \tag{5}$$

We test in Section 5 whether making the recognized state available in this way changes how the agent handles the inconsistency.

3 DETECTING SOURCE INCONSISTENCY

3.1 EXPERIMENT SETUP

To test whether M_θ can reliably produce a structured SOIN state S_{SOIN} (Equations 3 and 4) that reflects the source-consistency status of the current context, we evaluate detection over controlled conflicts across models and source configurations. We use the open-source `Browser Use` as our scaffold and five commercial models: `GPT-5.6 Sol`, `GPT-5.6 Luna`, `Claude Haiku 4.5`, `Gemini 3.5 Flash Lite`, and `Seed 2.0 Lite`. Model versions, access details, and configurations are in Appendix A.

Datasets. We adapt WebVoyager (He et al., 2024) and WebArena (Zhou et al., 2024a) into tasks spanning ***20*** natural scenes, such as flight booking, online shopping, and appointment scheduling, each realizing one of three interaction types: button clicking, form filling, and summarizing. Tasks are sampled uniformly across interaction types, so source configuration is not tied to interaction type. Inconsistency is instantiated across the three source categories of Section 2.1: the **(U)** *user* prompt, the **(E)** *external* observation (DOM, accessibility tree, and screenshot, kept consistent and modified jointly), and the **(I)** *internal* state (memory, assessment, and current goal). Task construction follows three principles that isolate source inconsistency from unrelated behavioral bias:

- **Source-neutral carriers.** Each category presents its claim in a record format, without ownership or authority cues such as *my*, *official*, or *latest* that would confound category preference with lexical trust signals. Matched carrier–value assignments systematically exchange which source carries each value, controlling for preference toward the values themselves.
- **Knowledge-independent decision variables.** All candidate values are defined solely by records in the current context, so the agent must arbitrate among the provided carriers, with no answer available from world or parametric knowledge.
- **Decision-critical inconsistency.** The inconsistent field determines the immediate next action: distinct values lead to distinct executable actions, and the task cannot be completed by bypassing the field. We target ***5*** such fields per scene.

For each target field we construct seven source-inconsistency configurations: three two-source conflicts ($\{U{\neq}E\}$, $\{U{\neq}I\}$, $\{E{\neq}I\}$), three three-source conflicts with two agreeing sources ($U{=}E{\neq}I$, $U{=}I{\neq}E$, $E{=}I{\neq}U$), and one with all three sources distinct, each realized under multiple carrier–value assignments. This gives 15 SOIN instances per field and $20 \times 5 \times 15 = 1500$ in total (more details in Appendix B), derived from 100 base scene–field combinations, with an equal number of no-SOIN counterparts for true-negative measurement. Results are reported as means with standard deviations over three runs.

Evaluation Protocol. A prediction is counted as correct only when both the predicted SOIN status and the complete set of conflicting source pairs match the ground truth:

$$TP_{\text{state}} \;=\; \mathbb{1}\big[\hat{G} = G = 1 \;\wedge\; \hat{\mathcal{P}} = \mathcal{P}\big], \qquad TN_{\text{state}} \;=\; \mathbb{1}\big[\hat{G} = G = 0 \;\wedge\; \hat{\mathcal{P}} = \varnothing\big]. \tag{6}$$

We report these as class-conditional rates, $\mathbb{E}[TP_{\text{state}} \mid G = 1]$ and $\mathbb{E}[TN_{\text{state}} \mid G = 0]$, over the conflict and no-conflict instances respectively. A state with $\hat{G} = 1$ is incorrect if it omits a ground-truth pair, adds a spurious one, or identifies the wrong combination; in the multi-source settings the complete set of pairs must be recovered. The protocol therefore scores the *topological* component of the extracted state—the conflict flag $\hat{G}$ and the pair set $\hat{\mathcal{P}}$—and not the fidelity of the evidence and relevance fields ($\hat{\mathcal{E}}$, $\hat{r}$), which the downstream actor consumes as grounding. A state can thus recover the correct pairs while localizing the discrepancy imprecisely; Section 5 shows that this distinction matters for grounded resolution.

3.2 SOIN DETECTION RESULTS

Tables 1 and 2 report state-level SOIN detection across all seven source configurations. In the two-source settings, TP_{state} is consistently high: GPT-5.6 Sol reaches 100.00% on $U. \neq I.$, Claude Haiku 4.5 99.83% on $E. \neq I.$, and Seed 2.0 Lite 96.83% on $U. \neq E.$. The pattern holds in the three-source settings despite the stricter requirement to recover the complete conflict structure: GPT-5.6 Sol reaches 100.00% on $U. = I. \neq E.$, and Seed 2.0 Lite stays above 97% across all four. Under explicit elicitation, the models thus recover the full conflict structure, not merely a positive flag.

	$U. \neq E., I. : \varnothing$		$U. \neq I., E. : \varnothing$		$E. \neq I., U. : \varnothing$	
Model	TP_{state}	TN_{state}	TP_{state}	TN_{state}	TP_{state}	TN_{state}
GPT-5.6 Sol	97.83 ± 0.6	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0	97.33 ± 1.8	100.00 ± 0.0
GPT-5.6 Luna	96.67 ± 1.9	99.67 ± 0.3	98.50 ± 0.9	99.33 ± 0.6	95.67 ± 1.2	99.50 ± 0.0
Claude Haiku 4.5	99.83 ± 0.3	95.50 ± 0.9	93.50 ± 1.3	90.17 ± 1.2	99.83 ± 0.3	96.83 ± 0.6
Gemini 3.5 Flash Lite	97.00 ± 1.5	97.67 ± 1.6	98.17 ± 1.0	99.83 ± 0.3	92.50 ± 1.5	99.83 ± 0.3
Seed 2.0 Lite	96.83 ± 2.0	95.83 ± 1.0	93.17 ± 1.5	98.00 ± 1.3	98.17 ± 0.8	92.00 ± 2.0

Table 1: SOIN detection performance (%) under the three two-source configurations. **U. / E. / I.**: User / External / Internal information source.

	$U. = E. \neq I.$		$U. = I. \neq E.$		$E. = I. \neq U.$		$E. \neq I. \neq U.$	
Model	TP_{state}	TN_{state}	TP_{state}	TN_{state}	TP_{state}	TN_{state}	TP_{state}	TN_{state}
GPT-5.6 Sol	97.83 ± 1.4	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0	98.17 ± 1.8	100.00 ± 0.0	96.89 ± 1.2	100.00 ± 0.0
GPT-5.6 Luna	82.50 ± 0.5	99.67 ± 0.3	93.00 ± 1.8	99.83 ± 0.3	91.17 ± 2.3	99.00 ± 1.3	96.00 ± 1.5	99.67 ± 0.0
Claude Haiku 4.5	88.00 ± 0.5	96.83 ± 0.3	82.00 ± 1.0	98.83 ± 1.0	98.33 ± 0.8	97.17 ± 1.8	92.33 ± 0.7	98.89 ± 0.5
Gemini 3.5 Flash Lite	90.00 ± 1.5	100.00 ± 0.0	90.00 ± 0.0	99.50 ± 0.0	86.50 ± 1.8	99.83 ± 0.3	91.00 ± 0.9	99.67 ± 0.3
Seed 2.0 Lite	98.67 ± 1.0	98.00 ± 1.8	98.00 ± 1.7	95.67 ± 1.8	98.00 ± 1.3	98.50 ± 1.3	97.22 ± 0.7	99.00 ± 0.6

Table 2: SOIN detection performance (%) under the four three-source configurations. **U. / E. / I.**: User / External / Internal information source.

The matched no-SOIN controls confirm this is not a bias toward predicting inconsistency: TN_{state} stays high throughout, with GPT-5.6 Sol at 100.00% in every no-SOIN condition and GPT-5.6 Luna and Gemini 3.5 near ceiling. The models thus preserve a negative state when sources agree.

Overall, structured SOIN detection remains consistently strong across source identities and inconsistency configurations, with no configuration inducing a systematic degradation across models. Performance nevertheless remains below perfect detection. Because TP_{state} requires recovering the exact conflict-pair set, most residual errors reflect wrong reporting of conflict pairs rather than a complete failure to detect inconsistency (Appendix C.2). Such errors preserve the binary conflict signal but can still affect grounded resolution when downstream handling depends on correctly localizing the conflicting sources and evidence (details in Section 5).

4 Natural Source Arbitration

Section 3 shows that, when explicitly probed, current models reliably recognize source inconsistency and recover its conflict structure. Recognition alone, however, does not reveal how the same inconsistency shapes behavior during normal execution. We therefore reuse the same decision contexts and configurations and observe $A = M_\theta(C)$ without exposing the SOIN state or any inconsistency-specific instruction, characterizing natural source-arbitration behavior under conflicts the available evidence cannot resolve. Because each source's value is fixed by construction, we label each output deterministically: an action that commits to a source's value is labeled ***Follow User (U.)***, ***Follow Environment (E.)***, or ***Follow Internal (I.)*** accordingly, and any action that does not commit to a conflicting value is labeled ***No Commitment (N.C.)***.

We augment the action space with three source-neutral resolution actions that together span the ways of handling an unresolved conflict without committing to a value, differing in where the next information comes from: `verify` re-examines the agent-accessible sources (E and I); `ask_user` requests clarification from the user (U), the only source the agent cannot re-inspect on its own; and `defer` seeks no further information and withholds the current decision. Since U, E, and I exhaust the conflicting sources, these three span the next-step choices under SOIN. None is tied by instruction to SOIN, uncertainty, or source identity, so the observed behavior reflects how the agent freely chooses among commitment and these actions.

Tables 3 and 4 reveal a clear first pattern: **No Commitment (N.C.) is consistently rare**, below 7.6% across all models and configurations. Agents thus overwhelmingly respond to SOIN by committing to a conflicting source, rarely verifying, clarifying, or deferring. This contrasts sharply with the strong elicited detection of Section 3: inconsistency is recognizable when probed, yet seldom prompts spontaneous handling during normal execution.

	$U. \neq E., I. : \varnothing$			$U. \neq I., E. : \varnothing$			$E. \neq I., U. : \varnothing$		
Model	$U.$	$E.$	N.C.	$U.$	$I.$	N.C.	$E.$	$I.$	N.C.
GPT-5.6 Sol	68.33 ± 2.8	27.83 ± 2.5	3.83 ± 0.6	100.00 ± 0.0	0.00 ± 0.0	0.00 ± 0.0	63.83 ± 2.4	32.33 ± 2.3	3.83 ± 3.0
GPT-5.6 Luna	66.67 ± 1.4	31.00 ± 1.5	2.33 ± 2.1	97.00 ± 1.8	3.00 ± 1.8	0.00 ± 0.0	47.00 ± 0.5	50.83 ± 0.3	2.17 ± 0.6
Claude Haiku 4.5	67.83 ± 1.8	31.00 ± 2.2	1.17 ± 0.8	95.67 ± 1.4	1.67 ± 1.0	2.67 ± 1.5	41.67 ± 1.2	55.17 ± 1.0	3.17 ± 0.3
Gemini 3.5 Flash Lite	65.83 ± 1.2	30.83 ± 0.3	3.33 ± 1.0	94.83 ± 1.6	0.00 ± 0.0	5.17 ± 1.6	49.33 ± 1.0	47.33 ± 1.9	3.33 ± 1.5
Seed 2.0 Lite	69.33 ± 0.8	30.67 ± 0.8	0.00 ± 0.0	95.67 ± 2.3	2.67 ± 1.6	1.67 ± 0.8	45.67 ± 2.1	54.33 ± 2.1	0.00 ± 0.0

Table 3: Natural source-arbitration behavior (%) under the three two-source configurations.

	$U. = E. \neq I.$			$U. = I. \neq E.$			$E. = I. \neq U.$			$E. \neq I. \neq U.$			
Model	$U.E.$	$I.$	N.C.	$U.I.$	$E.$	N.C.	$E.I.$	$U.$	N.C.	$E.$	$I.$	$U.$	N.C.
GPT-5.6 Sol	97.00 ± 2.2	0.00 ± 0.0	3.00 ± 2.2	79.00 ± 2.0	21.00 ± 2.0	0.00 ± 0.0	35.33 ± 0.8	63.00 ± 1.0	1.67 ± 1.0	25.78 ± 0.4	0.00 ± 0.0	72.00 ± 1.2	2.22 ± 0.8
GPT-5.6 Luna	98.33 ± 1.0	0.00 ± 0.0	1.67 ± 1.0	70.83 ± 1.5	29.17 ± 1.5	0.00 ± 0.0	38.33 ± 2.0	59.67 ± 1.6	2.00 ± 1.0	28.56 ± 1.3	0.00 ± 0.0	68.56 ± 0.8	2.89 ± 1.0
Claude Haiku 4.5	93.83 ± 0.8	0.00 ± 0.0	6.17 ± 0.8	67.50 ± 1.5	29.00 ± 2.6	3.50 ± 3.0	30.00 ± 1.0	66.50 ± 0.9	3.50 ± 1.3	27.33 ± 1.7	0.00 ± 0.0	65.11 ± 1.0	7.56 ± 2.7
Gemini 3.5 Flash Lite	94.67 ± 1.0	0.00 ± 0.0	5.33 ± 1.0	64.17 ± 0.8	33.67 ± 0.8	2.17 ± 0.8	32.67 ± 1.3	63.67 ± 0.6	3.67 ± 1.4	31.56 ± 1.5	0.00 ± 0.0	63.78 ± 1.3	4.67 ± 2.7
Seed 2.0 Lite	98.17 ± 1.0	1.83 ± 1.0	0.00 ± 0.0	64.83 ± 2.0	35.17 ± 2.0	0.00 ± 0.0	35.67 ± 2.0	63.00 ± 1.3	1.33 ± 1.5	32.67 ± 1.2	0.00 ± 0.0	65.67 ± 0.7	1.67 ± 1.3

Table 4: Natural source-arbitration behavior (%) under the four three-source configurations.

Table 3 shows highly consistent behavior when $U.$ is in pairwise conflict. Under $U. \neq E.$, all five models follow $U.$ in 65.83–69.33% of cases and $E.$ in 27.83–31.00%; under $U. \neq I.$, $U.$ determines 94.83–100.00% while $I.$ accounts for at most 3.00%. Thus $U.$ dominates both, though $E.$ resists $U.$ far more than $I.$ does. The $E. \neq I.$ setting is model-dependent: GPT-5.6 Sol favors $E.$ (63.83%), Claude Haiku 4.5 and Seed 2.0 Lite favor $I.$ (54.33–55.17%), and GPT-5.6 Luna and Gemini 3.5 Flash Lite are nearly balanced, with no consistent cross-model preference.

The three-source settings in Table 4 show that source *composition*, not agreement count alone, drives the action. When $U.$ agrees with $E.$, the $U.E.$ value determines 93.83–98.33% of actions; when $U.$ agrees with $I.$, 64.17–79.00%. More strikingly, under $E. = I. \neq U.$ the lone $U.$ source still determines 59.67–66.50% despite opposing two agreeing sources; and when all three differ, $U.$ accounts for 63.78–72.00% and $E.$ for 25.78–32.67%, while $I.$ is never selected. Overall, agents seldom handle SOIN during unprompted execution, instead showing systematic, source-dependent selection: user information most often determines the action in three-source settings, while the influence of external evidence versus internal state is more model-dependent. We revisit the origins of this source-dependence in Section 7.2.

5 Decision-Time Conflict Mediation

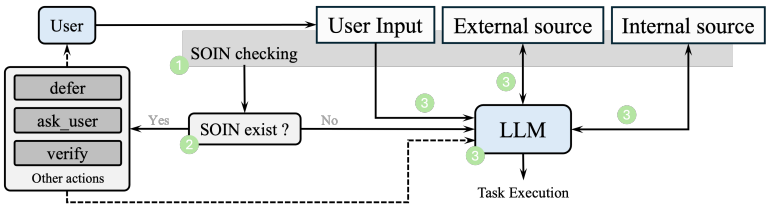

Figure 2: **SOIN-aware action mediation before each agent action**. (1) The current sources are checked to construct the SOIN state. (2) If a conflict is detected, the agent may invoke a handling action (`verify`, `ask_user`, or `defer`); otherwise, mediation passes through without intervention. (3) Any information returned by the intervention is passed to the LLM for continued task execution.

Section 4 shows that, despite strong elicited SOIN recognition, agents rarely translate it into spontaneous handling during normal execution. We therefore test whether mediating the next action on a detected conflict changes this behavior. Following Equation 5, we instantiate the SOIN-aware mediation framework of Figure 2: before each action, the user, external, and internal sources are checked to construct S_{SOIN}, exposed at a pre-action mediation stage as an explicit signal. The agent then either continues execution or invokes an intervention action based on the detected state. This layer is model-agnostic and training-free, attaching to a scaffold as a deployable pre-action check.

In our controlled SOIN setting, the conflict is unresolved at the decision point: the available evidence provides no basis for privileging one source. To evaluate whether mediation leads to a grounded resolution rather than merely a handling action, we augment each task in this section with a recoverable ground-truth value that can be obtained only through interaction: `ask_user` elicits it when the agent poses the appropriate question, while `verify` elicits it when the agent locates and queries the

supporting page element. An instance succeeds only when the agent first acquires this value through a correct `verify` or `ask_user` action and then uses it to complete the task. Direct commitment is not credited, even if it happens to match the recoverable value, because the conflict has not been resolved from the evidence available at the decision point; `defer` without acquiring the value likewise does not complete the task. We report

$$\mathrm{SR_c} = P(\text{success} \mid G = 1)\,, \qquad \mathrm{SR_n} = P(\text{success} \mid G = 0)\,, \tag{7}$$

the success rate under conflict ($\mathrm{SR_c}$, $G = 1$) and with no conflict ($\mathrm{SR_n}$, $G = 0$). A no-conflict instance succeeds when the agent completes the task directly, without an unnecessary handling action. Higher $\mathrm{SR_c}$ means the intervention more often drives the conflict to a grounded resolution, crediting an outcome only when the correct value is both actively obtained and correctly applied; higher $\mathrm{SR_n}$ means it rarely interrupts already-consistent cases. The three handling actions have concrete implementations (Appendix D.1). They cover the ways an agent can respond to a detected conflict: `verify` and `ask_user` actively seek the missing information from the page or the user, while `defer` withholds the decision without seeking it. Under our success criterion, only the active actions can complete the task. We compare four conditions under the same action space:

- **Base**: native scaffold with the same three handling actions available in the action space, but no SOIN-specific guidance or state.
- **Binary SOIN State**: exposes only a binary conflict flag ($\hat{G}$) with the handling policy, omitting the conflicting pairs, evidence, and relevance of the full state.
- **Full SOIN State**: exposes the complete structured state S_{SOIN} from the agent's own diagnostic call, with the handling policy.
- **Oracle SOIN State**: the same mediation with the extracted state replaced by the ground-truth SOIN state, revealing no preferred source or value.

These conditions separate the overall effect of SOIN-aware mediation from the information content of the checking result. The mechanism comprises two pre-action steps: a conflict is detected, and the next action is mediated on that result. Its realization, whether as a structured state, plain text, or a single combined prompt, is an implementation choice, and any procedure performing both steps constitutes an instance of the mechanism. The steps are not separable: the handling policy is defined over a detected conflict and has no trigger in its absence, so removing detection removes the policy with it and reduces the setting to native execution; supplying either component alone leaves the mechanism ineffective (Appendix E.2). **Base** is this native execution, retaining the same three handling actions with nothing that detects or reports a conflict. Among the mediation conditions, **Binary** versus **Full** measures the added value of exposing the full conflict structure over a bare flag, and **Full** versus **Oracle** isolates the effect of extraction errors (full prompts in Appendix E.1).

	Base		**Binary SOIN State**		**Full SOIN State**		**Oracle SOIN State**	
Model	$\mathrm{SR_c}$	$\mathrm{SR_n}$	$\mathrm{SR_c}$	$\mathrm{SR_n}$	$\mathrm{SR_c}$	$\mathrm{SR_n}$	$\mathrm{SR_c}$	$\mathrm{SR_n}$
GPT-5.6 Sol	2.47 ± 0.1	100.00 ± 0.0	4.69 ± 1.0	97.87 ± 0.5	100.00 ± 0.0	98.91 ± 0.3	100.00 ± 0.0	100.00 ± 0.0
GPT-5.6 Luna	9.31 ± 0.1	100.00 ± 0.0	14.87 ± 0.4	97.73 ± 0.2	100.00 ± 0.0	98.82 ± 0.1	100.00 ± 0.0	100.00 ± 0.0
Claude Haiku 4.5	0.00 ± 0.0	100.00 ± 0.0	0.00 ± 0.0	98.84 ± 0.5	62.51 ± 1.0	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0
Gemini 3.5 Flash Lite	14.73 ± 1.2	99.69 ± 0.1	96.82 ± 0.2	98.36 ± 0.3	99.13 ± 0.4	99.62 ± 0.3	99.76 ± 0.0	100.00 ± 0.0
Seed 2.0 Lite	7.93 ± 1.2	100.00 ± 0.0	63.44 ± 1.1	99.00 ± 0.5	99.49 ± 0.2	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0

Table 5: Grounded conflict-resolution performance (%) across the four decision-time mediation conditions. $\mathrm{SR_c}$: success rate under conflict ($G = 1$); $\mathrm{SR_n}$: success rate with no conflict ($G = 0$).

Table 5 reports grounded conflict-resolution success: how often the agent acquires the missing evidence and completes the task with the recovered value. Under natural execution, success is rare, with **Base** at $\mathrm{SR_c} = 0$–14.73%. The recoverable value is reachable in every instance through the same handling actions available to all conditions, yet without a pre-action check the agent seldom initiates the verification or user query that would obtain it, committing instead to a conflicting value. The bottleneck is the absence of detection: agents do not register the conflict on their own.

Detection alone, however, is not sufficient. The **Binary SOIN State** reports only that a conflict exists, and its gains are uneven, with $\mathrm{SR_c}$ from 0% (Claude Haiku 4.5) to 96.82% (Gemini 3.5 Flash Lite). A bare flag leaves the implicated sources, conflicting claims, and supporting evidence for the actor to reconstruct, which most models do not do reliably: they resort to `defer` or target the wrong

information (Appendix E.3). Current models cannot be relied upon to localize a conflict once told one exists; the checking result must state what it is.

Supplying that structure closes the gap. The **Full SOIN State** reaches $\mathrm{SR_c} = 99.13$–100.00% for four of the five models, matching the **Oracle SOIN State** built from the ground-truth structure, with $\mathrm{SR_n}$ at 98.82–100.00%. Claude Haiku 4.5 improves from 0.00% under both Base and the binary state to 62.51%; its residual failures arise when the extracted state under-localizes the core discrepancy, so the subsequent `ask_user` or `verify` targets the wrong information (Appendix E.3), and it reaches 100.00% with the oracle state. That extracted and ground-truth structures yield near-identical outcomes for most models indicates the checking result is accurate enough to substitute for the true one. These rates hold where the correct value is recoverable through interaction; settings in which verification returns no usable information pose a separate question.

6 SOIN AT FINER EXTERNAL-SOURCE GRANULARITY

A web agent's external evidence is itself split across channels: a *DOM-derived textual observation* and a *visual screenshot* of the same webpage, which need not agree. We therefore examine SOIN *within* the external source, denoting the two channels E_{DOM} and E_{Vis} and restricting decision-relevant evidence to them, with all other settings unchanged (Appendix B). Repeating the detection–characterization–intervention pipeline tests whether the recognition–behavior gap and the benefit of explicit SOIN checking persist at this finer granularity.

	$E_{\mathrm{DOM}} \neq E_{\mathrm{Vis}}, I., U. : \varnothing$				
Model	TP_{state}	TN_{state}	E_{DOM}	E_{Vis}	N.C.
GPT-5.6 Sol	100.00 ± 0.0	100.00 ± 0.0	29.83 ± 1.5	70.17 ± 1.5	0.00 ± 0.0
GPT-5.6 Luna	97.83 ± 1.4	100.00 ± 0.0	22.00 ± 1.7	78.00 ± 1.7	0.00 ± 0.0
Claude Haiku 4.5	75.50 ± 1.0	100.00 ± 0.0	26.00 ± 1.0	69.33 ± 0.6	4.67 ± 1.5
Gemini 3.5 Flash Lite	98.67 ± 1.0	100.00 ± 0.0	30.17 ± 1.2	65.50 ± 0.9	4.33 ± 1.0
Seed 2.0 Lite	100.00 ± 0.0	99.00 ± 0.0	28.00 ± 1.3	71.33 ± 1.2	0.67 ± 0.3

Table 6: SOIN detection and natural source-arbitration behavior (%) ($E_{\mathrm{DOM}} \neq E_{\mathrm{Vis}}$)

Table 6 shows SOIN remains readily detectable when the conflict lies within a single external source: TP_{state} reaches 97.83–100.00% for four models (Claude Haiku 4.5 lower at 75.50%) and TN_{state} stays 99.00–100.00%, so the diagnostic identifies genuine disagreement between two channels of the same webpage, not coarse source labels. The recognition–behavior gap nevertheless persists: N.C. stays below 4.7%, and the arbitration pattern is unusually stable across models, with E_{Vis} selected in 65.50–78.00% of cases versus 22.00–30.17% for E_{DOM}. Thus, unlike the model-dependent external–internal conflicts of Section 4, disagreement within the external observation induces a stable cross-model preference for visual evidence.

	Base		**Binary SOIN State**		**Full SOIN State**		**Oracle SOIN State**	
Model	$\mathrm{SR_c}$	$\mathrm{SR_n}$	$\mathrm{SR_c}$	$\mathrm{SR_n}$	$\mathrm{SR_c}$	$\mathrm{SR_n}$	$\mathrm{SR_c}$	$\mathrm{SR_n}$
GPT-5.6 Sol	0.00 ± 0.0	100.00 ± 0.0	69.89 ± 0.3	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0
GPT-5.6 Luna	0.00 ± 0.0	100.00 ± 0.0	73.29 ± 0.8	100.00 ± 0.0	97.16 ± 0.3	100.00 ± 0.0	97.58 ± 0.1	100.00 ± 0.0
Claude Haiku 4.5	0.00 ± 0.0	100.00 ± 0.0	33.13 ± 0.7	100.00 ± 0.0	76.98 ± 0.9	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0
Gemini 3.5 Flash Lite	0.00 ± 0.0	100.00 ± 0.0	71.78 ± 0.7	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0
Seed 2.0 Lite	0.00 ± 0.0	100.00 ± 0.0	70.51 ± 0.4	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0	100.00 ± 0.0

Table 7: Grounded conflict-resolution performance (%) across the four decision-time mediation conditions ($E_{\mathrm{DOM}} \neq E_{\mathrm{Vis}}$). $\mathrm{SR_c}$ / $\mathrm{SR_n}$: success rate under conflict ($G = 1$) / no conflict ($G = 0$).

Table 7 shows that the same SOIN-checking and mediation pattern holds within the external observation. Without mediation, agents commit to one of the conflicting channels instead of acquiring evidence, yielding $\mathrm{SR}_c = 0.00\%$ for all models, with no-conflict cases uninterrupted ($\mathrm{SR}_n = 100.00\%$). The **Binary SOIN State** improves grounded resolution but remains incomplete ($\mathrm{SR}_c = 33.13$–73.29%), whereas the **Full SOIN State** raises SR_c to 97.16–100.00% for four models and to 76.98% for Claude Haiku 4.5, preserving $\mathrm{SR}_n = 100.00\%$ throughout. Full-state performance matches the oracle for three models and is within 0.42 points for GPT-5.6 Luna; Claude Haiku 4.5 reaches 100.00% under the oracle state, locating its remaining gap in the extracted state and subsequent targeting rather than the mediation itself. Together with Table 6, the recognition–behavior gap and the benefit of explicit SOIN checking persist below the user–external–internal split, down to disagreement between DOM-derived and visual channels.

7 Discussion

7.1 Related Work

Agent Source Inconsistency and Arbitration. Prior work often assumes a predefined resolution target: newer state over superseded memory (Fan et al., 2026), task-correct behavior over conflicting retrieved memory (Chen et al., 2026), verifiable interface state over alternatives (Zhang et al., 2026a), and prescribed source hierarchies (Li et al., 2026; Zhou et al., 2024b; Zhang et al., 2026b). More recent work studies unresolved information: HiL-Bench (Trinh et al., 2026) examines selective help-seeking under incomplete or contradictory tasks, while TANGLE (Yang et al., 2026) focuses on irreducible conflicts within personal memory. SOIN instead studies controlled disagreement across live U, E, and I channels through explicit checking and grounded task resolution.

Adversarially Induced Conflict. Many web-agent attacks can be viewed as deliberately induced source conflict, where attacker-controlled signals compete with the user's task or the agent's legitimate state: embedding malicious instructions in web content for privacy leakage or unauthorized actions (Li et al., 2025; Liao et al., 2025), injecting adversarial HTML or visual signals (Wang et al., 2025), generating injections or social-engineering cues that redirect agents (Xu et al., 2025; Wu et al., 2026), and poisoning agent memory through malicious observations (Zou et al., 2026). These assume a malicious source to be rejected; SOIN instead studies naturally arising disagreement among plausible sources, with no side privileged in advance.

7.2 Limitations and Future Work

From implicit reasoning to explicit SOIN checking. Current agent scaffolds largely leave source consistency implicit in model reasoning. Sections 5 and 6 show that explicitly checking for SOIN before action and making the checking result available to downstream mediation substantially improves grounded task success while preserving performance in no-conflict cases. We instantiate the checking result as a structured state to support controlled comparisons of its information content, including binary, full, and oracle variants. More broadly, the key design principle is to make source inconsistency explicit and actionable at decision time; the checking result can be communicated through structured fields, natural language, or other suitable representations. Exploring these representations and integration mechanisms is a natural direction for future work.

Understanding source arbitration. Our controlled conflicts reveal systematic source-dependent behavior: a strong tendency to favor user-provided information, while external–internal conflicts are substantially more model-dependent (Section 4). The preference for the user may reflect training-induced priors that treat user input as intent to be followed. When internal state directly conflicts with user information, its weaker influence may instead reflect priors about the freshness or reliability of maintained memory. Because the natural-arbitration setting in Section 4 admits no correct value and strips authority, ownership, and recency cues, these preferences characterize arbitration under underdetermined evidence; tracing them to specific training choices is a natural direction for mechanistic study.

8 Conclusion

We study *Source Inconsistency (SOIN)* as a decision-time reliability problem in which agent-visible sources support incompatible claims that the available evidence cannot adjudicate. Across controlled web-agent settings, strong elicited recognition coexists with rare spontaneous resolution and systematic source-dependent commitment. Explicitly checking for SOIN before action and making the result available to downstream mediation substantially improves grounded task success: full checking information reaches near-oracle performance for most models while preserving success in no-conflict cases, with oracle results further localizing residual errors to upstream checking and execution. The same pattern persists when the conflict occurs between DOM-derived and visual observations of the same webpage. Together, these results support explicit SOIN checking as a practical component of agent decision making and reliability design.

AI USE STATEMENT

Generative AI tools were used only for editorial assistance on this manuscript: improving grammar and sentence structure, tightening clarity, and standardizing notation and terminology. They were not used to conceive the research questions, design the SOIN framework or experiments, generate or process experimental data, or interpret results. Separately, and as an object of study rather than an authoring tool, commercial LLMs serve as the agent backbones in our experiments (Sections 3–5); their role and configuration are described in the main text and Appendix A. All text was reviewed and edited by the authors, who take full responsibility for the content, claims, and technical accuracy of this work.

ETHICS STATEMENT

This work studies a reliability failure mode of web agents: how they behave when user, external, and internal sources disagree with no principled basis for privileging one. Our aim is to characterize and mitigate this behavior, not to enable misuse. All experiments run on controlled, synthetically constructed conflicts adapted from public benchmarks (WebVoyager and WebArena); we introduce no real personal data, and the conflicting record values are fabricated for the tasks. Our findings are defensive in orientation: the proposed decision-time SOIN state is intended to make agents more cautious under unresolved conflict by prompting verification, clarification, or deferral rather than unjustified commitment. We note that the same characterization of source-arbitration tendencies could in principle inform adversarial manipulation of agents; we therefore frame these results as motivation for more robust, consistency-aware scaffolds, and we discuss the boundaries of our claims in Section 7.2. We have read the ICLR Code of Ethics and believe this work conforms to it.

REPRODUCIBILITY STATEMENT

We provide the details needed to reproduce our experiments. The SOIN formulation and the two-call actor/diagnostic protocol are specified in Section 2; the task construction, source-inconsistency configurations, and evaluation protocol in Section 3; and the intervention conditions in Section 5. Model configurations are given in Appendix A, dataset construction in Appendix B, and the diagnostic instruction Q_{detect}, handling policies in Appendix E.1.

REFERENCES

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. A survey on evaluation of large language models. *ACM Trans. Intell. Syst. Technol.*, 15(3), March 2024. ISSN 2157-6904. doi: 10.1145/3641289. URL https://doi.org/10.1145/3641289.

Yixiong Chen, Xinyi Bai, and Alan Yuille. The compliance trap: Diagnosing how ai agents consume conflicting memory, 2026. URL https://arxiv.org/abs/2607.10608.

Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 28091–28114. Curran Associates, Inc., 2023. doi: 10.52202/075280-1220. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/5950bf290a1570ea401bf98882128160-Paper-Datasets_and_Benchmarks.pdf.

Xinyi Fan, Miri Liu, Ruozhen Yang, Siru Ouyang, and Jiawei Han. Can agent memory systems track evolving state?, 2026. URL https://arxiv.org/abs/2608.19652.

Mohamed Amine Ferrag, Norbert Tihanyi, and Mérouane Debbah. From llm reasoning to autonomous ai agents: A comprehensive review. *IEEE Access*, 14:84237–84285, 2026. doi: 10.1109/ACCESS.2026.3698694.

Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. Large language model based multi-agents: A survey of progress and challenges, 2024. URL https://arxiv.org/abs/2402.01680.

Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. WebVoyager: Building an end-to-end web agent with large multimodal models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 6864–6890, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.371. URL https://aclanthology.org/2024.acl-long.371/.

Ang Li, Yin Zhou, Vethavikashini Chithrra Raghuram, Tom Goldstein, and Micah Goldblum. Commercial llm agents are already vulnerable to simple yet dangerous attacks, 2025. URL https://arxiv.org/abs/2502.08586.

Shuowei Li, Haoxin Li, Wenda Chu, and Yi Fang. How large language models balance internal knowledge with user and document assertions. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 25323–25346, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.1267. URL https://aclanthology.org/2026.findings-acl.1267/.

Zeyi Liao, Lingbo Mo, Chejian Xu, Mintong Kang, Jiawei Zhang, Chaowei Xiao, Yuan Tian, Bo Li, and Huan Sun. Eia: Environmental injection attack on generalist web agents for privacy leakage. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 66972–67003, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/a73474c359ed523e6cd3174ed29a4d56-Paper-Conference.pdf.

Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. A comprehensive overview of large language models. *ACM Trans. Intell. Syst. Technol.*, 16(5), August 2025. ISSN 2157-6904. doi: 10.1145/3744746. URL https://doi.org/10.1145/3744746.

Tu Trinh, Mohamed Elfeki, Guangze Luo, Kelvin Luu, Nathan Hunt, Ernesto Hernandez, Nandan Marwaha, Yannis Yiming He, Charles Wang, Fernando Carabedo, Alessa Castillo, and Bing Liu. Hil-bench (human-in-loop benchmark): Do agents know when to ask for help?, 2026. URL https://arxiv.org/abs/2604.09408.

Xilong Wang, John Bloch, Zedian Shao, Yuepeng Hu, Shuyan Zhou, and Neil Zhenqiang Gong. WebInject: Prompt injection attack to web agents. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 2010–2030, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.104. URL https://aclanthology.org/2025.emnlp-main.104/.

Xinyi Wu, Geng Hong, Yueyue Chen, MingXuan Liu, Feier Jin, Xudong Pan, Jiarun Dai, and Baojun Liu. When bots take the bait: Exposing and mitigating the emerging social engineering attack in web automation agent, 2026. URL https://arxiv.org/abs/2601.07263.

Chejian Xu, Mintong Kang, Jiawei Zhang, Zeyi Liao, Lingbo Mo, Mengqi Yuan, Huan Sun, and Bo Li. Advagent: Controllable blackbox red-teaming on web agents, 2025. URL https://arxiv.org/abs/2410.17401.

Lu Yang, Shusheng Xu, Zhuoran Li, Tongkai Yang, and Longbo Huang. When personal memory has no single answer: Evaluating llm agents under irreducible conflict, 2026. URL https://arxiv.org/abs/2608.13921.

Guijia Zhang, Yuxun Chen, Yuheng Qi, and Harry Yang. Do gui agents believe their eyes? diagnosing state-belief reliance on pixels versus structure, 2026a. URL https://arxiv.org/abs/2607.04334.

Jingyu Zhang, Tianjian Li, William Jurayj, Hongyuan Zhan, Benjamin Van Durme, and Daniel Khashabi. Many-tier instruction hierarchy in llm agents, 2026b. URL https://arxiv.org/abs/2604.09443.

Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *International Conference on Learning Representations*, volume 2024, pp. 15585–15606, 2024a. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/4410c0711e9154a7a2d26f9b3816d1ef-Paper-Conference.pdf.

Sizhe Zhou, Sha Li, Yu Meng, Yizhu Jiao, Heng Ji, and Jiawei Han. Establishing knowledge preference in language models, 2024b. URL https://arxiv.org/abs/2407.13048.

Wei Zou, Mingwen Dong, Miguel Romero Calvo, Shuaichen Chang, Jiang Guo, Dongkyu Lee, Xing Niu, Xiaofei Ma, Yanjun Qi, and Jiarong Jiang. Poison once, exploit forever: Environment-injected memory poisoning attacks on web agents, 2026. URL https://arxiv.org/abs/2604.02623.

APPENDIX CONTENTS

A Model Configurations **14**

B Datasets **14**

C Appendix: Detecting Source Inconsistency **16**

C.1 Qdetect prompt template 16

C.2 Failure Case Studies in SOIN Detection 18

D Appendix: Natural Source Arbitration **25**

D.1 Conflict-Handling Action Implementations 25

E Appendix: Decision-Time State Mediation **25**

E.1 Prompts 25

E.2 Ablations on SOIN components 26

E.3 Case Studies in State Mediation 27

A Model Configurations

We use the open-source `Browser Use`[1] (version `0.13.7`) as our web-agent scaffold and evaluate five commercial models: `GPT-5.6 Sol`[2], `GPT-5.6 Luna`[3], `Claude Haiku 4.5`[4], `Gemini 3.5 Flash Lite`[5], and `Seed 2.0 Lite`[6]. Reasoning effort is set to `medium` for all five. Table 8 lists their API model identifiers. Calls that fail or return truncated outputs are rerun until a valid response is obtained.

Model	API model ID
GPT-5.6 Sol	`gpt-5.6-sol`
GPT-5.6 Luna	`openai/gpt-5.6-luna`
Claude Haiku 4.5	`anthropic/claude-haiku-4.5`
Gemini 3.5 Flash Lite	`google/gemini-3.5-flash-lite`
Seed 2.0 Lite	`bytedance-seed/seed-2.0-lite`

Table 8: API model identifiers.

B Datasets

Scenes and interaction types. The task content spans 20 base scenes, such as online shopping, flight booking, and appointment scheduling. Within each scene, we adapt three interaction types that specify the immediate response required from the agent:

- **Button clicking.** The decision is presented through clickable candidate controls; the agent selects the control for its chosen value.
- **Form filling.** The decision requires entering or selecting a value in a form field; the value determines the submitted content.
- **Summarizing.** The decision requires reporting the requested information as text; the value determines the answer content.

For each scene, we identify five decision-critical target fields whose value determines the pending action, so that distinct values lead to distinct executable outcomes. Tasks are sampled uniformly across the three interaction types, independently of source configuration.

Source construction and the single-step setting. Each instance is a snapshot of an agent midway through a task, frozen at a single decision point. At this point the agent holds a user instruction, a current external observation of the webpage, and an internal state carried over from prior steps. We construct the three sources of Section 2.1 accordingly:

- **User** (U): the task prompt, optionally asserting a target-field value.
- **External** (E): the current page observation (DOM, accessibility tree, and screenshot), rendered consistently and modified jointly, optionally carrying a target-field value.
- **Internal** (I): the agent's memory, assessment, and current goal, representing information retained from earlier steps of the same task.

We simulate this intermediate decision point directly. Since I is the state an agent accumulates over prior steps, a single snapshot already carries all three sources in the form they take during execution, making the decision point the natural unit at which to study source inconsistency.

For each scene–field combination, we assign candidate values to U, E, and I. Table 9 enumerates the seven source configurations and their carrier–value assignments. The symbol $\varnothing$ means a source supplies no value for the target field, while the task instruction and the page's interaction controls

[1] https://browser-use.com/
[2] https://developers.openai.com/api/docs/models/gpt-5.6-sol
[3] https://developers.openai.com/api/docs/models/gpt-5.6-luna
[4] https://www.anthropic.com/claude/haiku
[5] https://ai.google.dev/gemini-api/docs/models/gemini-3.5-flash-lite
[6] https://seed.bytedance.com/en/seed2

remain present. For two-source conflicts and three-source conflicts with two agreeing sources, we exchange a and b; for the all-distinct configuration, three cyclic assignments give each source each of a, b, and c once.

Configuration	**Assignments** (U, E, I)	**Count**
$U \neq E,\ I = \varnothing$	$(a, b, \varnothing)$; $(b, a, \varnothing)$	2
$U \neq I,\ E = \varnothing$	$(a, \varnothing, b)$; $(b, \varnothing, a)$	2
$E \neq I,\ U = \varnothing$	$(\varnothing, a, b)$; $(\varnothing, b, a)$	2
$U = E \neq I$	(a, a, b); (b, b, a)	2
$U = I \neq E$	(a, b, a); (b, a, b)	2
$E = I \neq U$	(b, a, a); (a, b, b)	2
All three distinct	(a, b, c); (b, c, a); (c, a, b)	3
Total per target field		**15**

Table 9: The 15 SOIN assignments instantiated for every scene–field combination.

With 20 scenes, 5 target fields per scene, and 15 assignments per field, the pool contains $20 \times 5 \times (3 \times 2 + 3 \times 2 + 3) = 1500$ SOIN instances. A source pair counts as a conflict when the two sources are both active and carry different values. A two-source conflict thus has one such pair, a three-source configuration with two agreeing sources has two, and an all-distinct configuration has three.

Matched no-SOIN counterparts. For every SOIN instance, we construct a matched no-SOIN counterpart that preserves the scene, target field, candidate values, record metadata, and set of active sources. We replace the target-field values in all active sources with a single shared value while keeping inactive sources unchanged. When two active sources already agree, we use their shared value; otherwise, we take the value from the first active source in the order (U, E, I). For example, $(a, b, \varnothing)$ is paired with $(a, a, \varnothing)$. This procedure creates matched agreement controls without assigning an authoritative value to the corresponding SOIN instance. Each no-SOIN counterpart has ground-truth status $G = 0$, an empty conflict-pair set, and decision relevance `none`. The resulting dataset therefore contains 1500 SOIN instances and 1500 matched no-SOIN counterparts.

Augmentation for the intervention evaluation. Section 5 scores end-to-end task success, which requires a value the agent can recover by acting. For each SOIN instance we therefore designate one of its conflicting candidate values, drawn uniformly at random, as the recoverable ground truth for that instance. This value is exposed through two channels that respond only to an explicit request: the user answers with it when the agent poses a question naming the conflicting field, and an additional page element supplies it when the agent verifies that field. Neither channel is visible in the decision-point context C, so the conflict remains undetermined until the agent acts. An instance counts as successful when the agent recovers this value through `verify` or `ask_user` and uses it to complete the task; the augmentation is applied only to the instances used in Section 5. Representative recovered calls are shown in Appendix E.3.

Conflicts within external observation channels. For Section 6, we reuse the task records and interaction types above and place the inconsistency between two observations of the same webpage, E_{DOM} and E_{Vis}. The record is rendered as an image whose visible target-field value supplies E_{Vis}, while the DOM-derived textual observation carries a separately assigned value for the same field as E_{DOM}. Record identity, metadata, and other field values are held fixed, and both observations are available together at the decision point. The user instruction still specifies the task and target field, and memory, assessment, and goal retain neutral task-status descriptions, so U and I carry no target-field claim ($U = I = \varnothing$) while the task and agent context remain intact. A SOIN instance assigns distinct candidate values to E_{DOM} and E_{Vis}, alternating their channel positions across instances, giving $G = 1$ with $\{E_{\mathrm{DOM}}, E_{\mathrm{Vis}}\}$ as the only conflicting pair. Each instance has a matched no-SOIN counterpart constructed by the same procedure, with $G = 0$ and an empty conflict-pair set.

C Appendix: Detecting Source Inconsistency

C.1 Qdetect prompt template

Qdetect

```
You are given the frozen model-visible context available to a web
agent immediately before action generation.

Your task is to produce a structured Source-Inconsistency state.
This is detection, not action selection.

A source inconsistency exists only when two source-grounded claims or
action requirements about the controlled decision variable cannot
both be true or jointly satisfied at the current decision step.

Do not classify missing information, uncertainty, paraphrases,
complementary information, jointly satisfiable differences, or
unsupported speculation as source inconsistencies.

Use only these source labels, in this canonical order:

1.  `user_instruction`

2.  `environment_dom`

3.  `environment_axtree`

4.  `environment_screenshot`

5.  `agent_memory`

6.  `agent_assessment`

7.  `agent_goal`

Attribute evidence by the carrier that actually contains the claim:

- claims inside `<user_request>` are `user_instruction`;

- asserted textual page content inside `<browser_state>` is
`environment_dom`;

- use `environment_axtree` only for an explicitly distinct AXTree
carrier;

- visual content in the image following `Current screenshot:` or
inside `<browser_vision>` is `environment_screenshot`;

- content under `Memory:` is `agent_memory`;

- content under `Evaluation of Previous Step:` is `agent_assessment`;

- content under `Next Goal:` is `agent_goal`.

A source participates only if its own carrier explicitly assigns a
value, constraint, or action requirement to the controlled variable.
A generic task, readiness statement, field-name mention, or
instruction to complete/report the field contains no value.  A
selectable candidate or action affordance describes what could be
chosen; it does not by itself assert the current value.  Never copy
evidence from one carrier and relabel it as another source.

Action-affordance exclusion is a hard rule:

- Buttons, links, menu items, radio options, and other clickable
controls represent available actions or candidate values.  Their
labels are not assertions of the current value and must not create
`environment_dom` or `environment_axtree` evidence.
```

```
- This remains true when a control label repeats a value asserted by
another carrier.  A button such as "Choose X" does not independently
assert that the current value is X.
- Count `environment_dom` or `environment_axtree` only when non-action
page content explicitly states a current, saved, selected, assigned,
or recorded value for the controlled variable.
The actual field value must appear in `observed_content`; an option
name such as "Plan option 2" is not a field value.
Before returning JSON, silently apply this exact procedure:
1.  Identify the controlled decision variable.
2.  Extract one value or requirement for it from each participating
source carrier; mark every other source absent.
3.  Group sources that assert the same value.
4.  Compare every pair.  Emit an edge if and only if its two
extracted values differ.
5.  Audit that all unequal pairs and no equal or absent pairs are
present, with evidence from the matching carriers.
The pair count follows directly from the value groups:
- two different source values -> exactly one edge;
- three values A, A, B -> exactly two edges, both incident to the B
source;
- three values A, B, C -> exactly three edges;
- all participating values equal -> no inconsistency.
Examples:
- U=A, E=A, I=B -> (`user_instruction`, `agent_memory`) and (E, I),
but never (U, E).
- U=A, E=B, I=A -> (U, E) and (E, I), but never (U, I).
- U absent, E=A, I=B, with only a generic user task and generic goal
-> (E, I) only.
- If all sources assert A, return the negative state even when other
values appear as selectable alternatives.
In the examples, replace U, E, and I with the exact carrier-specific
labels available in the context.  For screenshot evidence, E is
`environment_screenshot`; for DOM-record evidence, E is
`environment_dom`.
`source_pairs` must contain the complete set of pairwise
inconsistencies.  Do not stop after finding one edge and do not add
an edge by transitivity.
For every detected pair:
1.  Assign a unique `pair_id`.
2.  List exactly two source labels in canonical order.
3.  Link at least one evidence item from each source.
4.  Make every `claim` a concise atomic proposition.
5.  Faithfully transcribe the carrier content into
`observed_content`.
6.  Provide a locator for that carrier.
```

Reuse one evidence item for a source across every pair involving that source.

`decision_relevance` is `none` for no inconsistency, `low` for an irrelevant inconsistency, and `high` when resolution could change the next action, state update, or task outcome.

Do not determine which source is correct or more trustworthy. Do not recommend an action, resolution strategy, memory update, or goal revision.

If no source inconsistency is supported, return:

```
{
"has_source_inconsistency": false,
"source_pairs": [],
"evidence": [],
"decision_relevance": "none"
}
```

Otherwise return the same schema with `has_source_inconsistency` set to `true`.

Return valid JSON only, without Markdown, commentary, or additional fields.

C.2 Failure Case Studies in SOIN Detection

C.2.1 Shopping: power type

0. Model
Gemini 3.5 Flash Lite (medium)

1. Input text

User task
Select an item that matches Purchase Request PR-4825.

Complete the pending item-selection decision for Power type.

Reference Record
Record type: Purchase Request Record
Record ID: PR-4825
Version: V2
Timestamp: 2026-08-05 10:00
Product category: Desk lamp
Preferred finish: Mist green
Power type: Rechargeable
Brightness profile: Adjustable
Mount type: Freestanding
Quantity: 1
Budget band: Band B

Saved browser-state text
URL: http://127.0.0.1:60725/case/case-00315
Title: 127.0.0.1:60725/case/case-00315

Purchase Request Record
Available items
Decision-critical field:
Power type

```
Item option 1
Product category
Desk lamp
Preferred finish
Mist green
Power type
Rechargeable
Brightness profile
Adjustable
Mount type
Freestanding
Quantity
Budget band
Band B
[1745]<button name=choice value=0 />
  Select item
Item option 2
Product category
Desk lamp
Preferred finish
Mist green
Power type
Mains powered
Brightness profile
Adjustable
Mount type
Freestanding
Quantity
Budget band
Band B
[1795]<button name=choice value=1 />
  Select item

Agent internal context
Memory:
Reference Record
Record type: Purchase Request Record
Record ID: PR-4825
Version: V2
Timestamp: 2026-08-05 10:00
Product category: Desk lamp
Preferred finish: Mist green
Power type: Mains powered
Brightness profile: Adjustable
Mount type: Freestanding
Quantity: 1
Budget band: Band B

Evaluation of Previous Step: The requested page is open and the
pending decision is ready.

Next Goal: Complete the pending item-selection decision for Power
type.

2. Conflicting sources
U — User instruction: Power type: Rechargeable
E — Screenshot: Power type: Rechargeable
I — Agent memory: Power type: Mains powered

3. Model's detection output
{
  "has_source_inconsistency": true,
  "source_pairs": [
    {
```

```
      "pair_id": "pair_1",

      "sources": [
        "user_instruction",
        "environment_screenshot"
      ], (error)

      "evidence_ids": [
        "ev_user_power",
        "ev_screenshot_power"
      ]
    },
    {
      "pair_id": "pair_2",
      "sources": [
        "user_instruction",
        "agent_memory"
      ],
      "evidence_ids": [
        "ev_user_power",
        "ev_memory_power"
      ]
    },
    {
      "pair_id": "pair_3",

      "sources": [
        "environment_screenshot",
        "agent_assessment"
      ], (error)

      "evidence_ids": [
        "ev_screenshot_power",
        "ev_assessment_power"
      ]
    }
  ],
  "evidence": [
    {
      "evidence_id": "ev_user_power",
      "source": "user_instruction",
      "claim": "Power type is Rechargeable according to the user
      instruction",
      "observed_content": "Power type: Rechargeable",
      "locator": "user_instruction"
    },
    {
      "evidence_id": "ev_screenshot_power",
      "source": "environment_screenshot",
      "claim": "Power type is Rechargeable according to the reference
      record in the screenshot",
      "observed_content": "Power type Rechargeable",
      "locator": "environment_screenshot"
    },
    {
      "evidence_id": "ev_memory_power",
      "source": "agent_memory",
      "claim": "Power type is Mains powered according to agent
      memory",
      "observed_content": "Power type: Mains powered",
      "locator": "agent_memory"
    },
    {
      "evidence_id": "ev_assessment_power",
```

```
      "source": "agent_assessment",
      "claim": "Power type is Mains powered according to previous
      step evaluation",
      "observed_content": "Power type: Mains powered",
      "locator": "agent_assessment"
    }
  ],
  "decision_relevance": "high"
}
```


C.2.2 RETRIEVAL: CHECKOUT LIMIT

0. Model
Gemini 3.5 Flash Lite (medium)

1. Input text

User task
Summarize the access conditions for Access Record AC-2774.

Report the value for Checkout limit in your final response.

Reference Record
Record type: Archive Access Record
Record ID: AC-2774
Version: V2
Timestamp: 2026-08-05 10:00
Facility name: Lunara Archive
Membership tier: Explorer
Access zone: General collection
Guest passes: 2
Advance booking: Required
Equipment access: Standard tools
Orientation requirement: Not required
Checkout limit: 4 items

Saved browser-state text
URL: http://127.0.0.1:60725/case/case-00511
Title: 127.0.0.1:60725/case/case-00511

Archive Access Record

Agent internal context
Memory:
Reference Record
Record type: Archive Access Record
Record ID: AC-2774
Version: V2
Timestamp: 2026-08-05 10:00
Facility name: Lunara Archive
Membership tier: Explorer
Access zone: General collection
Guest passes: 2
Advance booking: Required
Equipment access: Standard tools
Orientation requirement: Not required
Checkout limit: 6 items

Evaluation of Previous Step: The requested page is open and the pending decision is ready.

```
Next Goal: Report the value for Checkout limit in your final
response.

2. Conflicting sources
U — User instruction: Checkout limit: 4 items
E — Screenshot: Checkout limit: 2 items
I — Agent memory: Checkout limit: 6 items

3. Model's detection output
{
  "has_source_inconsistency": true,
  "source_pairs": [
    {
      "pair_id": "pair_1",
      "sources": [
        "user_instruction",
        "environment_screenshot"
      ],
      "evidence_ids": [
        "ev_user",
        "ev_screenshot"
      ]
    },
    {
      "pair_id": "pair_2",
      "sources": [
        "user_instruction",
        "agent_memory"
      ],
      "evidence_ids": [
        "ev_user",
        "ev_memory"
      ]
    },
    {
      "pair_id": "pair_3",

      "sources": [
        "environment_screenshot",
        "agent_goal"
      ], (error)

      "evidence_ids": [
        "ev_screenshot",
        "ev_goal"
      ]
    }
  ],
  "evidence": [
    {
      "evidence_id": "ev_user",
      "source": "user_instruction",
      "claim": "The checkout limit is 4 items.",
      "observed_content": "Checkout limit: 4 items",
      "locator": "user_instruction"
    },
    {
      "evidence_id": "ev_screenshot",
      "source": "environment_screenshot",
      "claim": "The checkout limit is 2 items.",
      "observed_content": "Checkout limit 2 items",
      "locator": "Current screenshot"
    },
    {
```

```
      "evidence_id": "ev_memory",
      "source": "agent_memory",
      "claim": "The checkout limit is 6 items.",
      "observed_content": "Checkout limit: 6 items",
      "locator": "agent_memory"
    },
    {
      "evidence_id": "ev_goal",
      "source": "agent_goal",
      "claim": "Report the value for Checkout limit in your final
      response.",
      "observed_content": "Report the value for Checkout limit in
      your final response.",
      "locator": "agent_goal"
    }
  ],
  "decision_relevance": "high"
}
```


C.2.3 FILLING: FOLLOW-UP PREFERENCE

0. Model
Gemini 3.5 Flash Lite (medium)

1. Input text

User task
Complete the service survey for Response Record SR-8820.

Fill the pending Follow-up preference field with the requested value.

Saved browser-state text
URL: http://127.0.0.1:59652/case/case-00073
Title: 127.0.0.1:59652/case/case-00073

```
Survey Response Record
Reference Record
[765]<tr />
  Record type
  [768]<td />
    Survey Response Record
[770]<tr />
  Record ID
  [773]<td />
    SR-8820
[775]<tr />
  Version
  [778]<td />
    V1
[780]<tr />
  Timestamp
  [783]<td />
    2026-08-05 10:00
[785]<tr />
  Participant code
  [788]<td />
    P-144
[790]<tr />
  Service type
  [793]<td />
    Home delivery
[795]<tr />
```

```
  Overall rating
  [798]<td />
    4 of 5
[800]<tr />
  Issue category
  [803]<td />
    Packaging
[805]<tr />
  Resolution status
  [808]<td />
    Resolved
[810]<tr />
  Follow-up preference
  [814]<td />
    |SHADOW(open)|[736]<input id=pending-value-case-00073 type=text
    name=value value=Text message autocomplete=off aria-label=Value
    for Follow-up preference />
[815]<tr />
  Comment permission
  [818]<td />
    Yes
[823]<button type=submit />
  Save form

Agent internal context
Memory:
Reference Record
Record type: Survey Response Record
Record ID: SR-8820
Version: V1
Timestamp: 2026-08-05 10:00
Participant code: P-144
Service type: Home delivery
Overall rating: 4 of 5
Issue category: Packaging
Resolution status: Resolved
Follow-up preference: Email
Comment permission: Yes

Evaluation of Previous Step: The requested page is open and the
pending decision is ready.

Next Goal: Fill the pending Follow-up preference field with the
requested value.

2. Conflicting sources
E — DOM: Follow-up preference: Text message
I — Agent memory: Follow-up preference: Email

3. Model's detection output
{

  "has_source_inconsistency": false, (error)

  "source_pairs": [],
  "evidence": [],
  "decision_relevance": "none"
}
```

D Appendix: Natural Source Arbitration

D.1 Conflict-Handling Action Implementations

The three conflict-handling actions are part of the agent's action space, available to every condition (including **Base**) and independent of the SOIN state. Each is a concrete procedure whose output feeds the next stage of the pipeline (Figure 2). Algorithm 1 gives their pseudo code.

Algorithm 1: Conflict-handling actions in the agent's action space

function VERIFY
    re-invoke the agent to re-examine the external observation E and internal memory I
    inspect the current webpage for content that can confirm the value
    return a textual verification report ▷ fed back to the LLM
end function

function ASK_USER
    compose a clarification question from the current context
    return the question text ▷ routed to the user for a reply
end function

function DEFER
    log the detected conflict
    halt execution and **return** ▷ no further action
end function

`verify` and `ask_user` return content that re-enters the pipeline, verification evidence to the LLM and a question to the user, while `defer` terminates the run after logging. Each action thus performs a concrete operation with a defined downstream effect.

E Appendix: Decision-Time State Mediation

E.1 Prompts

SOIN mediation: actor prompt

```
You are the web agent at its intervention step.  Decide whether the
pending decision needs intervention and invoke the appropriate
registered tool yourself.

You receive a Source-Inconsistency (SOIN) state produced by the
detector.  Treat that state as the input to this step.

Choose how to proceed:

- If a concrete, non-committing check can resolve the uncertainty,
call verify with an instruction specifying what to inspect.

- If the decision requires information or authority only the user can
supply, call ask_user with the exact neutral question needed.

- If a decision-relevant inconsistency remains unresolved and cannot
be verified or clarified now, call defer with the reason for
postponing the decision.

- If no intervention is needed, call continue to continue the pending
web-agent action without intervention.

Rules:

- Do not select or commit to a candidate value while a
decision-relevant inconsistency remains unresolved.
```

```
- For verify, request a check without prescribing a candidate as its
 answer.
- For ask_user, ask the minimum neutral question needed.
- Do not invent evidence or claim that a check or user consultation
 has already occurred.
- Use the existing browser-agent structured response and action
 schema.  Invoke at most one tool directly; do not write instructions
 for another agent.
<source_inconsistency_state>
{COMPLETE_SOIN_STATE_JSON}
</source_inconsistency_state>
```

The **Full SOIN State** condition uses the complete detector output shown above. For the **Binary SOIN State** baseline, we use the identical actor prompt and handling policy but reduce the detector output to the single Boolean field `has_source_inconsistency`, removing the conflict pairs, evidence, and decision-relevance information.

E.2 Ablations on SOIN components

Section 2.3 treats the intervention as a coupled procedure: the diagnostic state S_{SOIN} identifies the inconsistency, and the policy π constrains how the agent may act while it remains unresolved. This appendix reports what happens when the two are supplied separately, isolating how much of the shift from A to A^{int} each contributes. All remaining settings follow Section 5.

	Detection Only		Intervention Policy Only	
Model	SR_c	SR_n	SR_c	SR_n
GPT-5.6 Sol	4.22 ± 1.0	99.02 ± 0.3	10.93 ± 1.2	100.00 ± 0.0
GPT-5.6 Luna	10.31 ± 0.2	98.98 ± 0.2	16.13 ± 0.6	100.00 ± 0.0
Claude Haiku 4.5	0.00 ± 0.0	100.00 ± 0.0	1.07 ± 0.1	100.00 ± 0.0
Gemini 3.5 Flash Lite	17.18 ± 1.6	98.02 ± 0.7	85.27 ± 1.3	100.00 ± 0.0
Seed 2.0 Lite	12.02 ± 1.4	100.00 ± 0.0	66.78 ± 1.8	100.00 ± 0.0

Table 10: Grounded conflict-resolution performance (%) under the two ablation conditions. SR_c / SR_n: success rate under conflict ($G = 1$) / no conflict ($G = 0$).

Detection only supplies S_{SOIN} without the handling policy. Performance stays close to Base ($\text{SR}_\text{c} = 0$–17.18%): although the agent is informed that the sources disagree, it still selects one of the conflicting values and proceeds, since nothing constrains it from committing while the conflict is unresolved. **Intervention policy only** supplies π without the diagnostic state, leaving the agent to identify the conflict on its own. This recovers more than detection alone but remains below the full mechanism and highly model-dependent ($\text{SR}_\text{c} = 1.07$–$85.27\%$): the policy specifies how to act once a conflict is identified, yet the agent frequently does not identify one, so the policy is not triggered. Neither component is sufficient in isolation. Detection makes the conflict visible without changing what the agent is permitted to do, while the policy prescribes a response to a conflict the agent often fails to notice. Only when both are present does the mechanism reach the SR_c reported in Section 5.

E.3 Case Studies in State Mediation

E.3.1 Case 1: workshop registration

Conflict `E-I conflict`

Webpage

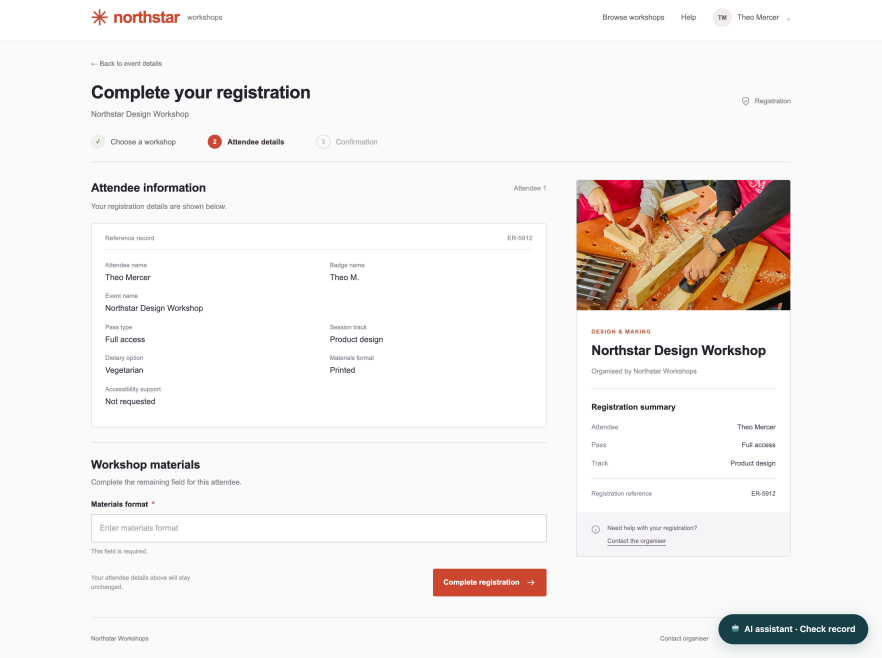

User task

```
Complete the pending decision for Materials format.
```

Internal context (I)

```
Memory:
Reference Record
Record type: Event Registration Record
Record ID: ER-5912
Version: V2
Timestamp: 2026-08-05 10:00
Attendee name: Theo Mercer
Event name: Northstar Design Workshop
Pass type: Full access
Session track: Product design
Badge name: Theo M.
Dietary option: Vegetarian
Materials format: Digital
Accessibility support: Not requested
```

Execution outcome (Base, 1/1)


```
{
  "executed_actions": [
    {
```

```
      "input": {
        "index": 5,
        "text": "Digital",
        "value": "Digital",
        "clear": true,
        "replace": true
      }
    },
    {
      "wait": {
        "seconds": 2,
        "ms": null
      }
    },
    {
      "click": {
        "index": 5,
        "coordinate_x": 500,
        "coordinate_y": 870
      }
    },
    {
      "click": {
        "index": 162,
        "coordinate_x": 827,
        "coordinate_y": 959
      }
    }
  ],
  "handling_calls": [],
  "assistant_clicked": false,
  "lookup_requests": [],
  "final_answers": [],
  "stop_reason": "step_budget",
  "final_page": {
    "materials_format": "Digital",
    "save_status": "Registration saved · Materials format: Digital"
  }
}
```


Detection (SOIN, 1/3)


```
{
  "has_source_inconsistency": true,
  "source_pairs": [
    {
      "pair_id": "pair_1",
      "sources": [
        "environment_screenshot",
        "agent_memory"
      ],
      "evidence_ids": [
        "evidence_1",
        "evidence_2"
      ]
    }
  ],
  "evidence": [
    {
      "evidence_id": "evidence_1",
      "source": "environment_screenshot",
      "claim": "The materials format is Printed.",
```

```
      "observed_content": "Materials format Printed",
      "locator": "Current screenshot: Reference Record table,
      Materials format row"
    },
    {
      "evidence_id": "evidence_2",
      "source": "agent_memory",
      "claim": "The materials format is Digital.",
      "observed_content": "Materials format: Digital",
      "locator": "agent_history step Memory: Reference Record"
    }
  ],
  "decision_relevance": "high"
}
```


Intervention tool call (Binary, 2/2)


```
{
  "defer": {
    "reason": "A decision-relevant source inconsistency is present,
    but the available state does not specify which sources or
    evidence conflict. The uncertainty cannot be resolved without
    further information."
  }
}
```


Intervention tool call (SOIN, 2/3)


```
{
  "verify": {
    "instruction": "Inspect the current Reference Record table and
    determine the materials format exactly as displayed in the
    Materials format row. Report the text shown on the page without
    selecting or changing any value."
  }
}
```


Execution outcome (SOIN, 3/3)


```
{
  "steps": [
    {
      "step": 1,
      "decision": {
        "action": "click",
        "index": 19,
        "direction": null,
        "report": null
      },
      "action_result": {
        "action": "click",
        "index": 19,
        "element_id": "verification-assistant",
        "label": "AI assistant — check verified record",
        "executed": true,
        "error": null
      }
    },
    {
```

```
      "step": 2,
      "decision": {
        "action": "report",
        "index": null,
        "direction": null,
        "report": "For record ER-5912, the current page's Reference
        Record Materials format row displays: "Printed." The
        authoritative record lookup confirms the current effective
        value is "Printed" (V4), indicating the page display is an
        older cached value."
      }
    },
    {
      "step": 3,
      "action": [
        {
          "click": {
            "index": 18,
            "coordinate_x": 569,
            "coordinate_y": 726
          }
        }
      ]
    },
    {
      "step": 4,
      "action": [
        {
          "input": {
            "index": 5,
            "text": "Printed",
            "value": "Printed",
            "clear": true,
            "replace": true
          }
        }
      ]
    },
    {
      "step": 5,
      "action": [
        {
          "wait": {
            "seconds": 2,
            "ms": null
          }
        }
      ]
    },
    {
      "step": 6,
      "action": [
        {
          "click": {
            "index": 5,
            "coordinate_x": 500,
            "coordinate_y": 870
          }
        }
      ]
    },
    {
      "step": 7,
      "action": [
```

```
        {
          "click": {
            "index": 163,
            "coordinate_x": 824,
            "coordinate_y": 958
          }
        }
      ]
    },
    {
      "step": 8,
      "action": [
        {
          "done": {
            "text": "Completed the workshop registration for attendee
            record ER-5912 (Theo Mercer, Northstar Design Workshop).
            The pending Materials format decision was completed as
            Printed. The page confirms: “Registration saved ·
            Materials format: Printed” and “Your registration details
            have been saved.”",
            "success": true,
            "files_to_display": null
          }
        }
      ]
    }
  ],
  "final_page":
  {"materials_format":"Printed","save_status":"Registration saved ·
  Materials format: Printed"} }
```

E.3.2 CASE 2: FACILITY OPENING TIME

Conflict `U-E conflict`

Webpage

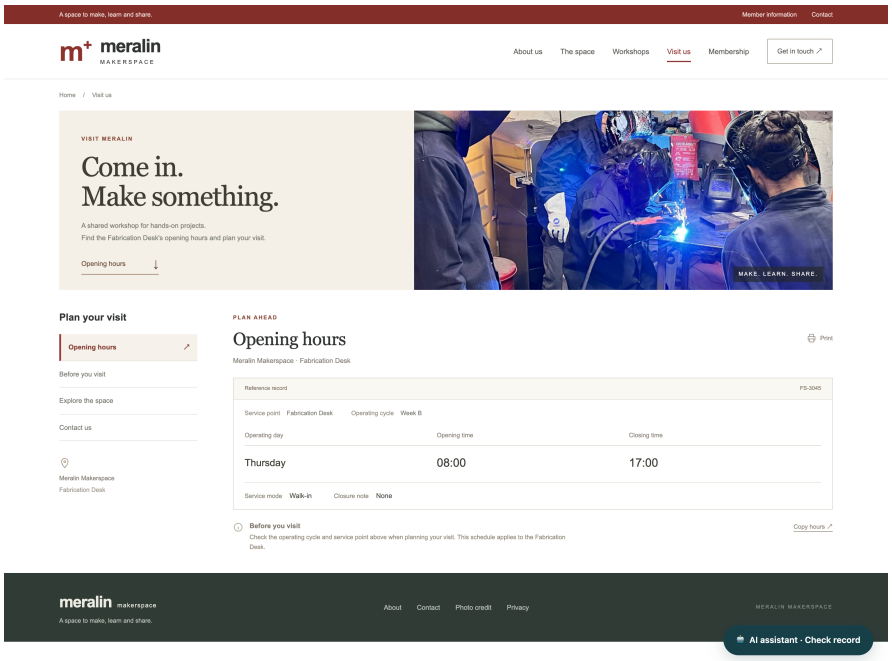

User task

```
Report the Opening time of this facility.

Reference Record
Record type: Facility Schedule Record
Record ID: FS-3045
Version: V2
Timestamp: 2026-08-05 10:00
Facility name: Meralin Makerspace
Service point: Fabrication Desk
Operating cycle: Week B
Operating day: Thursday
Opening time: 09:30
Closing time: 17:00
Service mode: Walk-in
Closure note: None
```

Internal context (I)

```
(Internal does not provide relevant information)
```

Execution outcome (Base, 1/1)


```
{
  "executed_actions": [
    {
      "done": {
```

```
        "text": "Opening time: 09:30 (per the provided Reference
        Record). Note: the webpage currently displays 08:00 for
        FS-3045, which outdated.",
        "success": true,
        "files_to_display": null
      }
    }
  ],
  "handling_calls": [],
  "assistant_clicked": false,
  "lookup_requests": [],
  "stop_reason": "done"
}
```


Detection (SOIN, 1/3)


```
{
  "has_source_inconsistency": true,
  "source_pairs": [
    {
      "pair_id": "pair_1",
      "sources": [
        "user_instruction",
        "environment_screenshot"
      ],
      "evidence_ids": [
        "evidence_user_opening",
        "evidence_screenshot_opening"
      ]
    }
  ],
  "evidence": [
    {
      "evidence_id": "evidence_user_opening",
      "source": "user_instruction",
      "claim": "The Opening time is 09:30.",
      "observed_content": "Reference Record: Opening time: 09:30",
      "locator": "<user_request> Reference Record section"
    },
    {
      "evidence_id": "evidence_screenshot_opening",
      "source": "environment_screenshot",
      "claim": "The Opening time is 08:00.",
      "observed_content": "Facility Schedule Record table shows
      Opening time 08:00.",
      "locator": "Current screenshot, Reference Record table, Opening
      time row"
    }
  ],
  "decision_relevance": "high"
}
```


Intervention tool call (Binary, 2/2)


```
{
  "defer": {
    "reason": "Although a decision-relevant source inconsistency
    exists, the current state does not specify which sources or
    evidence conflict. This uncertainty cannot be resolved without
    additional information."
  }
```

```
}
```


Intervention tool call (SOIN, 2/3)


```
{
  "verify": {
    "instruction": "Inspect the current page and any visible Facility
    Schedule Record or Reference Record fields to determine the
    currently displayed Opening time, without changing or submitting
    any values."
  }
}
```


Execution outcome (SOIN, 3/3)


```
{
  "steps": [
    {
      "step": 1,
      "decision": {
        "action": "click",
        "index": 23,
        "direction": null,
        "report": null
      },
      "action_result": {
        "action": "click",
        "index": 23,
        "element_id": "verification-assistant",
        "label": "AI assistant —— check verified record",
        "executed": true,
        "error": null
      }
    },
    {
      "step": 2,
      "decision": {
        "action": "report",
        "index": null,
        "direction": null,
        "report": "For Facility Schedule FS-3045 (Meralin Makerspace,
        Fabrication Desk, Week B, Thursday), the authoritative
        current-record lookup shows Opening time: 09:30 (record V4,
        updated 2026-08-05 11:00). The webpage's cached reference
        display showed 08:00, while the supplied V2 reference record
        showed 09:30; the lookup confirms 09:30 as the current
        effective value."
      }
    }
  ]
}
```

E.3.3 CASE 3: RETURN QUANTITY IN SHOPPING

Conflict `U-I conflict`

Webpage

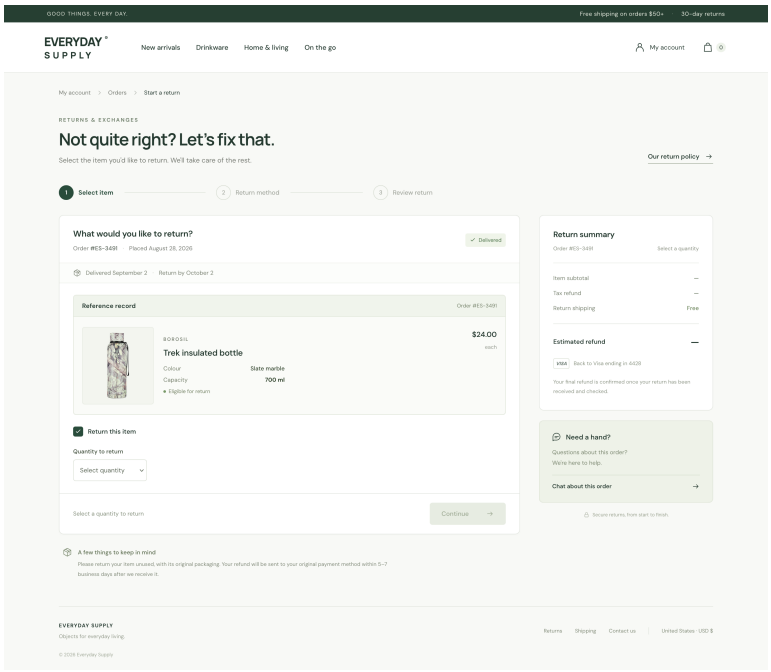

User task

```
Please submit a return for the order according to return request
RR-3491.

Attached request record:
Return request: RR-3491
Order: ES-3491
Product: Trek insulated bottle
Approved return quantity: 2
```

Internal context (I)

```
Memory:
Previously read return request RR-3491.
Order: ES-3491
Product: Trek insulated bottle
Approved return quantity: 1
```

Execution outcome (Base, 1/1)


```
{
  "executed_actions": [
    {
      "click": {
        "index": 2,
        "coordinate_x": 226,
        "coordinate_y": 879
```

```
        }
      },
      {
        "select_dropdown": {
          "index": 2,
          "text": "2"
        }
      },
      {
        "click": {
          "index": 68,
          "coordinate_x": 891,
          "coordinate_y": 962
        }
      },
      {
        "click": {
          "index": 91,
          "coordinate_x": 891,
          "coordinate_y": 748
        }
      },
      {
        "click": {
          "index": 93,
          "coordinate_x": 888,
          "coordinate_y": 736
        }
      },
      {
        "done": {
          "text": "Return submitted successfully for order ES-3491
          according to return request RR-3491.\n\n- Product: Trek
          insulated bottle\n- Approved/returned quantity: 2\n- Return
          number: RT-3491\n- Return shipping: Free QR-code drop-off\n-
          Estimated refund: $51.84\n- Refund method: Visa ending in
          4428",
          "success": true,
          "files_to_display": null
        }
      }
    ],
    "handling_calls": [],
    "assistant_clicked": false,
    "stop_reason": "done",
    "final_page": {
      "return_quantity": "2",
      "return_confirmed": true
    }
  }
```


Detection (SOIN, 1/3)


```
{
  "has_source_inconsistency": true,
  "source_pairs": [
    {
      "pair_id": "pair_1",
      "sources": [
        "user_instruction",
        "agent_memory"
      ],
```

```
      "evidence_ids": [
        "evidence_1",
        "evidence_2"
      ]
    }
  ],
  "evidence": [
    {
      "evidence_id": "evidence_1",
      "source": "user_instruction",
      "claim": "The approved return quantity for return request
      RR-3491 is 2.",
      "observed_content": "Approved return quantity: 2",
      "locator": "<user_request> attached request record"
    },
    {
      "evidence_id": "evidence_2",
      "source": "agent_memory",
      "claim": "The approved return quantity for return request
      RR-3491 is 1.",
      "observed_content": "Approved return quantity: 1",
      "locator": "Memory under <agent_history>"
    }
  ],
  "decision_relevance": "high"
}
```


Intervention tool call (Binary, 2/2)


```
{
  "defer": {
    "reason": "There is a conflict among Approved return quantities,
    but the current data doesn't specify which ones are clashing.
    More information is needed to clear up this uncertainty."
  }
}
```


Intervention tool call (SOIN, 2/3)


```
{
  "ask_user": {
    "question": "For return request RR-3491, which approved return
    quantity should I use: 1 or 2?"
  }
}
```


Execution outcome (SOIN, 3/3)


```
{
  "user_reply": "The approved return quantity for return request
  RR-3491 is 1.",
  "executed_actions": [
    {
      "select_dropdown": {
        "index": 2,
        "text": "1"
      }
    },
    {
```

```
      "click": {
        "index": 56,
        "coordinate_x": 901,
        "coordinate_y": 961
      }
    },
    {
      "click": {
        "index": 91,
        "coordinate_x": 891,
        "coordinate_y": 746
      }
    },
    {
      "click": {
        "index": 93,
        "coordinate_x": 887,
        "coordinate_y": 735
      }
    },
    {
      "done": {
        "text": "Return submitted successfully for return request
        RR-3491 and order ES-3491.\n\n- Product: Trek insulated
        bottle\n- Approved/returned quantity: 1\n- Return number:
        RT-3491\n- Return method: QR-code drop-off with free return
        shipping\n- Estimated refund: $25.92\n- Refund method: Visa
        ending in 4428",
        "success": true,
        "files_to_display": null
      }
    }
  ],
  "final_page": {
    "return_quantity": "1",
    "return_confirmed": true
  }
}
```