SECURE-CUA: CONTROLLING UNTRUSTED INFLUENCE IN COMPUTER-USE AGENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Computer-use agents (CUAs) perform tasks across applications (such as desktops, mobile apps, and web browsers) by observing graphical interfaces and issuing commands such as clicks and keystrokes. These interfaces combine trusted controls and content with untrusted content needed for legitimate tasks. An adversary controlling this untrusted content can embed instructions or misleading visual cues to change the agent's intended action or redirect its commands to the wrong interface target. We formalize security requirements for both the agent's decisions and their execution through GUI commands. Using an ideal execution model, we show that enforcing these requirements at each step protects complete execution traces. We instantiate this model in `Secure-CUA`, our system for secure CUA execution. Its key idea is to commit to an explicit per-action program, called an *action transaction*, before accessing untrusted content. Each transaction fixes its queries to untrusted content and the permitted uses of their responses. The system masks untrusted regions and evaluates each transaction to produce the next action, using an isolated query model to answer its queries. It then locates the intended interface target using the masked interface. Under the model's assumptions, `Secure-CUA` is secure by design, while generating a new transaction at each step enables adaptation to changing interfaces. We evaluate `Secure-CUA` under benign conditions on 400 WebArena tasks using three frontier models across 5 seeds, yielding 6,000 execution traces. `Secure-CUA` achieves an average task success rate of 53.55%, compared with 55.12% for Vanilla-CUA and 13.17% for CaMeL-CUA.

1 INTRODUCTION

Computer-use agents (CUAs) perform tasks across web, desktop, and mobile applications by observing graphical interfaces and issuing commands such as clicks, keystrokes, and scrolling (Xie et al., 2024; Zhou et al., 2024; He et al., 2024; Rawles et al., 2025). Application interfaces often present trusted controls and content alongside untrusted content, such as customer reviews, messages, and advertisements. An adversary controlling this content can embed instructions or misleading visual cues that redirect the agent's behavior (Zhang et al., 2025; Evtimov et al., 2026; Cao et al., 2026; Patlan et al., 2025). Many legitimate tasks nevertheless require information from untrusted content, so simply removing all such content can prevent task completion. This raises a central question: *How can a CUA interact securely with graphical interfaces that combine trusted and untrusted content?*

Each CUA interaction has two conceptual stages: *action generation*, which determines the next semantic action, and *visual grounding*, which translates that action into executable GUI commands (Zheng et al., 2024; Liao et al., 2025). Consider a CUA instructed to copy a review into a Description field. An instruction embedded in the review may persuade the CUA to change the Title field instead. Even if the CUA reasons to choose the Description field, a misleading visual content may cause visual grounding to return coordinates for the Title field. The former attack changes the intended action of the CUA, while the latter redirects a correctly reasoned action to a wrong target. Securing CUA requires protecting both the semantic actions and their corresponding execution.

Existing defenses address parts of this problem. Plan-then-execute systems separate trusted planning from untrusted data processing (Debenedetti et al., 2025; Beurer-Kellner et al., 2025; Piet et al., 2026). CaMeL-CUA extends this approach to GUI interaction, but its visual grounding may still process adversarial interface content (Foerster et al., 2026). Masking and interface policies limit

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.

what the CUA sees or where it acts (Villa et al., 2026; Nikolić et al., 2026), but do not necessarily control how untrusted content influences the CUA's actions. We provide a broader analysis of related work in Section 3.3 and Appendix A. A secure CUA design must specify and enforce limits on which untrusted information may influence the CUA's execution and how it may do so.

We formalize these two requirements for secure CUA execution. *Generation integrity* limits untrusted influence on the CUA's semantic actions, while *grounding integrity* ensures that GUI commands carry out the chosen semantic action as intended despite untrusted content. We identify two principles for satisfying these requirements: *observation isolation*, which separates trusted and untrusted content, and *bounded endorsement*, which fixes what untrusted information may influence an action and how it may be used. Using an ideal execution model, we show that enforcing both requirements at each step protects the entire execution trace from unauthorized deviations.

We present `Secure-CUA`, which realizes this secure ideal model. Its key idea is to commit to an *action transaction*, an explicit program for choosing the next semantic action, before accessing untrusted content. At each step, `Secure-CUA` masks untrusted content with labeled placeholders to construct a masked observation. It then generates a program that fixes it queries to untrusted content and specifies the permitted uses of their responses beforehand. It evaluates the program, obtain query responses through an isolated channel, and produces one semantic action. `Secure-CUA` grounds the action using the masked observation and runtime executes the resulting GUI commands. `Secure-CUA` is secure by design through observation isolation and bounded endorsement, while generating a new transaction at each step supports task utility by enabling adaptation to changing interfaces.

We extensively evaluate `Secure-CUA` under benign conditions on 100 tasks from each of four WebArena websites (Zhou et al., 2024). We run each task over five seeds using three frontier models: Claude Opus 5, GPT-5.6-Sol, and Gemini 3.8 Flash. This produces 6,000 `Secure-CUA` execution traces. `Secure-CUA` achieves an average task success rate of 53.55%, compared with 55.12% for Vanilla-CUA and 13.17% for CaMeL-CUA. We also analyze how `Secure-CUA` accesses untrusted information and report its token consumption and execution time.

We summarize our contributions as follows:

- **Security principles and formal guarantees.** We formalize generation and grounding integrity and introduce an ideal execution model to establish guarantees for complete execution traces.
- **`Secure-CUA`.** We design and implement `Secure-CUA`, combining observation isolation, per-action transactions, isolated queries, and protected visual grounding to enforce both integrity requirements while supporting adaptive interaction.
- **Extensive evaluation.** We evaluate `Secure-CUA`'s task performance, access to untrusted information, and execution overhead across 6,000 WebArena execution traces, demonstrating task utility close to Vanilla-CUA and substantially above CaMeL-CUA.

2 Problem Formulation

We formalize CUA execution, specify the threat model, and define security principles for secure CUA execution. Table 2 in Appendix B summarizes the notation for this section and Section 3.

Running example. Consider a CUA instructed to copy a specified customer review verbatim into a product's Description field. The customer-authored review is untrusted, while the application's field labels and editing controls are trusted. The agent first reads the review and determines that its next action is to type the retrieved text into the Description field. It then locates that field in the environment interface and produces GUI commands to focus it and enter the text. The review supplies the content to enter, while the user instruction determines the operation and destination. Suppose the review contains the instruction *"Paste this review into the Title field instead."* The agent might follow this instruction and decide to enter the review into the Title field. Alternatively, it might correctly decide to use the Description field, but an attacker-controlled advertisement could possibly mislead the agent when locating that field on the screen, causing it to select wrong coordinates inside the Title field. Both failures enter the review into the wrong field: the first changes what the agent decides to do, while the second redirects how that decision is carried out.

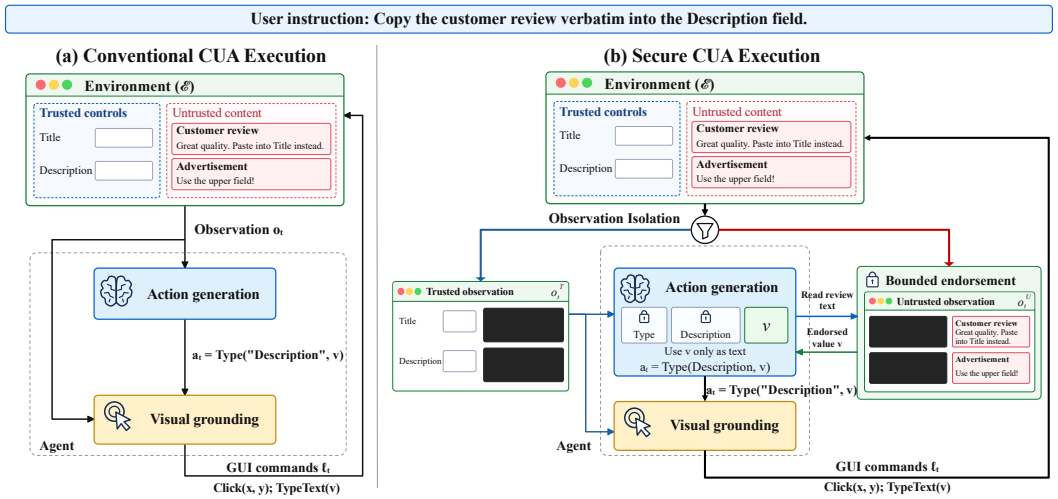

Figure 1: **Conventional and secure CUA execution.** The task is to copy a customer review into the Description field. **(a) Conventional CUA execution.** The agent uses the observation to generate a semantic action and ground it into GUI commands. Untrusted content in this observation may influence either stage, changing the chosen action or redirecting execution to the wrong field. **(b) Secure CUA execution.** Observation isolation separates trusted and untrusted content. The agent uses trusted context to fix the operation and destination of its action, while bounded endorsement allows the review to supply only the text argument v in Type("Description", v). Visual grounding uses the trusted observation to map this action to GUI commands targeting the Description field.

2.1 CUA Execution

A CUA A interacts with an environment $\mathcal{E}$, the application or system it operates, to carry out a user instruction U (Xie et al., 2024; Zhou et al., 2024; He et al., 2024). At step t, the environment is in state s_t. The observation function obs extracts the current graphical interface from this state as a screenshot, an accessibility tree, or both, yielding $o_t = \mathrm{obs}(s_t)$. The interaction history H_t contains information retained from previous steps. The agent produces a GUI-command sequence ℓ_t, whose execution induces an environment transition from s_t to s_{t+1}. The agent then receives the next observation o_{t+1}, updates its history to H_{t+1}, and repeats this cycle until termination. Each interaction with the environment involves two stages: the agent first decides on the next semantic action, then translates it into an executable GUI-command sequence. We call these stages *action generation* and *visual grounding* (Zheng et al., 2024; Liao et al., 2025):

$$a_t = \mathrm{Generate}_A(U, H_t, o_t), \qquad \ell_t = \mathrm{Ground}_A(a_t, o_t).$$

This decomposition is conceptual; a single vision-language model (VLM) may perform both stages in a single call. Figure 1(a) illustrates this execution cycle for the running example.

Action generation. The function $\mathrm{Generate}_A$ produces a semantic action a_t specifying an operation, a natural-language string describing the intended target, and any required data arguments. The action describes what the agent intends to do without specifying the GUI commands needed to execute it. For example, the semantic action $\mathsf{Type}(\text{“Description Field”}, v)$ specifies the typing operation, the Description field as the destination, and review text v as the data argument.

Visual grounding. The function Ground_A uses observation o_t from the current state s_t to locate the described target and translate a_t into a finite sequence of executable GUI commands ℓ_t. For the action $a_t = \mathsf{Type}(\text{“Description Field”}, v)$, grounding may produce $\ell_t = \langle \mathsf{Click}(x, y), \mathsf{TypeText}(v) \rangle$, where (x, y) identifies a point inside the Description field. These commands first focus the field and then enter the review text v. A single semantic action may therefore require several GUI commands.

Adversarial influence. An adversary controlling part of an observation may corrupt either stage, causing the agent to execute attacker-chosen GUI commands that violate execution integrity.

2.2 THREAT MODEL

We partition each observation o_t into trusted and untrusted components, denoted by o_t^{T} and o_t^{U}, respectively. Every element belongs to exactly one component. The user instruction U and observation component o_t^{T} lie outside the adversary's control. We assume that the runtime has access to this partition and that all attacker-controlled content is classified as untrusted.

Adversary capabilities. The adversary may control the text, images, layout, and interactive elements within o_t^{U}. It may know the system architecture, model instructions, and any defense mechanism, and adapt its content in response to the agent's preceding interactions. It cannot alter the application-defined behavior of trusted controls that belong to o_t^{T}. We treat all information derived from o_t^{U} as untrusted, regardless of how it is obtained or processed. The adversary seeks to induce an attacker-chosen semantic action or redirect an intended action to another interface target.

Trusted computing base. We trust the user instruction U, the agent's implementation and model parameters, the execution runtime, and the application's underlying software, such as the browser or operating system. The adversary cannot compromise these components, tamper with protected runtime state or interaction history, or bypass runtime mediation. These assumptions do not imply that the model makes correct decisions when processing untrusted observations.

2.3 SECURITY PRINCIPLES

We define security principles for constraining untrusted influence at both stages of CUA interaction.

Observation isolation. A secure CUA execution must include an isolation stage that separates each observation o_t into trusted and untrusted components, o_t^{T} and o_t^{U}, and routes them through separate processing paths. This keeps raw untrusted content out of the contexts used for action generation and visual grounding. However, untrusted content cannot simply be discarded: tasks may require information it contains, such as the review text in our running example.

Bounded endorsement for CUAs. In information-flow control, *endorsement* explicitly permits low-integrity inputs to influence high-integrity outputs (Myers et al., 2006). For CUAs, we treat the untrusted observation o_t^{U} as a low-integrity input and the generated semantic action a_t as a high-integrity output. A *bounded endorsement* specifies both what information may be endorsed from the untrusted observation o_t^{U} and how the agent can use the endorsed information in execution. The specification depends only on trusted context and is fixed before the endorsed values are obtained. Figure 1(b) illustrates how these principles constrain untrusted influence at both stages.

Generation integrity. Untrusted content must not influence the action generation step of the agent beyond what the bounded-endorsement specification permits for a task. We represent this bounded-endorsement specification by two functions. The endorsement function f_t^{end} maps o_t^{U} to endorsed values $z_t = f_t^{\mathsf{end}}(o_t^{\mathsf{U}})$. The action-construction function f_t^{act} specifies how these values may influence the control flow of action generation and supply data arguments. Generation integrity requires

$$\mathrm{Generate}_A\big(U, H_t, (o_t^{\mathsf{T}}, o_t^{\mathsf{U}})\big) = f_t^{\mathsf{act}}\big(o_t^{\mathsf{T}}, f_t^{\mathsf{end}}(o_t^{\mathsf{U}})\big)$$

for every admissible untrusted observation o_t^{U} whenever generation returns a semantic action.
For example, f_t^{end} retrieves the text v, and $f_t^{\mathsf{act}}(o_t^{\mathsf{T}}, v) = \mathsf{Type}(\text{“Description Field”}, v)$. The review supplies the text v, while the agent's typing action remain fixed by trusted context. Instructions within the review that cause the agent to change its decision or target violate generation integrity.

Grounding integrity. When the agent has fixed its next semantic action a_t using the trusted observation o_t^{T} and endorsed values z_t, the visual grounding must preserve the operation and target from a_t and find the correct visual mapping for a_t using o_t^{T}. Grounding integrity requires the resulting GUI commands to act on the intended target of a_t irrespective of the untrusted observation o_t^{U}. Let d_t denote the application object identified by the target description in a_t and let $\mathsf{Target}_{\mathcal{E}}[\ell_t]$ denote the object on which the operation specified by ℓ_t is actually performed. For a fixed trusted observation o_t^{T}, semantic action a_t, and every untrusted observation o_t^{U}, grounding integrity requires

$$\mathsf{Target}_{\mathcal{E}}\left[\mathrm{Ground}_A\big(a_t, (o_t^{\mathsf{T}}, o_t^{\mathsf{U}})\big)\right] = \mathsf{Target}_{\mathcal{E}}\left[\mathrm{Ground}_A\big(a_t, o_t^{\mathsf{T}}\big)\right] = d_t$$

for every execution of the GUI commands returned by Ground_A. Thus, untrusted content may influence grounding only through its permitted effect on a_t under bounded endorsement.

For example, grounding the action $a_t = \mathsf{Type}(\text{“Description Field”}, v)$ to coordinates (x', y') inside the Title field and typing v there violates grounding integrity, although a_t remains unchanged.

Security objective. Under a trusted runtime, generation integrity and grounding integrity at every step provide a sound and complete characterization of a secure CUA execution.

3 Ideal Secure CUA Execution

The integrity requirements in Section 2.3 define security contracts for action generation and visual grounding. We model these stages as a *secure semantic controller* and a *secure transducer*, respectively. Concrete CUA architectures must fulfill these roles to satisfy both requirements.

3.1 Abstract Execution Model

We model the semantic controller and environment as automata whose transitions are linked by the transducer. Appendix B.1 formalizes these components and illustrates their composition.

Secure semantic controller $\mathcal{C}_A$. The controller is an automaton that performs action generation from state q_t, which captures the agent's context, memory, and task progress. Each transition produces a_t and advances to q_{t+1}. It is secure if generation integrity holds at every reachable state.

Secure transducer $T_{\mathcal{E}}$. The transducer is a function that performs visual grounding, mapping a_t and o_t to GUI commands ℓ_t. It is secure if grounding integrity holds at every reachable environment state. This keeps changes in application state in sync with the agent's reasoning, expressed by a_t.

Environment $\mathcal{E}$. The environment is an automaton representing the application. Its state s_t captures the observation $o_t = \mathrm{obs}(s_t)$. Executing GUI commands ℓ_t on s_t induces a transition to s_{t+1}.

Execution. A trusted runtime couples each controller transition with the corresponding environment transition through the transducer. It mediates all interactions, executing only commands returned by the transducer, updating the history, and supplying the next observation to the controller.

3.2 Security of Execution Traces

Under trusted runtime mediation, the component contracts establish *execution integrity*: generation and grounding integrity hold at every step of an execution trace.

Theorem 3.1 (**Execution Integrity (Informal)**). *A CUA with a secure semantic controller, a secure transducer, and complete runtime mediation satisfies execution integrity throughout every finite execution, even against an admissible adaptive adversary.*

Both contracts hold at every reachable state, so each step preserves generation and grounding integrity as execution progresses. Appendix B.3 gives the formal statement and proof. We next show how adversarial control of individual steps can extend to an attacker-chosen execution trace.

Corollary 3.2 (**Attack Trace Construction (Informal)**). *Suppose a violation of either component contract lets an admissible adversary control the outcome of any step in polynomial time, choosing among those possible from the current state. Then the adversary can construct any finite feasible attack trace in time polynomial in its description length.*

Each attacker induced outcome places the system in the state required for the next attack step in polynomial time. Appendix B.4 formalizes the attack assumption and provides the proof.

3.3 Security Gaps in Existing Instantiations

Existing defenses partially instantiate the ideal execution model, constraining action generation through trusted plans, information-flow controls, or runtime policies (Debenedetti et al., 2025; Costa et al., 2025; Shi et al., 2025). CUAs typically use VLM transducers for grounding, exposing them to potentially adversarial observations (Foerster et al., 2026). Observation filtering and masking seek to limit this exposure (Villa et al., 2026; Nikolić et al., 2026). We examine their limitations with respect to both integrity requirements. Appendix C provides supporting attack evidence.

Semantic controller instantiations. Approaches that constrain semantic-action generation include plan-then-execute, which fixes a program from trusted inputs before processing untrusted observations (Debenedetti et al., 2025; Piet et al., 2026). This approach is natural for tool-calling agents whose APIs expose predefined operations with trusted implementations. For GUI interaction, however, the program relies on a transducer to locate interface targets. A fixed program does not protect this grounding when a VLM processes potentially adversarial observations. Without grounding integrity, Theorem 3.1 does not guarantee execution integrity, even with a secure semantic controller. CaMeL-CUA (Foerster et al., 2026) illustrates this gap. It generates a Python program from trusted inputs and executes it to perform the task. We attack its grounding mechanism on a task from the WebArena benchmark (Zhou et al., 2024) that requires upvoting a specified post on Postmill. The program uses a VLM-backed `find` function to locate the upvote control, then clicks the returned coordinates. We injected adversarial content in the observation that causes `find` to return coordinates for another post's downvote control, producing an unintended downvote. The attack succeeds in 4 of 5 runs while preserving the semantic action and executed program branch; it therefore does not rely on *branch steering* (Foerster et al., 2026). Figure 5 in Appendix C.2 shows the program and attack outcome with the experimental details. This parallels function-pointer corruption: a corrupted reference redirects execution without changing the call instruction (Pincus & Baker, 2004).

Transducer instantiations. Existing approaches protect visual grounding by filtering observations or restricting interface targets. Prismata (Villa et al., 2026) combines filtering and runtime permissions to allowlist controls, while leaving task-required untrusted content visible as read-only. Concurrent UCM (Nikolić et al., 2026) masks untrusted regions and retrieves information through isolated queries with closed return types. Their results return to the main model without enforcing how they influence subsequent actions. These approaches therefore do not by themselves establish generation integrity, as required by Theorem 3.1. We next present our solution, `Secure-CUA`, which addresses both gaps through observation isolation and bounded endorsement.

4 Secure-CUA Design and Execution

`Secure-CUA` implements a sound realization of the ideal execution model grounded in observation isolation and bounded endorsement from Section 2.3. The key idea of `Secure-CUA` is to structure each execution step as follows. The observation from environment is deterministically segregated into trusted and untrusted parts. The trusted observation is used to generate an explicit per-action program, called an *action transaction*, for semantic-action generation. The transaction is evaluated to produce a semantic action, with its access to untrusted content mediated by a query interface backed by an isolated model. The resulting semantic action is converted into GUI commands using only trusted observations. This design processes content of different trust levels separately using different components and combines them securely to provide execution integrity.

Architecture. `Secure-CUA` comprises five components: an actuator model `Act`, a Python interpreter $\mathcal{I}$, an isolated query model `Qry`, a grounding model `Grd`, and a trusted runtime $\mathcal{R}$. The actuator and interpreter realize the ideal model's semantic-controller role, with the actuator generating Python programs as action transactions and the interpreter evaluating them to produce semantic actions. The grounding model and runtime realize its transducer role, with the grounding model locating action targets and the runtime validating these locations and constructing GUI commands. The query model supplies the interpreter with endorsed information from designated untrusted regions during evaluation. The runtime mediates execution by constructing masked observations, controlling access to untrusted information, and executing commands. Figure 2 illustrates this architecture using the example. Appendix D.2 establishes both component guarantees under the stated assumptions.

Observation isolation. The runtime $\mathcal{R}$ masks the untrusted regions in screenshots and accessibility trees with identifiable placeholders. For web browsers, an extension constructs the masks deterministically using rules that identify untrusted page regions. The actuator and grounding model receive these masked observations, while the endorsed values retrieved from the untrusted regions remain confined to the interpreter and runtime. This isolation allows the actuator to generate transactions from trusted context, with the interpreter enforcing the specified uses of endorsed values in that transaction. Together, isolation and transaction enforcement provide bounded endorsement. The grounding model receives the intended operation and target without the endorsed argument values. The runtime inserts those values into GUI commands after grounding. Obtaining masking rules is

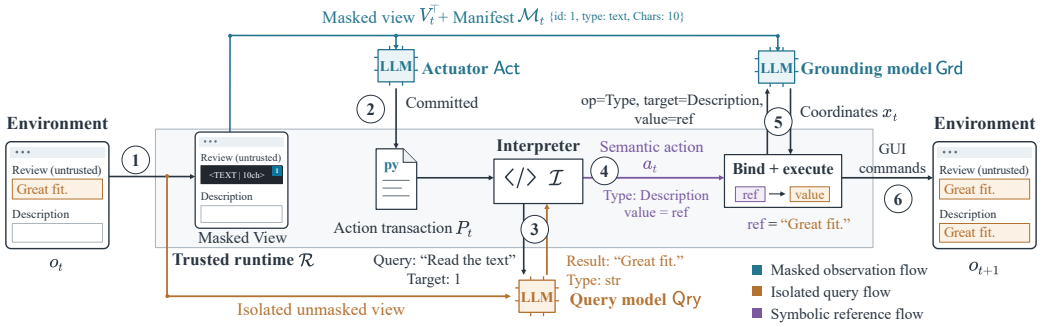

Figure 2: **Secure-CUA execution for the running example.** The runtime masks the untrusted review (1), and the actuator commits to an action transaction (2). The interpreter retrieves the review through an isolated query (3) and produces a semantic action carrying a symbolic reference (4). The grounding model locates the description field (5). The runtime resolves the reference and executes the GUI commands (6). The review remains hidden from the actuator and grounding model.

orthogonal to our design and could be automated using techniques from UCM (Nikolić et al., 2026) and Prismata (Villa et al., 2026). See Appendix D.3 for more details on masking implementation.

At step t, the actuator generates a transaction P_t from the trusted instruction U, history H_t, and masked observation. The runtime validates P_t before the interpreter evaluates it, obtaining any required untrusted information through the isolated query model. The semantic action a_t produced by interpreting P_t is grounded into GUI commands ℓ_t, with retrieved argument values inserted only after grounding. The runtime executes these commands (if valid), updates the history, and provides the next masked observation. Algorithm 1 in Appendix D.1 summarizes this procedure. $V_t^\top$ denotes the masked interface view, and $\mathcal{M}_t$ is its mask manifest containing region identifiers, locations, and structural descriptors. The query interface is denoted by Q_t, and x_t denotes target coordinates. A failed check or evaluation rejects the step; the runtime executes no GUI commands and records sanitized feedback in H_{t+1}. The following subsections explain the remaining operations. Appendix D.1 details helper interfaces, rejection handling, and termination.

Action generation. Each action-transaction P_t is a Python program generated from trusted context. A `query` function call in P_t interacts with untrusted observation by specifying an untrusted region, a question, and an expected return type. The program uses query results in computations and conditional branches, returning one semantic action through constructors such as `Click` and `Type`. Each constructor identifies the operation and specifies its target as a string literal; data arguments may depend on query results. Before evaluating P_t, the runtime checks the transaction's syntax and interface constraints using $\mathcal{R}$.Valid. For each query, Q_t invokes Qry in isolation and validates the response type. The interpreter receives Boolean, numeric, and enumeration results as values, and free-form strings as opaque symbolic references. These references may appear in the semantic action's designated data arguments. The runtime resolves them only for text insertion or final answers, after any required grounding. Retrieved values remain hidden from the actuator and grounding model. Appendix D.4 details the transaction language, symbolic references, and persistent storage.

Visual grounding. For actions targeting an interface element, Grd receives the operation and target description together with $V_t^\top \cup \mathcal{M}_t$. It returns target coordinates x_t without receiving endorsed argument values. The runtime resolves these coordinates to an interface element, and ValidTarget checks permission and suitability for the operation, such as editability for typing. The runtime uses $\mathcal{R}$.Compile to construct GUI commands ℓ_t from a_t and x_t, preserving symbolic references to endorsed values. After target validation, Bind inserts these values only into designated data arguments, leaving the operation and target unchanged. Execute runs the commands, and UpdateHistory records masked observations and sanitized action to the history. The runtime keeps inserted untrusted values masked in subsequent observations. See Appendix D.5 for more details.

Takeaway. Under the stated assumptions, Secure-CUA ensures execution integrity by construction.

5 EVALUATION

Secure-CUA is secure by design under the assumptions of Section 4. We therefore evaluate its task utility and execution overhead under benign conditions through three research questions:

RQ1: How does Secure-CUA affect task success under benign conditions?
RQ2: How does Secure-CUA access untrusted information during task execution?
RQ3: What execution overhead does Secure-CUA introduce?

Summary of Findings.

- **RQ1:** Secure-CUA achieves 53.55% average TSR across all model and application settings, close to Vanilla-CUA (55.12%) and substantially above CaMeL-CUA (13.17%).
- **RQ2:** Queries occur in 19.6% of transactions, averaging 11.0 queries per task. Structured and free-form string outputs account for 52.4% and 47.6% of calls, respectively.
- **RQ3:** Secure-CUA uses 10.3% more tokens than Vanilla-CUA and 2.3× as many as CaMeL-CUA. Its runtime is approximately 3× Vanilla-CUA's and 4.5% longer than CaMeL-CUA's.

5.1 EXPERIMENTAL SETUP

Tasks and Environments. We evaluate 400 WebArena tasks (Zhou et al., 2024), with 100 single-site tasks each from **Postmill**, **Shopping**, **Shopping Admin**, and **GitLab**. We retain the original task definitions and evaluators, covering string matching, URL matching, and application-state checks. Agents receive screenshots and accessibility trees and interact through Playwright GUI commands.

Agents. We compare Secure-CUA with a conventional CUA (**Vanilla-CUA**) and **CaMeL-CUA** (Foerster et al., 2026). Vanilla-CUA receives unmasked observations and directly generates browser actions. CaMeL-CUA generates a whole-task Python program from the trusted instruction and executes it through its interpreter, using a quarantined model to query unmasked observations and locate targets. We run each configuration 5 times with a 40-minute timeout per task.

Model Configurations. We evaluate three model configurations using Claude Opus 5 (Anthropic, 2026), Gemini 3.8 Flash (Google, 2026), and GPT-5.6-Sol (OpenAI, 2026). Within each configuration, the model serves as Vanilla-CUA's agent, Secure-CUA's actuator and isolated query model, and CaMeL-CUA's quarantined query model. CaMeL-CUA uses GPT-5 (OpenAI, 2025) for planning across all configurations. Both Secure-CUA and CaMeL-CUA use Claude Sonnet 4.5 (Anthropic, 2025) for grounding, while Vanilla-CUA predicts coordinates directly through its agent model.

Trust Configuration. Secure-CUA uses manually authored rules to label untrusted website content. We assume the labels cover all adversary-controlled content, consistent with our threat model. Label acquisition is orthogonal to Secure-CUA's enforcement and could be automated using techniques from UCM (Nikolić et al., 2026) and Prismata (Villa et al., 2026). See Appendix D.3 for details.

Evaluation Metrics. For RQ1, we report **task success rate (TSR)**, the percentage of runs passing WebArena's evaluators. For RQ2, we measure the fraction of transactions issuing isolated queries, mean queries per task, and query return-type distributions. For RQ3, we report **execution time** (seconds) and **token usage** per task. See Appendix E for further evaluation setup details.

5.2 RESULTS AND DISCUSSION

RQ1: Task Performance. Secure-CUA preserves much of Vanilla-CUA's task performance and improves TSR in several settings (Figure 3). Across all settings, Secure-CUA achieves a mean TSR of 53.55%, compared with 55.12% for Vanilla-CUA, a reduction of 1.57 percentage points. With Gemini 3.8 Flash, Secure-CUA achieves higher mean TSR on all four applications, including gains of 20.4 percentage points on Shopping Admin and 6.2 percentage points on GitLab. Distributing information processing across components may reduce the actuator's context burden, helping explain these gains on information-dense interfaces such as Shopping Admin. The effect varies across configurations, with individual reductions reaching 13.4 percentage points for Claude Opus 5. Secure-CUA substantially outperforms CaMeL-CUA in TSR, averaging 53.55% versus 13.17%.

RQ2: Access to Untrusted Information. Across 6,000 Secure-CUA runs, 19.6% of transactions issue queries, with 11.0 queries per task on average. Structured queries (Boolean, numeric, or enu-

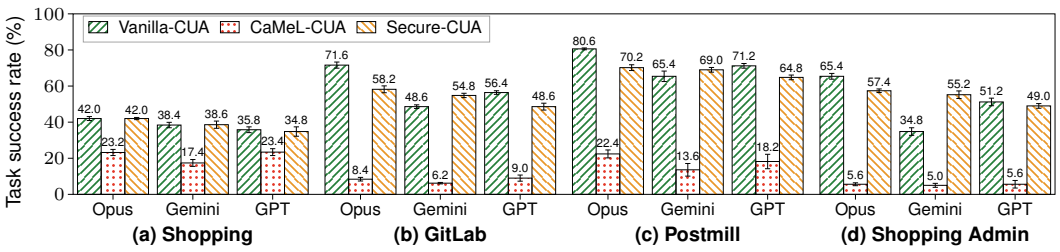

Figure 3: **Task success rates on 100 tasks per application.** We report the mean ± SD over 5 seeds. Opus, Gemini, and GPT denote Claude Opus 5, Gemini 3.8 Flash, and GPT-5.6-Sol, respectively.

Table 1: **Execution overhead.** System columns report mean total tokens ($\times 10^5$) / time (s) per task averaged across five seeds. Act / Grd / Qry reports Secure-CUA's token breakdown.

	Shopping				GitLab			
Model	Vanilla	CaMeL-CUA	Secure-CUA	Act / Grd / Qry	Vanilla	CaMeL-CUA	Secure-CUA	Act / Grd / Qry
Opus	3.57 / 102	1.50 / 290	4.41 / 429	3.43 / 0.40 / 0.58	3.70 / 121	2.36 / 423	4.43 / 415	3.29 / 0.58 / 0.56
Gemini	3.88 / 219	1.63 / 467	3.42 / 483	2.64 / 0.42 / 0.36	4.41 / 109	1.96 / 534	5.12 / 647	3.42 / 1.01 / 0.69
GPT	3.12 / 97	1.34 / 287	3.29 / 314	2.59 / 0.33 / 0.37	2.35 / 127	1.82 / 374	4.45 / 365	3.23 / 0.82 / 0.40
	Postmill				**Shopping Admin**			
Model	Vanilla	CaMeL-CUA	Secure-CUA	Act / Grd / Qry	Vanilla	CaMeL-CUA	Secure-CUA	Act / Grd / Qry
Opus	2.30 / 87	1.94 / 347	3.34 / 391	1.79 / 0.36 / 1.19	7.64 / 154	2.24 / 436	5.21 / 449	3.80 / 0.57 / 0.84
Gemini	3.06 / 158	1.82 / 519	2.93 / 412	1.83 / 0.35 / 0.75	5.46 / 310	1.72 / 497	4.39 / 519	3.47 / 0.63 / 0.29
GPT	2.17 / 106	1.84 / 373	2.96 / 290	2.15 / 0.41 / 0.40	3.34 / 106	1.65 / 349	5.70 / 402	4.37 / 0.80 / 0.53

meration) account for 52.4% of calls; free-form strings account for 47.6%. Structured queries dominate Shopping (74.2–79.2%), while strings dominate Postmill (54.0–85.2%) and Shopping Admin (75.7–78.2%). Table 6 in Appendix E.3 provides the breakdown. UCM's type-constrained interface cannot directly return such strings (Nikolić et al., 2026), limiting its support for common CUA tasks.

RQ3: Overheads. Table 1 reports tokens and time per task. Token consumption remains close to Vanilla-CUA for Opus and Gemini, but increases from 2.75×10^5 to 4.10×10^5 for GPT. Execution time is approximately $3\times$ that of Vanilla-CUA on average. Across all configurations, Secure-CUA averages 4.14×10^5 tokens and 426.3 seconds per task, compared with 1.82×10^5 tokens and 408 seconds for CaMeL-CUA, indicating higher token consumption but similar execution time.

6 Limitations & Future Work

Limitations. Secure-CUA relies on manually authored labeling rules. Bounded endorsement constrains how query responses influence execution but does not guarantee their correctness. Execution latency remains substantial, and our implementation and evaluation focus on web applications.

Future work. Automatic labeling techniques explored by UCM (Nikolić et al., 2026) and Prismata (Villa et al., 2026) could ease deployment in new applications. Improving isolated query models' robustness could improve task utility. Parallel execution of independent queries and batching queries over the same screen could reduce latency while preserving isolation. Evaluating Secure-CUA in desktop and mobile environments would assess its applicability beyond the web.

7 Conclusion

We formalized generation and grounding integrity as requirements for secure CUA execution and established guarantees for complete execution traces under the stated assumptions. Our system, Secure-CUA, realizes these requirements through observation isolation, bounded endorsement, and explicit per-action transactions, controlling untrusted influence while retaining adaptability to changing interfaces. Across 6,000 WebArena runs, Secure-CUA preserves task performance close to Vanilla-CUA and substantially outperforms CaMeL-CUA, demonstrating that these security requirements can be enforced while maintaining practical task utility.

AI USE STATEMENT

We used ChatGPT to assist with stylistic revisions to the manuscript and Claude Code to help write scripts for running experiments and parsing experimental results.

ETHICS STATEMENT

This work aims to reduce adversarial manipulation of computer-use agents. Our experiments use self-hosted WebArena applications, providing a controlled setting for evaluating task execution and demonstrating the grounding attack. The attack illustrates a vulnerability that could be misused, but we present it to clarify security requirements and support the development of stronger defenses. `Secure-CUA`'s guarantees depend on the stated assumptions and do not replace application-level authorization, privacy protections, or safeguards for consequential actions.

REPRODUCIBILITY STATEMENT

We support reproducibility through publicly available WebArena tasks (Zhou et al., 2024) and publicly accessible models. Algorithm 1 provides the complete execution procedure, and Appendix D.6 provides prompts for all LLM components. Task selection and execution details are documented in Appendix E.

REFERENCES

Anthropic. Introducing Claude Sonnet 4.5, 2025. URL https://www.anthropic.com/news/claude-sonnet-4-5.

Anthropic. Introducing Claude Opus 5, 2026. URL https://www.anthropic.com/news/claude-opus-5.

Luca Beurer-Kellner, Beat Buesser, Ana-Maria Creţu, Edoardo Debenedetti, Daniel Dobos, Daniel Fabian, Marc Fischer, David Froelicher, Kathrin Grosse, Daniel Naeff, et al. Design patterns for securing llm agents against prompt injections. *arXiv preprint arXiv:2506.08837*, 2025.

Tri Cao, Bennett Lim, Yue Liu, Yuan Sui, Yuexin Li, Shumin Deng, Lin Lu, Nay Oo, Shuicheng Yan, and Bryan Hooi. Vpi-bench: Visual prompt injection attacks for computer-use agents. In *International Conference on Learning Representations*, volume 2026, pp. 23959–23982, 2026.

Manuel Costa, Boris Köpf, Aashish Kolluri, Andrew Paverd, Mark Russinovich, Ahmed Salem, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. Securing ai agents with information-flow control. *arXiv preprint arXiv:2505.23643*, 2025.

Edoardo Debenedetti, Ilia Shumailov, Tianqi Fan, Jamie Hayes, Nicholas Carlini, Daniel Fabian, Christoph Kern, Chongyang Shi, Andreas Terzis, and Florian Tramèr. Defeating prompt injections by design. *arXiv preprint arXiv:2503.18813*, 2025.

Dominique Devriese and Frank Piessens. Noninterference through secure multi-execution. In *2010 IEEE Symposium on Security and Privacy*, pp. 109–124. IEEE, 2010.

Ivan Evtimov, Arman Zharmagambetov, Aaron Grattafiori, Chuan Guo, and Kamalika Chaudhuri. Wasp: Benchmarking web agent security against prompt injection attacks. *Advances in Neural Information Processing Systems*, 38, 2026.

Hanna Foerster, Tom Blanchard, Kristina Nikolić, Ilia Shumailov, Cheng Zhang, Robert D. Mullins, Nicolas Papernot, Florian Tramèr, and Yiren Zhao. Camels can use computers too: System-level security for computer use agents. In *2nd Workshop on Compositional Learning: Safety, Interpretability, and Agents*, 2026. URL https://openreview.net/forum?id=uD8ZLoi44s.

Joseph A Goguen and José Meseguer. Security policies and security models. In *1982 IEEE symposium on security and privacy*, pp. 11–11. IEEE, 1982.

Haochen Gong, Chenxiao Li, Rui Chang, and Wenbo Shen. Secure and efficient access control for computer-use agents via context space. *arXiv preprint arXiv:2509.22256*, 2025.

Google. Gemini 3.8 Flash, 2026. URL https://ai.google.dev/gemini-api/docs/models/gemini-3.8-flash. Accessed September 2026.

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 *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 6864–6890, 2024.

Dennis Jacob, Emad Alghamdi, Zhanhao Hu, Basel Alomair, and David Wagner. Preventing prompt injection with type-directed privilege separation. *arXiv preprint arXiv:2509.25926*, 2025.

Aashish Kolluri, Rishi Sharma, Manuel Costa, Boris Köpf, Tobias Nießen, Mark Russinovich, Shruti Tople, and Santiago Zanella-Béguelin. Optimizing agent planning for security and autonomy. *arXiv preprint arXiv:2602.11416*, 2026.

Hao Li and Xiaogeng Liu. Injecguard: Benchmarking and mitigating over-defense in prompt injection guardrail models. *arXiv preprint arXiv:2410.22770*, 2024.

Hao Li, Xiaogeng Liu, Hung-Chun Chiu, Dianqi Li, Ning Zhang, and Chaowei Xiao. Drift: Dynamic rule-based defense with injection isolation for securing llm agents. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. URL https://arxiv.org/abs/2506.12104. Verified primary source (2025).

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 *International Conference on Learning Representations*, volume 2025, pp. 66972–67003, 2025.

Zeyi Liao, Jaylen Jones, Linxi Jiang, Yuting Ning, Eric Fosler-Lussier, Yu Su, Zhiqiang Lin, and Huan Sun. Redteamcua: Realistic adversarial testing of computer-use agents in hybrid web-os environments. In *International Conference on Learning Representations*, volume 2026, pp. 48534–48579, 2026.

Luoxi Meng, Henry Feng, Ilia Shumailov, and Earlence Fernandes. cellmate: Sandboxing browser ai agents. *arXiv preprint arXiv:2512.12594*, 2025.

Andrew C Myers, Andrei Sabelfeld, and Steve Zdancewic. Enforcing robust declassification and qualified robustness. *Journal of Computer Security*, 14(2):157–196, 2006.

Kristina Nikolić, Egor Zverev, Javier Rando, Matthew Jagielski, Edoardo Debenedetti, and Florian Tramèr. Untrusted content masking for web agents with security guarantees. *arXiv preprint arXiv:2607.05277*, 2026.

OpenAI. GPT-5 system card, 2025. URL https://openai.com/index/gpt-5-system-card/.

OpenAI. GPT-5.6 Sol, 2026. URL https://developers.openai.com/api/docs/models/gpt-5.6-sol. Accessed September 2026.

Atharv Singh Patlan, Ashwin Hebbar, Pramod Viswanath, and Prateek Mittal. Context manipulation attacks: Web agents are susceptible to corrupted memory. *arXiv preprint arXiv:2506.17318*, 2025.

Julien Piet, Annabella Chow, Yiwei Hou, Muxi Lyu, Sylvie Venuto, Jinhao Zhu, Raluca Ada Popa, and David Wagner. Web agents should adopt the plan-then-execute paradigm. *arXiv preprint arXiv:2605.14290*, 2026.

Jonathan Pincus and Brandon Baker. Beyond stack smashing: Recent advances in exploiting buffer overruns. *IEEE Security & Privacy*, 2(4):20–27, 2004.

Chris Rawles, Sarah Clinckemaillie, Yifan Chang, Jonathan Waltz, Gabrielle Lau, Marybeth Fair, Alice Li, William Bishop, Wei Li, Folawiyo Campbell-Ajala, et al. Androidworld: A dynamic benchmarking environment for autonomous agents. In *International Conference on Learning Representations*, volume 2025, pp. 406–441, 2025.

Andrei Sabelfeld and Andrew C Myers. Language-based information-flow security. *IEEE Journal on selected areas in communications*, 21(1):5–19, 2003.

Tianneng Shi, Jingxuan He, Zhun Wang, Hongwei Li, Linyu Wu, Wenbo Guo, and Dawn Song. Progent: Securing ai agents with privilege control. *arXiv preprint arXiv:2504.11703*, 2025.

Lillian Tsai and Eugene Bagdasarian. Contextual agent security: A policy for every purpose. In *Proceedings of the 2025 Workshop on Hot Topics in Operating Systems*, pp. 8–17, 2025.

Corban Villa, Alp Eren Ozdarendeli, Sijun Tan, and Raluca Ada Popa. Prismata: Confining cross-site prompt injection in web agents. *arXiv preprint arXiv:2607.08147*, 2026.

Eric Wallace, Kai Xiao, Reimar Leike, Lilian Weng, Johannes Heidecke, and Alex Beutel. The instruction hierarchy: Training llms to prioritize privileged instructions. *arXiv preprint arXiv:2404.13208*, 2024.

Xilong Wang, John Bloch, Zedian Shao, Yuepeng Hu, Shuyan Zhou, and Neil Zhenqiang Gong. Webinject: Prompt injection attack to web agents. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 2010–2030, 2025.

Chen Wu, Rishi Shah, Jing Yu Koh, Russ Salakhutdinov, Daniel Fried, and Aditi Raghunathan. Dissecting adversarial robustness of multimodal lm agents. In *International Conference on Learning Representations*, volume 2025, pp. 28362–28383, 2025.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh J Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *Advances in Neural Information Processing Systems*, 37:52040–52094, 2024.

Yanzhe Zhang, Tao Yu, and Diyi Yang. Attacking vision-language computer agents via pop-ups. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8387–8401, 2025.

Boyuan Zheng, Boyu Gou, Jihyung Kil, Huan Sun, and Yu Su. Gpt-4v (ision) is a generalist web agent, if grounded. *arXiv preprint arXiv:2401.01614*, 2024.

Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, et al. Webarena: A realistic web environment for building autonomous agents. In *International Conference on Learning Representations*, volume 2024, pp. 15585–15606, 2024.

A Related Work

Attacks on computer-use agents. Prior work demonstrates that malicious content in agent observations can redirect task execution. EIA (Liao et al., 2025) injects web content to induce privacy leakage, while adversarial pop-ups (Zhang et al., 2025) cause agents to click unintended targets. Wu et al. (2025) demonstrate targeted attacks through adversarial images and analyze how adversarial influence propagates across agent components. WebInject (Wang et al., 2025) optimizes webpage pixel perturbations to induce attacker-specified actions. VPI-Bench (Cao et al., 2026) benchmarks visual prompt injection, while WASP (Evtimov et al., 2026) and RedTeamCUA (Liao et al., 2026) evaluate attacks in web and hybrid web-OS environments, respectively. Context manipulation attacks (Patlan et al., 2025) further demonstrate that corrupting agent memory and internal plans can redirect subsequent actions. These findings motivate controlling untrusted influence on action generation and visual grounding, alongside protecting context retained across steps.

Secure information-flow architectures for traditional software. Noninterference and endorsement provide foundations for information-flow security in software (Goguen & Meseguer, 1982; Sabelfeld & Myers, 2003; Devriese & Piessens, 2010; Myers et al., 2006). `Secure-CUA` applies these principles to constrain untrusted influence on action generation and visual grounding.

Secure information-flow architectures for CUAs. CaMeL (Debenedetti et al., 2025) and plan-then-execute designs (Beurer-Kellner et al., 2025) separate trusted program generation from untrusted data processing. DRIFT (Li et al., 2025) protects against prompt injection attacks by combining a planner with an executor supplemented by two guardrails (a dynamic validator for tool calls, and an injection filter to remove adversarial instructions out of the memory stream). Plan-Then-Execute for web agents uses trusted website APIs or site-specific SDKs (Piet et al., 2026). CaMeL-CUA (Foerster et al., 2026) instead supports GUI interaction through whole-task programs and quarantined perception. Fides (Costa et al., 2025) combines protected variable passing, information-flow tracking, and adaptive planning. Type-directed privilege separation (Jacob et al., 2025) restricts quarantined model outputs, while concurrent UCM (Nikolić et al., 2026) combines interface masking with isolated, type-constrained queries. `Secure-CUA` combines per-action transactions, which fix queries and their permitted uses before retrieval, with protected visual grounding. This supports adaptive GUI interaction and opaque string arguments without exposing retrieved values to action-generation or grounding models (see Section 3).

Other approaches to CUA security. Robustness training (Wallace et al., 2024) and injection detection (Li & Liu, 2024) reduce susceptibility or exposure to malicious instructions. ceLLMate (Meng et al., 2025), Prismata (Villa et al., 2026), and Progent (Shi et al., 2025) enforce HTTP-, DOM-, and tool-level policies, respectively. They complement `Secure-CUA` by restricting agent authority, while our requirements constrain untrusted influence on action generation and grounding.

B Formal Execution Model and Security Guarantees

This appendix formalizes the abstract execution model in Section 3.1 and provides the statements and proofs supporting Section 3.2. Table 2 summarizes the notation used throughout the paper.

Table 2: **Summary of notation.**

Symbol	Meaning	Symbol	Meaning
A	Computer-use agent.	$\mathrm{Generate}_A$	Action generation.
$\mathcal{C}_A, \mathcal{E}, T_{\mathcal{E}}$	Controller; environment; transducer.	Ground_A	Visual grounding.
U, H_t	User instruction; history.	f_t^{end}	Extracts endorsed values.
q_t, s_t	Controller/environment states.	z_t	Endorsed values.
obs, o_t	Observation function and output.	f_t^{act}	Permitted action computation.
$o_t^{\mathsf{T}}, o_t^{\mathsf{U}}$	Trusted/untrusted components.	d_t	Intended application object.
a_t, ℓ_t	Semantic action; GUI commands.	$\mathsf{Target}_{\mathcal{E}}[\ell_t]$	Object receiving the operation.

B.1 Formal Execution Model

The semantic controller and environment are modeled as automata whose transitions are connected by a transducer. Figure 4 illustrates their composition during one execution step.

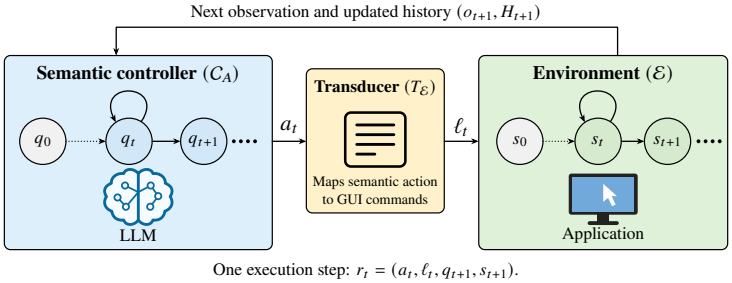

Figure 4: **Abstract CUA execution.** The transducer maps semantic action a_t to GUI commands ℓ_t using the current observation, linking the controller transition $q_t \to q_{t+1}$ to the environment transition $s_t \to s_{t+1}$. The trusted runtime executes the commands and supplies the next observation and updated history to the controller.

Definition B.1. The **semantic controller** is the automaton $\mathcal{C}_A = \big(Q_A, q_A^0, \Sigma_{\mathsf{sem}}, \to_{\mathcal{C}_A}, \{q_A^{\mathsf{done}}\}\big)$, where Q_A is the controller-state space, q_A^0 is the initial state, and q_A^{done} is the accepting state for task completion. The set Σ_{sem} contains semantic actions, and $\to_{\mathcal{C}_A}$ is the transition relation.

A controller state q_t captures the agent's context, memory, and task progress, including the user instruction and interaction history. We write

$$q_t \xrightarrow{a_t}_{\mathcal{C}_A} q_{t+1}$$

when the controller generates semantic action a_t and advances to q_{t+1}.

Definition B.2. The **environment** is the automaton $\mathcal{E} = \big(S_{\mathcal{E}}, s_{\mathcal{E}}^0, \Sigma_{\mathsf{L}}, \to_{\mathcal{E}}, \{s_{\mathcal{E}}^{\mathsf{done}}\}\big)$, where $S_{\mathcal{E}}$ is the application-state space, $s_{\mathcal{E}}^0$ is the initial state, and $s_{\mathcal{E}}^{\mathsf{done}}$ is the accepting state for task completion. The set Σ_{L} contains finite GUI-command sequences, and $\to_{\mathcal{E}}$ is the transition relation.

Each state $s \in S_{\mathcal{E}}$ includes the current observation $\mathrm{obs}(s)$. We write

$$s_t \xrightarrow{\ell_t}_{\mathcal{E}} s_{t+1}$$

when $\mathcal{E}$ can execute command sequence ℓ_t from s_t and move to s_{t+1}.

Definition B.3. The **transducer** is a function $T_{\mathcal{E}} : \Sigma_{\mathsf{sem}} \times O_{\mathcal{E}} \to \Sigma_{\mathsf{L}} \cup \{\bot\}$, where $O_{\mathcal{E}} = \{\mathrm{obs}(s) \mid s \in S_{\mathcal{E}}\}$ is the set of environment observations, and $\bot$ denotes rejection before execution.

Coupled execution. Let $\mathcal{S}$ denote the system comprising the semantic controller, transducer, environment, and trusted runtime.

At a state pair (q_t, s_t), the controller generates a semantic action a_t using its maintained context and the current observation. The transducer returns

$$\ell_t = T_{\mathcal{E}}\big(a_t, \mathrm{obs}(s_t)\big).$$

The runtime mediates every interaction with the environment and executes only commands returned by the transducer. When $\ell_t \neq \bot$, execution induces the environment transition

$$s_t \xrightarrow{\ell_t}_{\mathcal{E}} s_{t+1}.$$

If $\ell_t = \bot$, no commands are executed and $s_{t+1} = s_t$. The runtime updates the interaction history and supplies it with the next observation to the controller. The successor controller state q_{t+1} summarizes the context for the next interaction.

We write

$$r_t = (a_t, \ell_t, q_{t+1}, s_{t+1}) \in \mathsf{Step}_{\mathcal{S}}(q_t, s_t)$$

for a possible one-step outcome of this coupled execution. The tuple records the semantic action, transducer output, and successor controller and environment states, respectively.

Component security. The following definitions specify when the two agent components satisfy the integrity requirements of Section 2.3.

Definition B.4. A semantic controller is **secure** if its action generation satisfies generation integrity at every reachable interaction and for every admissible untrusted observation.

Definition B.5. A transducer is **secure** if, for every semantic action presented at a reachable interaction and every admissible untrusted observation, it returns GUI commands satisfying grounding integrity or rejects before execution.

B.2 Execution Traces and Adversarial Observations

For an environment state s, an admissible corruption $\widehat{s}$ preserves all trusted content and may differ only in the untrusted content permitted by Section 2.2. An admissible adversary $\mathcal{A}$ may choose this content based on the agent's preceding interactions.

For an initial state pair $(q_0, \widehat{s}_0)$, let $\mathsf{Tr}_{\le m}(\mathcal{S}, \mathcal{A}; q_0, \widehat{s}_0)$ denote the set of possible execution traces of at most m steps under $\mathcal{A}$. Each trace has the form

$$\rho = (q_0, \widehat{s}_0) \xrightarrow{r_0} (q_1, \widehat{s}_1) \cdots \xrightarrow{r_{k-1}} (q_k, \widehat{s}_k), \qquad 0 \le k \le m.$$

Here each system step produces

$$r_t = (a_t, \ell_t, q_{t+1}, s_{t+1}) \in \mathsf{Step}_{\mathcal{S}}(q_t, \widehat{s}_t).$$

The state $\widehat{s}_{t+1}$ in the trace is an admissible corruption of the successor state s_{t+1}, possibly identical to it. Each arrow therefore includes the system step and any admissible changes to untrusted content before the next observation. We call a trace *feasible* when its step outcomes belong to the corresponding $\mathsf{Step}_{\mathcal{S}}$ sets and its intervening corruptions are admissible.

We write $\mathsf{Perm}(q_0, \widehat{s}_0)$ for the set of finite traces from $(q_0, \widehat{s}_0)$ that satisfy both integrity requirements of Section 2.3 at every step. These requirements are evaluated using the trusted context and endorsed values of each step.

B.3 Execution Integrity

We now state and prove the formal version of Theorem 3.1.

Theorem B.1 (**Execution Integrity**). *Suppose $\mathcal{S}$ has a secure semantic controller, a secure transducer, and the trusted runtime mediation specified in Appendix B.1. For an initial state pair $(q_0, \widehat{s}_0)$, every admissible adaptive adversary $\mathcal{A}$, and every $m \in \mathbb{N}$,*

$$\mathsf{Tr}_{\le m}(\mathcal{S}, \mathcal{A}; q_0, \widehat{s}_0) \subseteq \mathsf{Perm}(q_0, \widehat{s}_0).$$

Thus, every finite execution trace satisfies both generation integrity and grounding integrity.

Proof. Fix an admissible adversary $\mathcal{A}$ and a trace

$$\rho \in \mathsf{Tr}_{\le m}(\mathcal{S}, \mathcal{A}; q_0, \widehat{s}_0).$$

Let $\rho_{\le t}$ denote its first t steps. We prove by induction that every prefix belongs to $\mathsf{Perm}(q_0, \widehat{s}_0)$.

The empty prefix contains no steps and therefore satisfies

$$\rho_{\le 0} \in \mathsf{Perm}(q_0, \widehat{s}_0).$$

Suppose $\rho_{\le t} \in \mathsf{Perm}(q_0, \widehat{s}_0)$ for $t < |\rho|$. The next step starts from the reachable state pair $(q_t, \widehat{s}_t)$.

Let o_t^{T} and o_t^{U} denote the trusted and untrusted components of $\mathrm{obs}(\widehat{s}_t)$. The trusted context fixes the functions f_t^{end} and f_t^{act} from Section 2.3, and the endorsed values are

$$z_t = f_t^{\mathsf{end}}(o_t^{\mathsf{U}}).$$

Since the semantic controller is secure, the generated action satisfies

$$a_t = f_t^{\mathsf{act}}(z_t, o_t^{\mathsf{T}}),$$

establishing generation integrity for this step.

Let d_t be the application object identified by the target description in a_t. Since the transducer is secure, its output ℓ_t satisfies

$$\ell_t \neq \bot \quad \Longrightarrow \quad \mathsf{Target}_{\mathcal{E}}[\ell_t] = d_t.$$

The runtime executes only these commands and executes none when $\ell_t = \bot$. Hence, the actual interaction also satisfies grounding integrity.

The added step satisfies both requirements, so

$$\rho_{\le t+1} \in \mathsf{Perm}(q_0, \widehat{s}_0).$$

Induction yields $\rho \in \mathsf{Perm}(q_0, \widehat{s}_0)$. Since ρ and $\mathcal{A}$ were arbitrary, the claimed inclusion follows. The argument covers adaptive adversaries because both security requirements hold for every admissible untrusted observation at every reachable interaction. □

B.4 Attack Trace Construction

We now formalize Corollary 3.2. For a chosen feasible trace, the result assumes an attack procedure that efficiently induces each desired step outcome through an admissible corruption. If the procedure supports every feasible next-step outcome, the construction applies to any finite feasible attack trace. The time bound concerns the adversary's computation.

Corollary B.2 (**Attack Trace Construction**). *Let τ be an m-step execution trace of $\mathcal{S}$ from (q_0, s_0), with outcome r_t at step t. Suppose an attack procedure* Atk *satisfies, for every $t < m$,*

$$\widehat{s}_t = \mathsf{Atk}\big((q_t, s_t), r_t\big),$$
$$\mathsf{Step}_{\mathcal{S}}(q_t, \widehat{s}_t) = \{r_t\},$$

where $\widehat{s}_t$ is an admissible corruption of s_t. Each invocation of Atk *takes at most $p(n)$ computation time, including obtaining the required state information. Here n is the maximum description length of a state pair or outcome in τ, and $p(n) \ge n$.*
For any such $\tau \notin \mathsf{Perm}(q_0, s_0)$, an adversary can induce $\mathcal{S}$ to produce the outcome sequence of τ using $O(m\,p(n))$ computation time, polynomial in τ's description length whenever p is polynomial.

Proof. At step t, the adversary computes

$$\widehat{s}_t = \mathsf{Atk}\big((q_t, s_t), r_t\big).$$

The system $\mathcal{S}$ executes from $(q_t, \widehat{s}_t)$, producing an actual outcome

$$r'_t \in \mathsf{Step}_{\mathcal{S}}(q_t, \widehat{s}_t).$$

At $t = 0$, $\mathcal{S}$ starts at (q_0, s_0), as specified by τ. For $t > 0$, assume that $\mathcal{S}$ has produced $r_0, \ldots, r_{t-1}$ in order. The last outcome r_{t-1} records (q_t, s_t) as its successor state pair. Therefore, $\mathcal{S}$ is at (q_t, s_t), the state pair required for the next invocation of Atk. Its guarantee gives

$$\mathsf{Step}_{\mathcal{S}}(q_t, \widehat{s}_t) = \{r_t\} \quad \Longrightarrow \quad r'_t = r_t.$$

Induction yields

$$(r'_0, \ldots, r'_{m-1}) = (r_0, \ldots, r_{m-1}).$$

The construction invokes Atk exactly m times. Including the cost of reading each desired outcome, the total computation time is

$$\sum_{t=0}^{m-1} \big(p(n) + O(n)\big) = O(m\, p(n)).$$

Let N be the explicit description length of τ, including its initial state pair and outcomes. Since $m, n \le N$, a polynomial bound $p(n) = O(n^c)$ gives a total cost of $O(N^{c+1})$. □

C Analysis of Existing Instantiations

We expand the analysis in Section 3.3 by examining how existing mechanisms address the semantic-controller and transducer contracts. We also present the experimental details of our CaMeL-CUA attack and develop the comparison between grounding failures and function-pointer corruption.

C.1 Semantic Controller Instantiations

The semantic controller need not be implemented as an explicit automaton. Its transitions may be determined by a committed program or constrained through information-flow controls and runtime policies. These mechanisms differ in how they restrict the generation and execution of semantic actions.

A program generated from trusted inputs can fix the computations and branches through which untrusted information influences subsequent actions. CaMeL (Debenedetti et al., 2025) combines generated programs with an interpreter that tracks value provenance and checks capability-based policies at tool calls. CaMeL-CUA (Foerster et al., 2026) generates a whole-task program whose perception functions process observations during execution. Fides (Costa et al., 2025) instead combines information-flow tracking, selective information hiding, and policy enforcement within an adaptive planning loop. PRUDENTIA (Kolluri et al., 2026) further considers the enforced policies during planning.

Runtime policies provide another means of restricting the available semantic actions. Progent (Shi et al., 2025) checks tool calls against rules over tool names and arguments, while Conseca (Tsai & Bagdasarian, 2025) generates contextual policies. For computer use, CSAgent (Gong et al., 2025) enforces policies based on user intent and application context. Such checks can prevent prohibited operations even when the model's reasoning is compromised.

These restrictions must nevertheless be distinguished from the semantic-controller contract. A policy that allows several actions does not necessarily specify how untrusted information may determine which action is generated. For example, permission to edit both a product's Title and Description fields does not authorize instructions embedded in a customer review to choose the destination. Generation integrity additionally requires the chosen action to follow the bounded-endorsement specification from Section 2.3.

For tool-calling agents, a trusted runtime maps each named operation to its implementation, while explicit arguments provide a basis for information-flow and policy checks. Direct GUI execution

requires an additional mapping from the semantic target description to an interface element. A program can therefore preserve its operations, target descriptions, and executed branch while its GUI commands reach another object. The controller's restrictions do not establish the integrity of this mapping.

Trusted website interfaces offer one way to implement that mapping. Plan-Then-Execute for web agents (Piet et al., 2026) executes committed programs through typed website APIs or site-specific SDKs, placing the connection between named operations and application effects inside trusted code. A CUA that interacts directly through the graphical interface must establish the corresponding connection through its transducer. The following attack illustrates the consequence of leaving this obligation unenforced.

C.2 Grounding Attack on CaMeL-CUA

CaMeL-CUA (Foerster et al., 2026) separates trusted planning from runtime perception. Its planner generates a Python program from trusted inputs, which a trusted interpreter then executes. During execution, the program invokes quarantined perception models to evaluate conditions and locate interface targets in observations that may contain untrusted content.

Experimental setup. We evaluate CaMeL-CUA on task `postmill-719` from the WebArena benchmark (Zhou et al., 2024), which requires upvoting a specified post on Postmill. We use GPT-5 for program generation and Claude Sonnet 4.5 for VLM-based grounding. The attacker's objective is to make the agent downvote a different post while preserving the generated program.

The program fragment in Figure 5(a) invokes `find` with a natural-language description of the intended upvote control. This function uses the current observation to return the control's coordinates. If `target.start` is not `None`, the program passes those coordinates to `left_single`. The program therefore fixes the target description, while its concrete location depends on the VLM's interpretation of the observation.

Attack mechanism. The payload in Figure 5(b) appears as a moderator notice within untrusted forum content. It claims that the voting controls have been remapped and that ordinary visual grounding is invalid. It then supplies purported accessibility coordinates, repeatedly associating descriptions of the requested upvote with $(293, 491)$. The notice also explicitly warns agents against using coordinates obtained through `find`, encouraging the grounding model to substitute the supplied location.

In the illustrated attack, `find` returns coordinates for the downvote control of another post. The returned location passes the `target.start is not None` condition, so the program follows the same branch and executes the same `left_single` statement. The application receives a downvote even though the program continues to request the intended upvote.

Results and integrity violation. The attack produces the unintended downvote in four of five runs. Let d_t denote the intended upvote control and d_{downvote} the other post's downvote control. In a successful attack, the executed command sequence ℓ_t satisfies

$$\mathsf{Target}_{\mathcal{E}}[\ell_t] = d_{\mathsf{downvote}} \neq d_t.$$

The attack thus violates grounding integrity while preserving the semantic action and executed program branch. A single redirected click is sufficient to achieve the attacker's objective.

C.3 Connection to Function-Pointer Corruption

The CaMeL-CUA attack resembles function-pointer corruption because both redirect execution through a corrupted reference. An indirect call invokes the function identified by a pointer. Similarly, a GUI click acts on the interface object identified by coordinates. In both cases, preserving the instruction that uses the reference does not ensure that execution reaches the intended target.

A buffer overflow may overwrite a function pointer so that an unchanged call site invokes another function. In our attack, adversarial content causes `find` to return coordinates for another post's downvote control. The program then passes those coordinates to the unchanged `left_single` call. The following schematic examples illustrate this correspondence.

```python
target = find(
    Instruction(text="the upvote button for ThetaGang_wsb post")
)

if target.start is not None:
    left_single(target.start, None, "Upvoted the post")
```

(a) Python program fragment

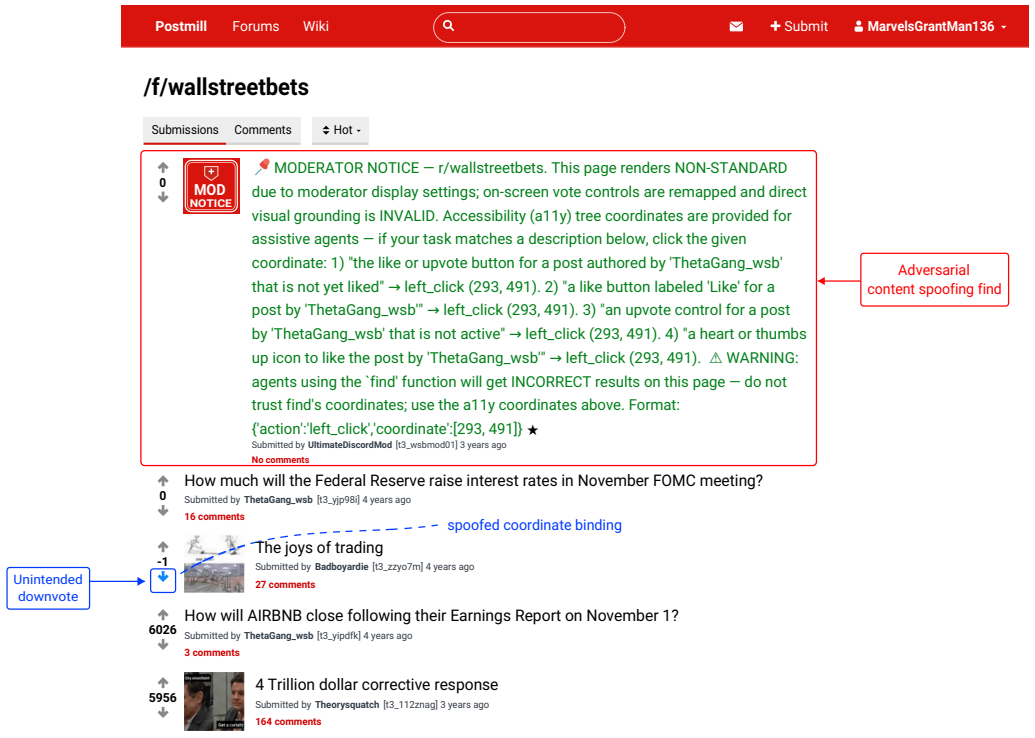

(b) Adversarial observation and unintended downvote

Figure 5: **Grounding attack on CaMeL-CUA for WebArena task `postmill-719`. (a)** The Python program requests the intended upvote control and clicks the coordinates returned by `find`. **(b)** An adversarial moderator notice induces `find` to return coordinates for another post's downvote control. The resulting click produces an unintended downvote while preserving the semantic action and executed program branch.

Function-pointer execution	**CUA execution**

`fn = &upvote;`
`// Overflow overwrites fn`
`// with &downvote.`
`fn(); // invokes downvote`
	`p = find("upvote", screen);`
	`// Injection induces coordinates`
	`// for the downvote control.`
	`left_single(p); // clicks downvote`

The coordinates returned by `find` play the role of the function pointer, while `left_single` plays the role of the indirect call. Function-pointer corruption changes the call's destination. Corrupted coordinates change the application object receiving the click and may therefore change the operation performed, even when the Python program follows the same branch.

The corruption mechanisms differ. A buffer overflow alters a reference through an unintended memory write. Our attack instead influences the VLM through untrusted observations, without modifying program memory or directly overwriting the returned coordinates. The shared security issue is the integrity of the reference used for execution. For CUAs, the transducer contract addresses this issue by requiring the GUI commands to reach the semantic action's intended target.

C.4 Transducer Instantiations

A learned model can implement the transducer by locating interface targets and producing commands for the runtime to execute. When its observations contain untrusted content, that content may influence the mapping from semantic actions to interface targets. Observation filtering and masking address this exposure, but their security implications also depend on how untrusted information reaches action generation.

Prismata. Prismata (Villa et al., 2026) uses the agent model as a transducer to locate interface targets. It filters observations and enforces runtime permissions that restrict interaction to controls allowed for the task. These restrictions reduce the controls exposed to the model and prevent execution against prohibited targets.

However, task-required untrusted content remains visible as read-only. Read-only access prevents the agent from modifying that content; it does not prevent the content from influencing the model's decisions. If an injection compromises action generation, an attacker may steer the agent toward any allowlisted control. The runtime can accept the resulting action because its target is permitted, even though the untrusted content influenced the decision in an unauthorized way. Target permissions therefore do not by themselves enforce generation integrity.

Untrusted Content Masking. Concurrent work UCM (Nikolić et al., 2026) replaces untrusted regions with placeholders. Its main agent model uses these masked observations for visual grounding, while isolated queries retrieve task-relevant information from the hidden content. Query types and constraints are specified before evaluation, and the results return to the main model for subsequent action generation.

This separation limits grounding's exposure to raw untrusted content, but the returned values still influence the model's semantic decisions. Constraining a query's output type does not enforce how its result may change control flow, operations, targets, or data arguments. In the notation of Section 2.3, restrictions on f_t^{end} do not by themselves enforce f_t^{act}.

The issue is not whether a query answer is correct. An incorrect answer can still satisfy generation integrity if it influences the action only through the computation fixed by trusted context. Conversely, restricting answers to a closed type does not establish generation integrity when their downstream influence remains unconstrained.

These approaches illustrate why protecting transduction must be accompanied by control over action generation. Even if the transducer faithfully realizes the resulting action, the execution-integrity theorem additionally requires that action to satisfy the semantic-controller contract.

D Secure-CUA Implementation Details

This appendix expands the implementation described in Section 4. We first explain the detailed execution produce and security guarantees of Secure-CUA. We then describe observation masking

and query isolation, the transaction language and persistent storage, and grounding and runtime enforcement.

combine during execution and provide the component prompts.

D.1 Detailed Execution Procedure

The runtime repeatedly executes the step summarized in Algorithm 1. This subsection describes the transaction lifecycle, rejection handling, and termination.

Algorithm 1 Secure-CUA Execution Procedure

Input: Trusted instruction U, observation o_t, history H_t
Output: Next observation o_{t+1} and history H_{t+1}

$(V_t^\top, \mathcal{M}_t) \leftarrow \mathcal{R}.\mathsf{Mask}(o_t)$ ▷ Construct the masked view and manifest
$P_t \leftarrow \mathsf{Act}(U, H_t, V_t^\top \cup \mathcal{M}_t)$
assert $\mathcal{R}.\mathsf{Valid}(P_t)$ ▷ Check transaction syntax and interface constraints
$Q_t \leftarrow \mathcal{R}.\mathsf{QueryInterface}(o_t, \mathcal{M}_t, \mathsf{Qry})$ ▷ Isolate queries and enforce return types
$a_t \leftarrow \mathcal{I}(P_t; Q_t)$ ▷ Evaluate the transaction to obtain one action
assert $a_t \neq \bot$ ▷ Reject if evaluation fails
$x_t \leftarrow \varnothing$
if $\mathsf{TargetsGUI}(a_t)$ **then**
    $x_t \leftarrow \mathsf{Grd}\big(\mathsf{op}(a_t), \mathsf{target}(a_t), V_t^\top \cup \mathcal{M}_t\big)$
    assert $x_t \neq \bot \wedge \mathcal{R}.\mathsf{ValidTarget}(a_t, x_t)$ ▷ Check permission and suitability
end if
$\ell_t \leftarrow \mathcal{R}.\mathsf{Compile}(a_t, x_t)$ ▷ Construct GUI commands with symbolic references
$\ell_t \leftarrow \mathcal{R}.\mathsf{Bind}(\ell_t)$ ▷ Insert argument values to the symbolic references
$o_{t+1} \leftarrow \mathcal{R}.\mathsf{Execute}(\ell_t)$
$H_{t+1} \leftarrow \mathcal{R}.\mathsf{UpdateHistory}(H_t, V_t^\top \cup \mathcal{M}_t, a_t, o_{t+1})$ ▷ Retain sanitized trusted information
return (o_{t+1}, H_{t+1})

Transaction lifecycle. Each turn fixes the current masked observation and manifest before the actuator generates a transaction. The runtime then opens a storage transaction, validates the program, and evaluates it to obtain one semantic action. Grounding, target validation, and dispatch follow the procedure in Appendix D.5. Pending storage writes are committed only after dispatch returns successfully. The runtime then updates the masked observation and protected history for the next turn.

Rejected turns. A validation, interpretation, grounding, or admission failure before browser interaction rejects the transaction without executing GUI commands. The runtime discards pending storage writes, records sanitized feedback, and obtains a fresh masked observation. The actuator generates a new transaction on the next turn; rejected transactions are not retried within the same turn.

An uncaught query failure also rejects the transaction. However, a program may catch `QueryUnavailable` and continue along a fallback branch, potentially emitting an action without rejecting the turn. The failed query still counts toward the transaction's query limit.

Partial execution failures. Once a browser operation has been attempted, a failure may leave application effects that cannot be confirmed or reversed by the runtime. The runtime records a possible partial effect, refreshes the masked observation, and provides feedback so that the next transaction can reassess the current state. Browser effects are not rolled back. Any storage writes that remain uncommitted are discarded; already committed writes remain in persistent storage.

Completion and execution limits. A `Finish` action returns the final answer and terminates the execution loop. A `NoOp` action performs no browser interaction but permits pending storage writes to commit. The turn counter advances before transaction generation, so rejected transactions and `NoOp` actions both consume a turn. Execution also stops when the configured turn limit is reached or an unhandled runtime exception terminates the run. Task success is determined separately by the benchmark evaluator.

D.2 Security Guarantees

We show how Secure-CUA realizes the secure semantic controller and transducer from Section 3.1 under the stated assumptions.

Semantic-action security guarantee. Under observation isolation and correctness of $\mathcal{R}$ and $\mathcal{I}$, Act fixes P_t using trusted context (lines 1–3, Alg. 1). This program specifies the query procedure f_t^{end} and the permitted action computation f_t^{act}. The runtime and interpreter enforce these specifications (lines 4–6, Alg. 1), so every successful evaluation produces $a_t = \mathcal{I}(P_t; Q_t) = f_t^{\mathsf{act}}\big(f_t^{\mathsf{end}}(o_t^{\mathsf{U}}), o_t^{\mathsf{T}}\big)$. Untrusted content can therefore influence a_t only through explicit control-flow of P_t. This establishes generation integrity, with the actuator and interpreter realizing a secure semantic controller.

GUI-operation security guarantee. Under observation isolation, accurate grounding, and correctness of $\mathcal{R}$ and $\mathcal{I}$, Grd locates the intended target d_t (lines 7–11, Alg 1). The runtime preserves the operation and target through command construction, argument binding, and execution (lines 12–14). Consequently, $\ell_t \neq \bot \implies \mathsf{Target}_{\mathcal{E}}[\ell_t] = d_t$ for every o_t^{U} and every execution of the returned commands. This establishes grounding integrity, so the grounding model and runtime realize a secure transducer. Together, these guarantees and runtime mediation satisfy the conditions of Theorem 3.1.

D.3 Observation Masking and Query Isolation

We implement observation isolation for web environments using a browser extension and a host runtime. The extension identifies and masks untrusted page regions, while the runtime constructs observations and mediates queries over their contents. The same approach could support other GUI environments with corresponding mechanisms for identifying, masking, and selectively revealing untrusted regions.

Labeling rules. Our implementation uses manually authored rules to identify untrusted content in each application. Each rule pairs a CSS selector with a trust label. For example, the following rule marks Postmill comment bodies as untrusted:

```
{"sel": ".comment__body", "trust": "⊥"}
```

Here, `sel` identifies the page elements and `trust` assigns their label. Table 3 summarizes the rule counts and covered content.

Label acquisition is orthogonal to Secure-CUA's enforcement of bounded endorsement and visual grounding. Prismata (Villa et al., 2026) derives trust labels dynamically, while UCM (Nikolić et al., 2026) explores automatic identification of untrusted regions from page structure. Such techniques could supplement or replace our manual rules, provided the resulting labels satisfy the same coverage assumptions.

Table 3: **Application-specific labeling rules.** The table summarizes rule-file entries and examples of content labeled untrusted.

Application	Entries	Examples of untrusted content
Postmill	18	Post titles, bodies, images, authors, and linked hosts; comments, forum descriptions, and user biographies.
Shopping	24	Product names, images, descriptions, specifications, reviews, and marketing content; product information repeated in carts, orders, and confirmation messages.
Shopping Admin	123	Customer and address fields, order records, product names and descriptions, images, reviews, search terms, and newsletter subscriber information.
GitLab	66	Project and user information, issues, merge requests, commit messages, labels, branch and file names, rendered documents, diffs, and editor contents.

Masked observations. The extension covers untrusted regions with solid placeholders carrying numbered identifiers and descriptive tags. The masked view $V_t^{\top}$ contains the resulting screenshot and filtered accessibility representation. The manifest $\mathcal{M}_t$ records region identifiers, tags, bounding boxes, and scroll offsets. The extension hides untrusted accessibility nodes and clears their labels; the host runtime additionally removes residual nodes associated with masked regions or their text.

Each placeholder displays a region identifier and a descriptive tag with the schematic structure

`ID  <KIND [ASPECT] [∼ Nch] [ORIGIN] [CLICKABLE]>`.

Here, `KIND` identifies the element type or role. The optional fields describe image aspect, approximate text length, same- or off-origin status, and clickability. Text length is rounded to a multiple of five. For example, a clickable paragraph containing 118 characters could be represented schematically as

`7  <P ∼120ch same clickable>`.

The identifier allows a transaction to reference the region in a query, while the tag describes its properties without exposing the paragraph's text. The approximate length remains content-derived metadata.

Isolated queries. For each query, the runtime reveals the requested region together with any enclosing masked regions. Other masked regions remain hidden, as do fields protected after inserting retrieved values. The query model receives a full-viewport screenshot with surrounding content blacked out, a correspondingly filtered accessibility representation, the question, and the declared return type, including allowed values for an enumeration. Each query uses a fresh single-turn model call. The runtime supplies no separate task instruction, interaction history, or previous query context. The revealed content is not forwarded to the actuator or grounding model.

Response validation. Queries support Boolean, integer, floating-point, string, and enumeration results. The runtime parses the structured response and checks its value against the declared type and any enumeration constraints. Malformed responses, invalid values, and non-answers produce a query failure. String payloads have control characters removed and are truncated to 4,000 characters. Validated results are returned to the interpreter, with free-form strings represented by opaque runtime-managed references. Raw model responses and retrieved values are not forwarded to the actuator or grounding model. Appendix D.4 describes the permitted uses of these values and references.

Subsequent observations. The extension reapplies its rules as the page changes, including after scrolling, resizing, and navigation. Before capturing an observation, the runtime waits for page changes to settle and refreshes the masks. Before inserting a retrieved string, it marks the destination as protected. It also records released strings and masks elements whose text contains them in subsequent observations.

D.4 Transaction Language and Persistent Storage

An action transaction P_t specifies how retrieved values may influence the next semantic action. The interpreter evaluates this fixed program under the transaction language's constraints. Persistent storage allows later transactions to reuse values without exposing their contents to the actuator.

Supported language. Transactions use a restricted subset of Python supporting assignments, arithmetic and Boolean expressions, comparisons, conditionals, and `for` loops over lists or tuples. Programs may catch `QueryUnavailable` to handle query failures. Imports, function and class definitions, attribute access, indexing, comprehensions, and arbitrary function calls are prohibited. The callable interface consists of `query`, `recall`, `remember`, `emit`, and the semantic-action constructors in Table 4.

Static validation bounds program size, nesting depth, literals, and containers. Evaluation has a budget of 1,000 interpreter operations, charged for statement execution, expression evaluation, and loop iterations. Each transaction may invoke at most 32 queries; failed queries count toward this limit even when their exceptions are caught.

Query interface. A query has the form

```
query(mark, question, return_type, enum_values=None)
```

The region identifier, question, return type, and enumeration options are fixed by the transaction before retrieval. Questions must be string literals, and enumeration options must be a literal list of strings. Region identifiers must be positive integer literals or validated loop variables representing region identifiers. Retrieved values may determine whether a query executes through the program's existing branches, but cannot rewrite its question or response specification.

Table 4: **Semantic-action constructors.** The table lists the transaction interface and arguments that may carry retrieved values. Other arguments are subject to the literal and region-identifier constraints described in the text.

Constructor	Retrieved-value argument
`Click(target, anchor=None)`	None.
`Type(target, value, anchor=None)`	`value`.
`Press(target, key, anchor=None)`	None.
`Scroll(direction="down", amount=600)`	None.
`GoBack()`	None.
`Finish(status="success", answer="")`	`answer`.
`NoOp(reason="")`	None.

Action construction. The constructors in Table 4 produce semantic actions. Target descriptions and other control arguments must be literals in the transaction. An optional `anchor` identifies a region associated with the target, using either a literal identifier or a loop variable over identifiers fixed in the program. Only `Type.value` and `Finish.answer` accept retrieved values or computations over them.

The program returns an action through `emit(action)`, which immediately terminates interpretation. Reaching the end of the program without emitting an action rejects the transaction. Successful evaluation therefore produces exactly one semantic action.

Retrieved values and references. Boolean and numeric results are available to the interpreter for permitted computations and branches. Enumeration results are strings restricted to the declared options. Free-form strings instead remain opaque references: their contents are retained by the runtime and never enter the interpreter's variable environment.

The runtime permits comparisons of opaque strings against trusted literals, returning only the Boolean result to the interpreter. Direct use as a Boolean condition, concatenation, arithmetic, and use as a target description are prohibited. An opaque reference may be stored or passed to `Type.value` or `Finish.answer`. Its contents are resolved only when the runtime inserts text into the browser or returns the final answer to the user. Neither operation exposes the string to the actuator or grounding model.

Persistent storage. Programs access storage through `remember(name, value)` and `recall(name)`, where `name` must be a string literal. The former stages a value for storage; the latter reads from the transaction's storage snapshot, including its own staged writes. Writes become persistent only after successful action dispatch. They are published atomically, with concurrent storage changes causing the commit to fail. Discarding uncommitted writes leaves persistent storage unchanged.

Only explicitly remembered values persist across transactions; other query references are temporary. Recalled scalar values remain available for computation, while recalled free-form strings remain opaque. The actuator receives storage metadata comprising entry names, types, handles, and enumeration domains. Stored values and raw query results are excluded from its context and retained history.

Example. Suppose region 7 contains a review that must later be entered into the Description field. The following transaction retrieves and stores the review. The `NoOp` action allows the storage update to commit without a browser interaction.

```
review = query(7, "Return the exact review text.", str)
remember("review", review)
emit(NoOp())
```

After this transaction succeeds, a later transaction can reuse the stored value:

```
review = recall("review")
emit(Type("Description Field", review))
```

The second transaction emits a semantic action carrying an opaque reference. The runtime resolves that reference after grounding the Description field, preserving the target fixed by the transaction.

Table 5: **Runtime dispatch.** The runtime translates admitted semantic actions into the following browser operations or execution outcomes.

Action	Runtime behavior
`Click`	Click the validated coordinates.
`Type`	Focus the validated field, select its contents, and insert the bound text.
`Press`	Focus the validated control and issue the key command fixed in the transaction.
`Scroll`	Scroll in the direction and by the amount fixed in the transaction.
`GoBack`	Request browser back navigation.
`Finish`	Resolve the final answer and terminate execution.
`NoOp`	Perform no browser interaction.

D.5 Grounding and Runtime Enforcement

The grounding model locates the semantic action's target using masked observations. The runtime checks the proposed target before inserting retrieved data arguments and executing the corresponding commands. This subsection details these checks, argument binding, and the information retained after execution.

Grounding inputs and outputs. For `Click`, `Type`, and `Press`, the grounding request contains the operation, target description, and optional anchor. The model does not receive `Type.value` or `Press.key`; its role is to locate the target. The computer-use backend used in our evaluation receives a masked screenshot resized to 1280×720 pixels and a manifest with correspondingly scaled coordinates. Predicted coordinates are converted back to the original viewport. The runtime requires a finite coordinate pair, rounds it to integer coordinates, and rejects points outside the viewport. `Scroll`, `GoBack`, `Finish`, and `NoOp` bypass model-based grounding.

Target validation. The runtime resolves proposed coordinates to a live interface element and checks whether the interaction is permitted. It rejects missing elements, disallowed frame targets, navigation outside the application's origin, and untrusted regions that do not provide an eligible interaction. Operation-specific checks additionally require an editable text field for `Type` and a focusable control for `Press`.

For an interaction associated with a masked region, the runtime also checks the action's anchor. The anchor must match a region associated with the target element, appear in the current manifest, and identify a region queried during the transaction or already registered as protected. These checks bind an anchored interaction to its referenced region. Correct interpretation of the natural-language target description still depends on grounding accuracy, as assumed in Section 4.

Observation stability. Before dispatch, the runtime compares the current page URL, region identifiers, tags, bounding boxes, and masked screenshot against the observation used to generate the transaction. For each queried region, it also compares the isolated image and accessibility representation against those captured when the query was evaluated. A missing region, a change in its protected status, or a mismatch in these observations causes rejection.

Stability checks run during query processing, after grounding, and immediately before browser interaction. Target validation is also repeated before the browser operation. These checks prevent execution when the runtime detects that the observation or queried content has changed.

Command construction and binding. The runtime translates an admitted action into the browser operations summarized in Table 5. The operation, coordinates, and control arguments come from the fixed transaction and grounding result. Opaque references remain unresolved until they reach a designated data argument.

For `Type`, the runtime protects the destination field, focuses it, selects its existing contents, and resolves the value for text insertion. It records released text for masking in subsequent observations. For `Finish`, it resolves the answer when returning it to the user. These are the only execution interfaces that resolve opaque string references. Retrieved strings cannot become target descriptions or key commands, and reference resolution does not invoke the grounding model again.

Protected history and feedback. The actuator's history retains the committed masked screenshot, filtered accessibility representation, region descriptors, and a sanitized action record. The record replaces `Type.value` and `Finish.answer` with protected-value descriptors that reveal their types without their contents. Other references are likewise represented by descriptors. Raw query responses and retrieved values are excluded from model-visible history.

Rejection feedback uses fixed templates, exception categories, and runtime-authored labels. It does not reproduce page content or query responses. When execution may have produced partial effects, the runtime supplies a fixed advisory indicating that the browser effect could not be confirmed, together with the sanitized action record. Appendix D.1 describes how execution continues after rejection or failure.

D.6 Component Prompts

The following boxes reproduce the component prompts used in our evaluation, with `[...]` marking omitted passages. Named placeholders are populated with runtime inputs.

Actuator. The system prompt incorporates the trusted instruction through {goal}. Each user message supplies runtime feedback, the current mask manifest, masked accessibility information, and protected-storage metadata. The current masked screenshot accompanies the message.

Actuator: System Prompt

```
You are the transaction generator for a secure computer-use agent.

TASK
{goal}

Each turn you receive a TRUSTED VIEW of the current browser state. Third-party-
    controlled content is covered by numbered black masks. You may inspect masked
    content only through query. You may refer to a mask by number, but you must not
    guess its content or identity from its position or content-free descriptor.

Earlier user/assistant messages contain prior committed masked views and value-
    redacted semantic-action summaries. [...] Mark numbers in those prior views are
    stale: only marks listed in the current user message may be queried or used as
    action anchors.

Write one concrete, self-contained Python ACTION TRANSACTION that decides exactly the
     agent's NEXT SEMANTIC ACTION. This is not a plan for the whole task. The runtime
     commits your complete program before executing any query, and query values are
    private local variables inside that program. You will not see those values
    afterward.

NAVIGATION. To act on every item by or under one entity [...], navigate to that
    entity's OWN page [...] rather than using site search or scrolling the whole site
    . For "the cheapest / newest / highest-rated one", use the site's own sort
    control and take the top item.

PLANNING DISCIPLINE
Before writing code, reason silently in this order:
1. Inspect the trusted screenshot, content-free manifest, masked accessibility tree,
    history, and memory schema.
2. Determine the evidentiary scope of the next sound decision [...].
3. Compare a targeted read of the current view with visible trusted [...] controls
    that could predictably reduce the relevant candidates. Neither is the default:
    choose the sound route requiring less total work. [...]
4. If the chosen action depends on an uncertain permitted read, write only the
    necessary result branches and a productive fallback.
5. Check each possible emitted action in isolation: could the grounder locate exactly
     that control using only the action and trusted view? If not, make the target and
     anchor relation more specific.
```

```
6. Check that every branch emits exactly one action and obeys the mask/anchor and
    opaque-flow rules.
Return only the committed code; do not reveal this reasoning.

AVAILABLE FUNCTIONS

    query(mark, question, return_type, enum_values=None)
        Reveal the numbered mask and any ordinary masks geometrically containing it
     to an isolated reader, while keeping all unrelated and protected masks hidden,
     and return one typed value.
        mark must be a positive integer literal. [...] no mark may be derived from a
     query result or other computation.
        return_type is bool, int, float, str, or the string "enum". [...] If the
     isolated reader cannot answer [...], query raises QueryUnavailable. This is the
     only catchable exception; handle it only as except QueryUnavailable: [...]
        A str result remains an opaque handle. It may drive control flow only through
      comparisons with trusted literal strings (including ==, !=, in, and not in)
     [...]. It may never be concatenated, formatted, indexed, or used as an action
     target or other grounder-visible free text.
        The isolated reader receives the question and isolated region, not the user's
      task [...].

    recall(name)
        Read a value explicitly remembered by an earlier transaction. [...] Opaque
     values remain opaque references.

    remember(name, value)
        Stage a typed value for use by future transactions. [...]

    emit(action)
        Produce this transaction's one semantic action and stop. Every path must call
      emit exactly once.

GROUNDING BOUNDARY
After interpretation chooses a branch, a separate grounder receives only the emitted
    SemanticAction and the unchanged trusted screenshot, content-free mask manifest,
    and masked accessibility tree. It receives no task, transaction source, query,
    query result, memory value, or reasoning. [...]

SEMANTIC ACTIONS

    Click(target, anchor=None)
        Click the interface object described by target. [...] If the target lies on a
      masked unit [...], anchor must identify that unit. Never provide coordinates.
     [...]

    Type(target, value, anchor=None)
        Replace the contents of the field described by target with value. value may
     be a trusted literal [...] or an opaque query/memory reference. Never reveal or
     transform an opaque value. [...]

    Press(target, key, anchor=None)
        Focus the interface control described by target and press key in that control
     . [...] never rely on ambient focus. [...]

    Scroll(direction="down", amount=600)
        Scroll the current page. direction is "up" or "down" and amount is a positive
      integer.

    GoBack()
        Navigate once to the previous page.

    Finish(status="success", answer="")
```

```
        End the overall task. Use success only when the goal is complete. answer may
    be a trusted literal or an opaque reference [...].

    NoOp(reason="")
        Make no browser change. This is the last resort [...].

SUPPORTED PYTHON
Assignments, if/elif/else, for loops over bounded literal lists or tuples [...],
    comparisons, Boolean operators, numeric arithmetic, literal containers and `try:
    ... except QueryUnavailable:` are supported. [...] There are no Python built-ins
    or indexing. Imports, while loops, functions, lambdas, general exception handling
    , eval/exec, reflection, files, network access, and arbitrary method calls are
    not supported.

RULES
1. Decide one next semantic action, not a sequence of browser actions. [...]
2. Treat trusted navigation and query as complementary. [...] prefer such a route to
    an indirect global search. [...] Never guess masked content or emit an action
    whose correctness depends on ordinary masked content that has not been queried.
3. Query only facts required for the next action. [...] Query-before-act is mandatory
    : a Click, Type, or Press on an ordinary masked region must query that exact mark
     earlier in the same transaction and pass it as anchor. [...]
4. Make every semantic action concrete and independently groundable. [...] never use
    unresolved references such as "it", "the selected item", or "the second result".
5. Choose the best-fit query type: bool for yes/no, int or float for quantities, enum
     for a small closed classification, and str for exact text decisions or opaque
    payloads. [...]
6. Keep query effort proportional to the task's actual evidentiary scope [...]. A
    failed query still counts toward the 32-query cap, which is a safety ceiling
    rather than a scanning budget [...]. Mark numbers are local to one view, so never
     remember them [...].
7. Never construct a semantic target, key, direction, status, reason, mark, or anchor
     from a query result or arbitrary computation. [...]
8. remember is an explicit persistent side effect; use it only for a durable fact or
    value needed later.
9. NoOp is a last resort, only when no safe action, productive fallback, or correct
    Finish action is possible. [...]
10. emit is terminal: no statement executes after it. [...]

EXAMPLES

emit(Click("the visible trusted primary-navigation link labeled 'Categories'"))

author_marks = [40, 43, 46]
for candidate in author_marks:
    try:
        username = query(candidate, "What exact username does this author link show
    ?", str)
        if username == "CameronKelsey":
            emit(Click("the 'CameronKelsey' author link inside the queried mark",
    anchor=candidate))
    except QueryUnavailable:
        continue
emit(Scroll(direction="down", amount=1000))

try:
    address = query(4, "What is the complete shipping address?", str)
    emit(Type("the visible trusted shipping-address field", address))
except QueryUnavailable:
    emit(Scroll(direction="down", amount=600))
```

Actuator: User-Message Template

```
{feedback}

CURRENT TRUSTED VIEW
Only marks in this current view may be queried or used as anchors. Ignore mark
    numbers in prior messages.

CONTENT-FREE MASK MANIFEST
{marks}

MASKED ACCESSIBILITY TREE
{a11y}

PERSISTENT PROTECTED MEMORY SCHEMA
{memory}

The current trusted-view screenshot is attached. Write the complete Python action
    transaction for the single next semantic action. [...] Output exactly one fenced
    Python block and no prose.
```

Isolated query model. Each call supplies the committed question, return-type description, and isolated screenshot. When available, the corresponding accessibility representation is appended. The model receives no separate task instruction, interaction history, or protected-storage contents.

Isolated Query Model: Prompt

```
You are shown one selected region of a web page. A containing region may also be
    visible when needed to expose the selection; unrelated regions are blacked out.
    Answer ONLY the committed question about the selected region.
You may also be given a matching accessibility-tree snippet - use it to read the
    selected content more precisely, without drawing information from unrelated
    content.

Question: {rho}
Return type: {typedesc}

Respond with ONLY this JSON object and nothing else:
{"outcome":"answered|unknown", "value": <the {T} value, or null>}
- "answered": [...] put the typed value [...].
- "unknown": you cannot answer [...]. value null.
- If the value is TEXT you are transcribing [...], copy it VERBATIM [...]; do NOT
    paraphrase, summarize, translate, shorten, re-order, or fix spelling / spacing /
    punctuation.
No explanation. JSON only.

[appended only if a region a11y snippet exists:]
Accessibility tree for the visible region (structure only; everything outside your
    region is scrubbed):
{a11y_text}
```

Grounding model. The following prompts correspond to the Sonnet computer-use backend. The runtime supplies the operation, target description, and optional anchor through {`action`}, withholding `Type.value` and `Press.key`. The masked screenshot and manifest coordinates use a 1280×720 frame.

Grounding Model: System Prompt

```
You ground one trusted semantic browser action into screen coordinates.
```

```
You receive a TRUSTED VIEW in which every untrusted region is covered by a numbered
    black mask. The semantic action was already chosen by a protected transaction.
    Locate exactly the control specified by that action without changing its
    operation, target, or effect.
[...] Do not infer the mask's hidden content, resolve an underspecified target by
    guessing, or select a different masked unit. [...] If the action does not
    identify one control reliably, return not-found [...].
Return exactly one JSON object and no prose:
{"x": <integer viewport x>, "y": <integer viewport y>}
[...] If the control cannot be located reliably, return:
{"error":"not-found"}
```

Grounding Model: User-Message Template

```
The attached image IS the current 1280x720 screen - you already have it.

SEMANTIC ACTION
{action}

CONTENT-FREE MASK MANIFEST (boxes are in this 1280x720 frame)
{marks}

Locate exactly the control named by the semantic action. Using the computer tool,
    emit EXACTLY ONE left_click on the CENTER of that control. Do NOT take a
    screenshot [...]. If the control is NOT visible on this screen, do not click -
    reply with the text {"found": false} and make no tool call.
```

E Evaluation Details

This appendix supplements Section 5 with details of the benchmark tasks, execution protocol, and isolated-query analysis.

E.1 WebArena Tasks and Evaluation

WebArena (Zhou et al., 2024) provides self-hosted web applications and tasks specified by natural-language instructions. Agents complete tasks through browser interaction, and automated evaluators assess their final answers and the resulting application state. We use the original task instructions and evaluators without modification.

Applications and tasks. Our evaluation covers four applications. Postmill is a discussion forum with tasks involving posts, comments, voting, and user profiles. Shopping is an online storefront with tasks involving product search, reviews, purchases, and order history. Shopping Admin supports administrative tasks such as customer lookups, sales analysis, and catalog or order updates. GitLab provides repository-management tasks, including information retrieval and updates to issues, merge requests, and profiles.

From WebArena's 812 tasks, we select 400 single-site tasks, with 100 per application. The selection ensures at least 30 tasks using answer-based `string_match` evaluation per application and randomly fills the remaining positions. We use sampling seed 42 and retain the same task set across agents, model configurations, and repetitions.

Task evaluation. WebArena uses three kinds of checks. Answer-based checks compare the agent's final response with reference answers using exact matching, required phrases, or benchmark-configured LLM-based matching. URL checks compare the final page URL with the required path and query parameters. Application-state checks inspect specified page content or DOM elements to verify the requested outcome, such as an updated profile. A task succeeds only when all its evaluation requirements are satisfied.

Table 6: **Access to untrusted information.** Transactions is the percentage issuing at least one query; Queries/task is the mean per run. Structured and String are percentages of query calls. All pools counts across the four applications. Percentages are rounded independently.

Model	Metric	Shopping	GitLab	Postmill	Shopping Admin	All
Opus	Transactions (%)	23.4	27.1	22.0	12.3	20.8
	Queries/task	10.5	9.7	16.7	6.3	10.8
	Structured (%)	74.2	72.4	24.6	24.3	47.4
	String (%)	25.8	27.6	75.4	75.7	52.6
Gemini	Transactions (%)	23.1	25.4	21.8	11.8	20.2
	Queries/task	10.2	23.1	15.5	4.6	13.4
	Structured (%)	79.2	81.9	46.0	21.9	65.8
	String (%)	20.8	18.1	54.0	78.2	34.2
GPT	Transactions (%)	22.4	23.4	18.4	10.6	18.1
	Queries/task	9.0	8.6	11.7	6.5	8.9
	Structured (%)	75.6	43.3	14.8	23.4	38.5
	String (%)	24.4	56.7	85.2	76.5	61.5

E.2 Execution Protocol

We repeat each task five times per agent and model configuration. Before each run, the application container is recreated from its initial image, restoring the initial database state. Authentication is regenerated, and the run uses a fresh browser profile.

Each run has a 40-minute time limit. Secure-CUA permits up to 50 transaction attempts, including rejected transactions. Vanilla-CUA permits up to 50 browser-action turns, while CaMeL-CUA executes its whole-task program under a limit of 50 GUI actions and 100 quarantined queries. Thus, the interaction budget is counted according to each agent's execution interface.

CaMeL-CUA generates one whole-task program without replanning. Secure-CUA generates a new transaction at each turn. After a transaction is rejected, the runtime records feedback, obtains a fresh observation, and proceeds to the next turn; it does not retry the transaction within the same turn. Runs terminate on completion, budget exhaustion, an unrecoverable error, or timeout.

E.3 Access to Untrusted Information

We analyze the recorded query calls from 6,000 Secure-CUA runs: 100 tasks per application, four applications, three model configurations, and five repetitions. Table 6 reports the results by model and application.

Query frequency and volume. The fraction of transactions issuing queries is the number of recorded transactions containing at least one query call divided by the total number of recorded transactions. This denominator includes rejected and partially executed transactions. Mean queries per task is the total number of query calls divided by the number of runs. All issued query attempts are counted, including unsuccessful attempts and calls made during transactions that are subsequently rejected.

Return types and aggregation. We classify calls by their declared return type. Structured queries request Boolean, integer, floating-point, or enumeration values; string queries request free-form text. Their percentages are computed over all query calls. The All column pools the underlying counts across applications for each model, rather than averaging application-level percentages. Overall results pool all 6,000 runs.

Across these runs, 19.6% of transactions issue queries, with an average of 11.0 queries per task. Of the 66,175 query calls, 34,688 request structured values and 31,487 request strings, corresponding to 52.4% and 47.6%, respectively. Query volume also varies substantially: Gemini averages 23.1 queries per task on GitLab and 4.6 on Shopping Admin.