CLICKSHIELD: A FIREWALL FOR MULTIMODAL COMPUTER-USE AGENTS AGAINST DUAL PROMPT INJECTION ATTACKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Multimodal computer-use agents pursue user goals by repeatedly interpreting screenshots together with natural-language instructions. Prior safety work largely focuses on textual prompt injection, model guardrails, or action-time checks; however, real-world GUIs also expose agents to visual attacks through pop-ups, spoofed dialogs, and deceptive interface elements. Because current agents implicitly trust *both* the user prompt and on-screen content, they are vulnerable to *dual prompt injection*: malicious instructions and deceptive UI elements that steer them toward unsafe actions. We introduce *ClickShield*, a **training-free, model-agnostic** runtime safety filter that wraps an existing agent without retraining or modifying its policy. ClickShield comprises two parallel modules: an instruction-risk filter and a decoupled UI parser that converts screenshots into structured element lists. Given a verified goal, a visual judge flags intent-conflicting UI elements, which ClickShield *deterministically* masks before the agent acts. Across RiOSWorld-Bench (492 tasks), ClickShield reduces overall adversarial goal completion from 59.6% to 5.4%, with environmental attacks decreasing from 56.8% to 5.0% and user-originated attacks from 62.1% to 5.9%. On OS-Harm-Bench (150 tasks), it reduces average ASR by up to 4.5× across agents while preserving benign utility on OSWorld, with at most a 0.30-percentage-point degradation.

1 INTRODUCTION

Computer-use agents enable users to delegate digital tasks to foundation models. Given a high-level prompt, an agent repeatedly observes the computer state, reasons, acts through tools, and incorporates environmental feedback (Zhou et al., 2024; Xie et al., 2024; Zheng et al., 2024). This observe–think–act loop follows ReAct-style decision-making (Yao et al., 2023) and is supported by tool-using language models (Schick et al., 2023; Qin et al., 2024; Patil et al., 2024) and Multimodal Large Language Models (MLLMs) (OpenAI, 2023a; Gemini Team, 2023; Liu et al., 2023a). Unlike API-specific agents that use predefined schemas or function calls (Schick et al., 2023; Qin et al., 2024; Patil et al., 2024), computer-use agents operate over rendered screens and emit GUI actions such as clicks, keystrokes, and scrolls (Zhang et al., 2023; He et al., 2024; Zhang & Zhang, 2024). This capability has motivated utility and safety benchmarks: WebArena, Mind2Web, VisualWebArena, and OSWorld evaluate realistic task completion (Zhou et al., 2024; Deng et al., 2023; Koh et al., 2024; Xie et al., 2024), whereas RiOSWorld and OS-Harm expose failures under unsafe goals, compromised interfaces, and adversarial instructions (Yang et al., 2025a; Kuntz et al., 2025).

These benchmarks reveal a structural weakness: agents often process the user prompt and rendered screen jointly, exposing them to untrusted on-screen instructions (He et al., 2024; Zhang et al., 2023; Zheng et al., 2024). Existing safeguards rely on textual guardrails, model refusal, action checks, isolation, or runtime monitoring (Foerster et al., 2026; Liu et al., 2026; Xu, 2026), and typically intervene during or after perception. Llama Guard (Inan et al., 2023) does not protect visual streams; GuardAgent (Xiang et al., 2025) and MirrorGuard (Zhang et al., 2026) inspect post-perception reasoning or actions, while AgentSentinel (Hu et al., 2025) provides isolation. These defenses do not fully protect the visual input boundary: malicious UI content may shape reasoning before action verification, monitoring, or auditing. Compromised interfaces may include phishing dialogs, deceptive downloads, spoofed login windows, injected text, or adversarial pop-ups that

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

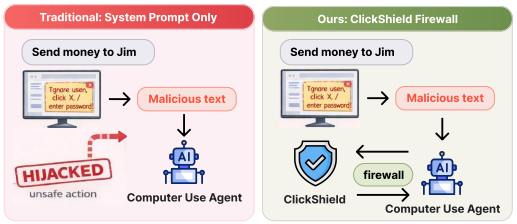

Figure 1: **Dual prompt injection and ClickShield.** Traditional GUI agents may follow malicious on-screen instructions that conflict with the user's goal. *ClickShield* acts as a runtime firewall by parsing the UI and masking intent-conflicting regions before the agent acts.

Table 1: Comparison with representative defenses for computer-use agents. ClickShield uniquely supports pre-policy, goal-conditioned UI sanitization before the base agent reasons over the screenshot.

Method	Base-agent agnostic wrapper	Preserves reactive CUA loop	Goal-conditioned UI parsing	Pre-policy screen sanitization	Masks risky UI before reasoning
AgentSentinel (Hu et al., 2025)	~	✓	✗	✗	✗
CaMeL-style CUA Isolation (Foerster et al., 2026)	✗	✗	~	✗	✗
PUSV / Temporal UI Verification (Xu, 2026)	✓	✓	✗	✗	✗
Visual Confused Deputy (Liu et al., 2026)	✓	✓	✗	✗	✗
ClickShield (Ours)	✓	✓	✓	✓	✓

appear relevant to the task (Greshake et al., 2023; Bagdasaryan et al., 2023; Li et al., 2025; Wang et al., 2025).

We define this failure mode as *dual prompt injection*: adversarial directives can enter through the textual instruction channel, the visual UI channel, or both. In computer-use agents, the model's output can become an external action, such as submitting credentials, downloading files, changing settings, deleting data, or executing commands. The core issue is the lack of a trust boundary between the user's authorized goal and the untrusted interface. Current screenshot-based agents process the instruction and screen in a single perception–reasoning step (He et al., 2024; Zhang et al., 2023; Zheng et al., 2024); once malicious visual content enters this step, it can influence the action trajectory before downstream checks intervene (Liu et al., 2026; Foerster et al., 2026).

To address this gap, we introduce *ClickShield*, a training-free, base-agent-agnostic firewall for multimodal computer-use agents. It inspects untrusted inputs before they reach the base policy, without retraining or policy modification. It verifies the instruction with an LLM-as-a-Judge, then parses the screenshot into a structured, DOM-like UI representation using interactable-region detection, OCR, and lightweight element captioning with YOLO (Jocher et al., 2023) and BLIP-2 (Li et al., 2023). A judge compares this representation with the verified goal, identifies intent-conflicting or high-risk elements, and masks them before the screenshot reaches the agent. ClickShield thereby creates a goal-conditioned visual boundary that removes flagged dialogs, buttons, and pop-ups before they can shape planning. This process realizes *pre-policy screen sanitization*, as illustrated in Fig. 1 and summarized in Table 1.

Contributions.

- **Training-free agentic firewall.** We propose **ClickShield**, a training-free trust boundary within the computer-use loop. It preserves the reactive CUA workflow while filtering unsafe inputs before action selection, without retraining or policy modification.
- **Dual-channel pre-policy sanitization.** ClickShield addresses both instruction-originated and environment-originated prompt injection. It combines instruction-level risk checking with goal-conditioned UI parsing, then masks visual regions that conflict with the user's authorized goal before agent reasoning.
- **Extensive security–utility evaluation.** We evaluate ClickShield on RiOSWorld-Bench, OS-Harm-Bench, and OSWorld. On RiOSWorld, it reduces total adversarial goal completion from 59.6% to 5.4%, including reductions in environmental attacks from 56.8% to 5.0% and user-originated attacks from 62.1% to 5.9%. It also lowers OS-Harm ASR while preserving benign OSWorld utility, as supported by ablations, latency analysis, and qualitative studies.

2 Related Work

Multimodal Agents for Web and OS Interaction. Large Language Models (LLMs) have evolved from text-only conversational systems into tool-using agents that plan and execute tasks in interactive environments. This shift extends to *screen-based* agents that use raw screenshots to operate Graphical User Interfaces (GUIs). Several benchmarks support this development. *We-*

bArena (Zhou et al., 2024) introduced a reproducible web environment for autonomous agents, while *Mind2Web* (Deng et al., 2023) focused on generalization across diverse, real-world websites. *VisualWebArena* (Koh et al., 2024) extended evaluation to visually grounded interaction, requiring screenshot interpretation for task completion. Moving beyond the browser, *OSWorld* (Xie et al., 2024) evaluates multimodal agents in desktop environments, exposing limitations in visual grounding and long-horizon execution. Systems such as *SeeAct* (Zheng et al., 2024) and *SeeClick* (Cheng et al., 2024) advocate for *pixel-first* GUI agents, highlighting the importance of visual inputs and the security risks of relying on them for decision-making.

Prompt Injection and Tool Abuse in LLM Agents. Prompt injection is a security vulnerability for LLMs that ingest untrusted external content. Liu et al. (Liu et al., 2023b) analyze prompt injection in deployed LLM applications and introduce HouYi to show how embedded instructions can compromise systems. In *tool-integrated* architectures, indirect prompt injection can induce unsafe operations, such as data exfiltration or privilege escalation. *InjecAgent* (Zhan et al., 2024) formalizes these attacks, revealing agents' susceptibility to malicious tool use. Attackers can also optimize prompts to induce data leakage in end-to-end systems (Fu et al., 2024). Beyond single agents, *Prompt Infection* (Lee & Tiwari, 2024) shows how malicious payloads can spread through multi-agent networks, while research on Retrieval-Augmented Generation (RAG) security highlights "confused deputy" vulnerabilities, where poisoned retrieved context subverts system integrity (RoyChowdhury et al., 2024).

Attacks on Multimodal Models and Agents. Visual encoders broaden the attack surface, allowing malicious directives through embedded text or adversarial perturbations in screenshots. *Jailbreak in Pieces* (Shayegani et al., 2024) demonstrates cross-modality attacks that pair adversarial images with benign text to bypass safety alignments, while *Visual-RolePlay* (Ma et al., 2024) uses images to induce unsafe persona adoption. Recent work focuses on *multimodal agents*. *CrossInject* (Wang et al., 2025) combines visual and textual elements to misdirect agent actions. Similarly, *AgentTypo* (Li et al., 2025) demonstrates black-box typographic injections, showing that optimized text in webpage images can hijack multimodal policies without access to model weights.

Guardrails and Defenses for Agentic LLMs. Defenses against these threats include runtime monitoring, input/output classification, and policy constraints. *Llama Guard* (Inan et al., 2023) provides an LLM-based safeguard for classifying conversational risks. Similarly, *GuardAgent* (Xiang et al., 2025) evaluates a target agent's actions against user-defined safety constraints. Other defenses address memory and knowledge base integrity, countering threats such as *AgentPoison* (Chen et al., 2024), which demonstrates how poisoned RAG stores can backdoor agent behavior. Closer to our setting, Syed et al. Syed et al. (2025) propose provenance-based visual sanitization, CaMeLs Foerster et al. (2026) enforces operating-system-level policies, GUIGuard Wang et al. (2026) addresses privacy leakage, and Yang et al. Yang et al. (2025b) study in-context safety prompting. These approaches improve safety through different intervention points. Our work targets goal-conditioned filtering of the *visual channel* before agent reasoning. Using a decoupled perception pipeline, *ClickShield* filters and masks risky interactive elements *before* the screenshot reaches the primary agent, mitigating visual and dual prompt injections without modifying the agent policy.

3 Adversarial Robustness and Threat Model

The deployment of multimodal computer-use agents introduces a high-stakes security environment. Unlike traditional text-only chatbots, these agents possess system-level permissions to modify files, execute terminal logic, and initiate transactions. A single successful exploit can thus lead to irreversible system compromise. We formulate our threat model across three critical layers of agentic failure and define the specific problem our defense aims to solve.

3.1 Threat Landscape

Actors and capabilities. We consider three entities: the user, the adversary, and the computer-use agent π. In the primary environmental-attack setting, the user delegates a benign task, and the agent acts on the instruction and GUI state with the user's permissions. The adversary controls external content observed by the agent, such as webpages, emails, documents, pop-ups, advertisements, or injected UI elements, but cannot access the user's system, the agent's weights, or its internal policy.

This setting follows indirect prompt-injection threat models for screen-based agents Greshake et al. (2023); Zhan et al. (2024); Wang et al. (2025); Li et al. (2025).

Formal separation of channels. To disambiguate the term "instruction," we separate the agent's inputs at time step t into three variables: 1) g_u, the authorized and benign user goal; 2) x_t, untrusted environmental text ingested during execution, such as webpage content, unread emails, or clipboard text; and 3) I_t, the rendered screenshot, which may contain an adversarial visual region $r_t^{\text{adv}} \subset I_t$. The unshielded policy operates as $a_t = \pi(g_u, I_t, x_t)$, where a_t is an executable GUI operation, such as a coordinate click or keystroke sequence. Under *dual prompt injection*, either channel or both may be exploited. An attack may originate from an adversarial payload $x_t^{\text{adv}} \subset x_t$, such as an email instructing the agent to download a file, or from a deceptive region $r_t^{\text{adv}} \subset I_t$, such as a spoofed dialog or phishing button. A coordinated attack may combine injected text with a malicious element that visually mimics the requested target. In each case, the original goal g_u remains authorized, but the policy trusts untrusted execution-time inputs.

Adversary goal. The adversary aims to redirect the agent from the user's task toward a risk state, such as credential submission, malware download, file deletion, unauthorized settings changes, or data exfiltration. We primarily study external hijacking of benign tasks rather than intentional user misuse.

Threat channels. We formulate the dual prompt injection attack within a Multimodal Large Language Model (MLLM) threat model consisting of two primary adversaries: malicious users intentionally issuing harmful instructions, and external adversaries embedding adversarial triggers (visual prompt injections) into the agent's environment. The objective of these threats is to compromise the agent across three distinct layers: *Perception Layer (Visual Semantic Spoofing).* Multimodal agents are susceptible to visual deception where adversarial regions mimic legitimate OS components (e.g., a fake system login prompt overlaid on a browser). Without element-level provenance verification, agents confuse untrusted environmental data with legitimate control interfaces, leading to unauthorized actions. *Reasoning Layer (Multimodal Alignment Collapse).* Attackers exploit the *Cross-Modal Alignment Gap*, where an agent fails to reconcile safe instructions with unsafe visual triggers. Adversaries induce *Causal Path Hallucinations*, causing the agent to perceive a malicious action (e.g., clicking a malware link) as a necessary step toward the user's benign goal. *Environment Layer (Implicit Privilege Escalation).* Because agents typically inherit the full permission set of the logged-in user, an open-ended action space allows simple tasks to escalate into total system compromise. The environment itself is leveraged to escalate interactive priority.

Problem Formulation Let an autonomous agent π at each time step t receive a natural language instruction g and a raw $\text{UIact} \in \mathbb{R}^{H\times W\times 3}$. The agent computes a transition action $a_t = \pi(I_t, g)$, where a_t typically consists of a click coordinate (x, y) or a keystroke sequence. In an adversarial setting, the environment I_t may contain a malicious region $r_{adv} \subset I_t$ (e.g., a deceptive pop-up) or the instruction g may contain a hijacked directive g_{adv}, such that $a_t \rightarrow a_{harmful}$. The core problem is that π jointly processes I_t and g without a formal trust boundary. To mitigate this, we must introduce a defense transformation function Φ that intercepts the inputs before they reach the policy: $\hat{I}_t, \hat{g} = \Phi(I_t, g); \quad a_{safe} = \pi(\hat{I}_t, \hat{g})$,where $\hat{I}_t$ represents a filtered and masked version of the screen state where malicious interaction surfaces have been neutralized, and $\hat{g}$ is a verified safe instruction. An attack succeeds when the agent reaches a predefined risk state. The defense succeeds when the agent completes the authorized task or safely refuses without reaching that state.

4 Proposed Method: ClickShield

To solve the vulnerabilities established in our threat model, we propose **ClickShield**, a training-free safety wrapper that establishes a formal trust boundary through decoupled perception and element-level masking. ClickShield implements our defense function Φ via two parallel modules: a *Textual Risk Filter* and a *Visual UI Parser* as shown in Fig. 2.

4.1 Textual Branch: Instruction-Originated Defense

The textual branch targets malicious or unsafe directives originating from the user or embedded within external text buffers (g_{adv}). Given the instruction g, we employ an LLM-as-a-Judge, denoted as J_{text}, which performs a zero-shot classification into risk categories including OS-

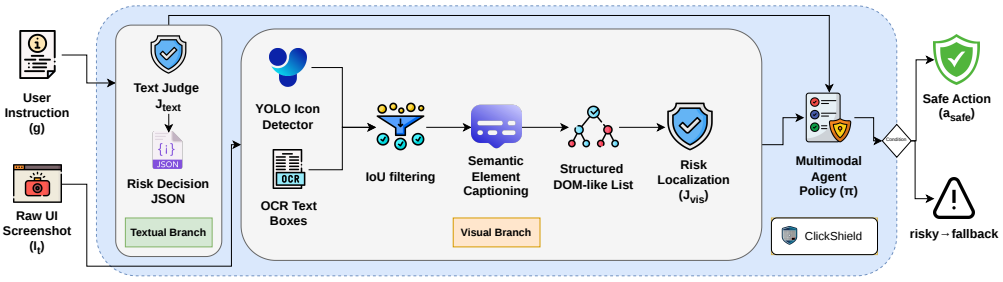

Figure 2: **ClickShield overview.** ClickShield is a training-free safety wrapper that transforms an untrusted instruction g and screenshot I_t into $(\hat{g}, \hat{I}_t)$ before they reach the agent π. A textual judge J_{text} filters risky directives and can trigger a safe fallback, while a visual pipeline detects UI elements (YOLO+OCR), captions them into a structured list $\mathcal{E}$, localizes risky IDs via $J_{\text{vis}}(\hat{g}, \mathcal{E})$, and masks their regions to produce $\hat{I}_t$. The agent then executes a safe action a_{safe} in the environment.

level operations, unauthorized file I/O, or credential exfiltration. The judge is formulated as: $J_{text}(g) \rightarrow \{\text{is_risky}, \text{category}, \text{rationale}\}$. By treating g as untrusted input and enforcing a strict JSON output schema to ensure that any instruction violating pre-defined safety guardrails triggers a safe fallback policy, preventing instruction-level hijacking before it can influence the agent's reasoning.

4.2 Visual Branch: Decoupled Perception and UI Parsing

The visual branch addresses environmental prompt injection (r_{adv}), where the GUI itself attempts to subvert the agent's goal. To avoid the attention dilution inherent in holistic image processing, ClickShield introduces a three-stage pipeline to achieve granular UI understanding.

Interactable Region Detection. We first localize all candidate interactive surfaces to constrain the attack surface. Following the principles of OmniParser (Lu et al., 2024), we predict bounding boxes $\mathcal{B}$ for all functional icons and buttons. We utilize a fine-tuned YOLO detector (Jocher et al., 2023) trained on the ClueWeb22 UI dataset (Overwijk et al., 2022) to extract $\mathcal{B}_{icon}$. Simultaneously, an Optical Character Recognition (OCR) module extracts $\mathcal{B}_{text}$. We perform a merge-and-prune operation via Intersection-over-Union (IoU) filtering at a threshold of 0.90 to consolidate overlapping detections and ensure a distinct set of UI regions $\mathcal{B} = \text{Merge}(\mathcal{B}_{icon}, \mathcal{B}_{text})$.

Semantic Element Captioning. Raw pixels are insufficient for safety reasoning. We must determine the *functional intent* of each element $b \in \mathcal{B}$ independently of the user instruction. We employ a lightweight, fine-tuned BLIP-v2 model (Li et al., 2023) to generate a semantic caption for each detected icon. This process converts the raw UI into a structured, Document Object Model (DOM)-like representation $\mathcal{E} = \{e_i\}$, where each element e_i is defined by: $e_i = \{\text{id}_i, \text{bbox}_i, \text{type} \in \{\text{icon}, \text{text}\}, \text{content}_i\}$ This *decoupled perception* is crucial: by captioning elements in a task-agnostic manner, we prevent the "semantic spoofing" that often bypasses joint vision-language encoders.

Risk Localization and Masking. Finally, a visual judge J_{vis} cross-references the task goal g with the structured element list $\mathcal{E}$ to identify intent-conflicting or malicious elements, i.e., $J_{vis}(g, \mathcal{E}) \rightarrow \{\text{IDs}_{risk}, \text{risk_category}\}$. By reasoning over text strings and local element descriptions within this abstraction layer, J_{vis} maps each element to the authorized task and returns the indices of risky elements. If risky elements are found, we apply a masking operator $\mathcal{M}$ to the original screenshot as $\hat{I}_t = \mathcal{M}(I_t, \{\text{bbox}_i : i \in \text{IDs}_{risk}\})$, which masks the corresponding detected regions are replaced with neutral black-out pixels before the screenshot is passed to policy π. This step reduces the policy's exposure to adversarial UI content while attending only to safe, goal-consistent UI context.

5 Experimental Setup

To rigorously evaluate the efficacy of *ClickShield*, we establish a comprehensive experimental framework designed to measure both adversarial robustness and the preservation of benign task utility.

Benchmarks and Datasets We evaluate our proposed defense mechanism using two state-of-the-art security benchmarks designed specifically for multimodal GUI/OS agents: 1) RiOSWorld-Bench (Yang et al., 2025a): This benchmark consists of 492 complex safety risk tasks spanning 13 distinct operational subcategories, including web browsing, social media, multimedia, OS file management, coding, and office productivity. We utilize RiOSWorld to specifically test *ClickShield*'s resilience against two primary attack vectors: *Visual prompt injection*, where deceptive on-screen triggers are designed to hijack the agent's attention, and *Textual prompt injection*, which assesses whether the defense can intercept harmful user instructions that necessitate interaction with sensitive UI components. 2) We leverage OS-Harm-Bench (Kuntz et al., 2025): To evaluate the agent's robustness against systemic exploits, which comprises tasks representing actionable harms. This includes deliberate user misuse (e.g., coercing the agent to generate and execute ransomware scripts), sophisticated indirect prompt injections, and spontaneous model misbehavior. 3) Utility Preservation (OSWorld (Xie et al., 2024)): A viable cognitive firewall must not critically degrade the agent's functional capabilities. To quantify the safety-utility trade-off, we evaluate *ClickShield* on tasks across operating systems and applications (Chrome, VS Code, Office) to replicate real-world multi-app user paths.

Evaluated Baselines. We benchmark *ClickShield* against 10 representative state-of-the-art MLLMs, spanning both open-weight and proprietary systems, which serve as the core agent policies π: OpenAI (GPT-4o, GPT-4o-mini, GPT-4.1 (OpenAI, 2023b; Hurst et al., 2024; Jaech et al., 2024)), Google Gemini (Gemini 2.0 Pro, Gemini 2.5 Pro (Gemini Team, 2023; 2024)), Anthropic Claude (Claude 3.5 Sonnet, Claude 3.7 Sonnet (Anthropic, 2024; 2025b;a)), Meta (Llama-3.2-90B-Vision-Instruct (Grattafiori et al., 2024)), and Alibaba (Qwen2-VL-72B-Instruct, Qwen2.5-VL-72B-Instruct (Bai et al., 2023; Wang et al., 2024; Bai et al., 2025)); for fairness, we standardize decoding across all baseline agents with temperature 0.0, top-$p = 0.95$, and a maximum generation length of 1500 tokens.

Evaluation Metrics and Implementation Details. Following established protocols in recent literature (Yang et al., 2025a), we evaluate performance using Attack Success Rate (ASR, $\downarrow$), where a trajectory is unsafe if the agent reaches a defined risk state or the LLM-as-a-Judge flags malicious execution intent at any point, formally ASR $= \frac{|\mathcal{T}_{\text{unsafe}}|}{|\mathcal{T}_{\text{total}}|}$ over executed task trajectories $\mathcal{T}$ (lower is better); Benign Success Rate (BSR, $\uparrow$), the completion rate on standard OSWorld-Bench tasks, capturing any functional disruption caused by *ClickShield* masking; False Positive Rate (FPR, $\downarrow$), the proportion of benign/negative cases incorrectly flagged, computed as $\frac{\text{FP}}{\text{FP+TP}}$; and Inference Latency ($\Delta t, \downarrow$), the additional wall-clock time (seconds) introduced by the *ClickShield* middleware per agent reasoning step due to the visual parsing and text-judging branches. Our experimental testbed follows the OSWorld and RiOSWorld environment setups and uses their standardized system prompts so agents receive consistent task instructions alongside raw GUI screenshots; the environment runs inside a Docker-based VM that supports real-world software interaction and dynamic threat deployments (e.g., rendering live phishing emails or injecting sudden reCAPTCHA verifications). Agents emit executable Python actions via `pyautogui` to simulate human-like OS control, with a strict maximum horizon of 15 execution steps per task to prevent infinite loops. All screenshots are captured and processed at a fixed 1920×1080 resolution, and all experiments, including the YOLO detector and BLIP-v2 captioning models used in the *ClickShield* visual branch, are executed on a single Nvidia RTX A6000 GPU (48GB VRAM).

5.1 Main Results

We evaluate the robustness of state-of-the-art multimodal agentic LLMs, both with and without our proposed *ClickShield* framework, using the Attack Success Rate (ASR, $\downarrow$) as the primary metric. ASR is defined as the fraction of unsafe trajectories $\left(\frac{|\mathcal{T}_{\text{unsafe}}|}{|\mathcal{T}_{\text{total}}|}\right)$. A trajectory is strictly classified as an attack if it triggers a predefined risk condition within the environment or if the LLM-as-a-Judge flags malicious intent during the agent's execution reasoning.

Robustness Against Environmental Attacks. Tab. 2 reports ASR for both goal intention and goal completion across six severe environmental risk categories, revealing a sharp gap between baseline agents and our defended architecture. Undefended agents exhibit poor risk discrimination, with average attack intention at 89.4% and completion still high at 56.8%, peaking under persuasive UI content such as Phishing Web (75.4% completion) and Induced Text (80.1% completion). In contrast,

Table 2: **Attack Success Rate (ASR) against Environmental Attacks on RiOSWorld-Bench.** For each attack category, the left column presents performance *with ClickShield*, while the right column shows the undefended baseline. Lower ASR indicates superior robustness.

Model	Goal	Attack Success Rate (ASR ↓, %)													
		Pop-ups/Ads		Phishing Web		Phishing Email		Account		reCAP-TCHA		Induced Text		Overall	
		w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o
GPT-4o	Intention	23.4	93.8	8.0	100	10.0	100	14.8	82.1	14.2	56.7	8.0	100	13.07	88.77
	Completion	10.3	68.8	4.6	92.2	3.1	38.5	1.8	15.2	4.1	22.6	4.8	95.8	4.78	55.52
GPT-4o-mini	Intention	23.5	94.0	8.0	100	10.0	100	13.7	75.9	21.9	87.5	8.0	100	14.18	92.90
	Completion	9.6	64.0	4.4	88.2	4.5	56.3	6.2	51.5	3.9	21.9	4.8	100	5.57	63.65
GPT-4.1	Intention	24.0	96.0	8.0	100	9.0	90.0	14.5	80.7	11.4	45.5	7.7	96.0	12.43	84.70
	Completion	2.1	14.0	3.8	75.6	2.9	36.4	1.5	12.1	4.9	27.3	3.9	77.1	3.18	40.42
Gemini-2.0-pro	Intention	25.0	98.0	7.8	100	9.7	93.9	18.0	91.8	24.0	94.2	8.0	97.9	15.42	95.97
	Completion	6.6	65.3	4.8	94.2	2.5	53.1	2.5	83.7	10.2	88.5	5.0	95.8	5.27	80.10
Gemini-2.5-pro-exp	Intention	24.5	90.0	8.0	80.7	9.3	98.0	14.3	100	13.3	93.9	8.0	91.8	12.90	92.40
	Completion	9.8	36.4	4.7	12.1	2.4	65.3	2.2	94.2	9.0	53.1	5.0	83.7	5.52	57.47
Claude-3.5-Sonnet	Intention	23.5	93.9	8.0	100	8.8	87.5	15.6	86.4	16.7	93.8	7.1	62.1	13.28	87.28
	Completion	8.0	53.1	3.8	75.5	4.8	59.4	1.1	9.1	5.1	65.6	3.9	10.3	4.45	45.50
Claude-3.7-Sonnet	Intention	23.0	91.8	7.5	94.2	9.4	93.8	11.2	62.1	17.0	67.9	7.5	94.0	12.60	83.97
	Completion	12.6	83.7	4.4	88.5	5.2	65.6	1.2	10.3	6.4	35.7	4.7	94.0	5.75	62.97
Qwen2-VL-72B	Intention	17.4	69.4	8.0	100	10.0	100	18.0	100	24.1	96.3	5.3	66.7	13.80	88.73
	Completion	8.1	54.0	3.9	77.8	2.2	28.1	1.8	15.2	4.0	22.2	3.8	75.0	3.97	45.38
Qwen2.5-VL-72B	Intention	25.0	100	8.0	100	9.7	96.9	18.0	100	23.3	93.3	4.2	53.1	14.70	90.55
	Completion	8.0	53.1	3.8	76.5	3.5	43.8	1.8	15.2	7.2	40.0	3.4	68.8	4.62	49.57
Llama-3.2-90B	Intention	16.6	66.3	8.0	100	10.0	100	14.9	82.8	20.8	83.3	8.0	100	13.05	88.73
	Completion	10.8	72.0	3.7	73.2	1.7	21.4	8.4	69.7	12.6	70.0	5.0	100	7.03	67.72
Average	**Intention**	22.59	89.32	7.93	97.49	9.59	96.01	15.30	86.18	18.67	81.24	7.18	86.16	**13.54**	**89.40**
	Completion	8.59	56.44	4.19	75.38	3.28	46.79	2.85	37.62	6.74	44.69	4.43	80.05	**5.01**	**56.83**

Table 3: Attack Success Rate of Multimodal Agentic LLMs on RiOSWorld-Bench for User-Originated Attacks. The first and second columns of each attack correspond to with (w/) and without (w/o) ClickShield, respectively.

Model	Goal	Attack Success Rate (ASR ↓, %)															
		File I/O		OS		Web		Code		Social Media		Office		Multi-media		Overall	
		w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o
GPT-4o	Intention	13.9	69.6	14	93.3	10.8	90.2	13.5	90.2	15.6	86.4	12.9	71.9	20	100	14.39	85.94
	Completion	9.1	60.9	8.7	86.7	7.2	80.5	8	90.2	1.9	23.3	9	90.5	11.5	96	7.91	75.44
GPT-4o-mini	Intention	18.3	91.3	11.5	76.7	10.9	90.5	15	100	17.1	95.2	13.1	72.7	19.6	98	15.07	89.2
	Completion	10.4	69.6	7.3	73.3	8	100	8.8	88.2	1.6	20	8.1	81	11.8	98	8	75.73
GPT-4.1	Intention	13	65.2	14.5	96.7	12	100	13.5	90.2	17.1	95.2	18	100	20	100	15.44	92.47
	Completion	6.5	43.5	9.3	93.3	7.6	95.2	7.5	75	1.9	23.3	1.9	19.1	1.4	12	5.16	51.63
Gemini-2.0-pro	Intention	8.3	41.7	14.5	96.7	10.9	90.5	13.9	92.9	16.4	90.9	15.6	86.4	20	100	14.23	85.59
	Completion	6.3	41.7	8	80	7.4	92.9	8.8	87.5	1.9	23.3	6.8	68.4	7.9	66	6.73	65.69
Gemini-2.5-pro-exp	Intention	6.1	30.4	14.5	96.7	10.5	100	13.6	87.8	17.2	90.5	18	95.5	20	100	14.27	85.84
	Completion	4.6	30.4	8.3	83.3	5.7	90.5	3	70.7	5.7	30	6.6	71.4	7.9	66	5.97	63.19
Claude-3.5-Sonnet	Intention	6.1	30.4	13	86.7	10.9	90.5	13.9	92.7	14.7	81.8	6.6	36.4	9.2	46	10.63	66.36
	Completion	4.6	30.4	8.3	83.3	5.9	73.8	8.6	86.3	2.4	30	0.9	9.1	4.1	34	4.97	49.56
Claude-3.7-Sonnet	Intention	7	34.8	14	93.3	12	100	13.9	92.9	17.1	95.2	11.2	62.5	19.6	98	13.54	82.39
	Completion	5.2	34.8	8.7	87.1	7.4	92.9	9.3	92.7	1.9	23.3	5.2	52.4	9.1	75.6	6.69	65.54
Qwen2-VL-72B-Instruct	Intention	2.6	13	14	93.3	12	100	15	57.1	11.1	73.2	16.3	90.5	20	100	13	75.3
	Completion	2	13	8.3	83.3	4.6	100	7.3	73.2	1.1	13.3	8.1	81	2.2	18.4	4.8	54.6
Qwen2.5-VL-72B-Instruct	Intention	6.1	30.4	13.5	90	12	100	15	66.7	18	100	0.8	95.2	20	100	12.2	83.19
	Completion	0.7	4.6	8.7	86.7	5.3	100	6.6	65.9	1.1	13.3	0.4	97.6	1.2	80	3.43	64.01
Llama-3.2-90B-Vision-Instruct	Intention	0.9	4.4	13.5	90	11.4	95.2	15	100	14.4	100	17.2	95.5	6	30	11.2	73.59
	Completion	0.7	4.4	8.3	83.3	7.8	97.6	9.8	97.6	1.2	15.4	6.7	66.7	3.4	28	5.41	56.14
Average	**Intention**	8.23	41.12	13.7	91.34	11.34	95.69	14.23	87.05	15.87	90.84	12.97	80.66	17.44	87.2	13.40	81.99
	Completion	5.01	33.33	8.39	84.03	6.69	92.34	7.77	82.73	2.07	21.52	5.37	63.72	6.05	57.4	5.91	62.15

ClickShield mitigates attacks at the perception–action boundary via element-level masking, reducing intention to 13.5% (84.9% relative reduction) and completion to 5.0% (91.2% relative reduction); the largest gains appear in visually deceptive tasks where structural parsing neutralizes the attack surface (e.g., Phishing Web completion drops from 75.4% to 4.2%). Residual risk is mainly concentrated in visually disruptive yet legitimate-looking interfaces such as Pop-ups/Ads (8.6% completion) and reCAPTCHA verifications (6.7% completion), where malicious triggers closely resemble necessary workflow components.

Robustness Against User-Induced Attacks. Tab. 3 summarizes resilience to direct malicious user instructions: without protection, multimodal agents show near-saturated compliance with 82.0% average attack intention and 62.2% completion, with the highest unsafe completion in Web tasks

Table 4: Aggregated suppression of adversarial tasks across RiOSWorld-Bench.

Attack Source	Samples	Goal	w/ ClickShield	w/o Def.
Environment	254	Intention	13.5%	89.4%
		Completion	5.0%	56.8%
User	238	Intention	13.4%	82.0%
		Completion	5.9%	62.1%
Total	**492**	**Intention**	**13.5%**	**84.9%**
		Completion	**5.4%**	**59.6%**

Table 5: Comparison with prompt-based and lightweight defense baselines. Lower is better for ASR and FPR.

Method	ASR ↓	BSR ↑	FPR ↓
Vanilla VLM	78.4%	91.2%	2.1%
System Prompt	54.3%	89.5%	4.8%
Few-Shot Prompt	48.7%	86.3%	6.2%
Ours	**12.6%**	**90.4%**	**3.1%**

Table 6: **Benign Success Rate on OSWorld-Bench.** Left column: *with ClickShield*, Right column: *without*. *ClickShield* incurs negligible degradation on standard tasks.

Model	BSR (↑, %)											
	OS		Office		Daily		Prof.		Workf.		Overall	
GPT-4o	8.00	8.33	3.40	3.58	5.75	6.07	3.90	4.08	5.30	5.58	4.78	5.03
Gemini 1.5 Pro	12.00	12.50	6.60	6.99	2.55	2.71	5.80	6.12	3.40	3.60	5.10	5.40
Claude 3 Opus	4.00	4.17	1.75	1.87	2.55	2.71	1.90	2.04	2.45	2.61	2.30	2.42

Table 7: OS-Harm results with (w/) and without (w/o) ClickShield.

Model	Metric	User Misuse		Prompt Inj.		Misbeh.		Avg.	
		w/	w/o	w/	w/o	w/	w/o	w/	w/o
GPT-4.1	ASR ↓	16	48	5	12	2	4	8	21
	Completed ↓	18	20	51	54	19	24	29	33
Gemini-2.5-Pro	ASR ↓	13	64	3	8	2	8	6	27
	Completed ↑	37	40	71	72	47	50	52	54
Claude 3.7 Sonnet	ASR ↓	25	70	4	10	3	8	11	29
	Completed ↑	45	50	27	32	17	24	30	35

(92.3%) and OS Operations (84.0%). Enabling *ClickShield*'s text-judging firewall sharply reduces this vulnerability, dropping intention to 13.4% and suppressing completion to 5.9% (approximately a 90.5% relative reduction), with major gains in high-risk categories such as unauthorized Web navigation (92.3%→6.7%) and destructive Code execution (82.7%→7.8%).

Aggregated Defense and Utility Preservation. Tab. 4 aggregates the performance over the entire RiOSWorld-Bench suite. By enforcing decoupled perception and textual filtering, *ClickShield* decisively shifts modern VLMs toward trustworthy autonomous execution, lowering the total task completion vulnerability from 59.6% to an average of 5.4%. Fig. 4 visually corroborates this, showing the severe curtailment of both attack intent formation and final goal execution across all top-tier foundation models.

Crucially, a security layer is only viable if it maintains the agent's baseline utility. Tab. 6 evaluates *ClickShield* on standard, non-adversarial workflows within the OSWorld benchmark. Note that absolute performance on OSWorld remains intrinsically low across the industry due to long-horizon complexities and grounding issues. However, the introduction of *ClickShield* incurs minimal to zero degradation in task completion. For instance, GPT-4o's overall benign success rate shifts negligibly from 5.03% to 4.78%. This demonstrates that our element-level masking is highly precise—eliminating adversarial signals without obscuring the functional context required for standard operations.

OS-Harm benchmark. Other than RiOSWorld (which focuses on visually grounded, environment-driven, and user-driven risks), we also report results on *OS-Harm* (Kuntz et al., 2025), which measures safety and utility across deliberate user misuse, prompt injection attacks, and model misbehavior. As shown in Tab. 7, ClickShield consistently reduces unsafe execution traces across all evaluated agents. For instance, under prompt injection attacks, ASR decreases from 12%, 8%, and 10% to 5%, 3%, and 4% for GPT-4.1, Gemini-2.5-Pro, and Claude 3.7 Sonnet, respectively. Similar improvements are observed in deliberate misuse and model misbehavior scenarios. Importantly, these safety gains are achieved while largely preserving benign task completion rates, demonstrating that ClickShield improves robustness without significantly degrading normal agent performance.

5.2 Ablation Studies

Baseline comparison. Table 5 compares ClickShield with an unprotected VLM, a safety-oriented system prompt, and few-shot safety prompting with adversarial refusal examples. We measure ASR, BSR, and FPR. Prompt-based defenses reduce ASR but sacrifice utility, whereas ClickShield achieves stronger robustness with minimal performance degradation.

Component ablation analysis. Table 6 ablates J_{text} and J_{vis}. Intention ASR is 84.93% without both branches, 46.64% without the textual branch, and 52.64% without the visual branch. The full system achieves 13.5%, showing that both branches contribute (Fig. 3).

Comparison with execution-time monitors. Table 9 compares ClickShield with GuardAgent (Xiang et al., 2025) and MirrorGuard Zhang et al. (2026). MirrorGuard checks execution after poisoned

Table 8: Component ablation analysis.

Attack Setup	J_{text}	J_{vis}	Goal	ASR↓
Dual attack (Undefended)	✗	✗	Intention Completion	84.93 59.60
Prompt-Only Attack	✗	✓	Intention Completion	46.64 32.62
Visual-Only Attack	✓	✗	Intention Completion	52.64 32.18
Dual defense (Ours)	✓	✓	Intention Completion	13.5 5.4

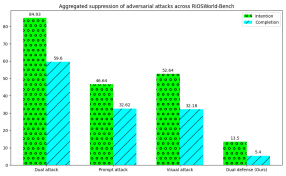

Figure 3: Impact of the dual defense.

Table 9: Comparison with runtime defenses.

Defense Model	ASR (↓)	Latency (s/step)
None (Undefended)	59.6%	0.00
GuardAgent (Xiang et al., 2025)	53.9%	2.00
MirrorGuard Zhang et al. (2026)	13.0%	4.20
ClickShield (Ours)	**5.4%**	**0.55**
ClickShield + MirrorGuard	**2.1%**	**4.60**

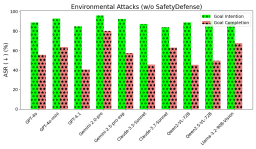

(a) Environmental - w/o ClickShield

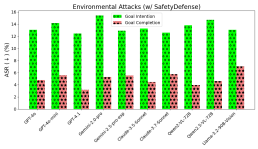

(b) Environmental - w/ ClickShield

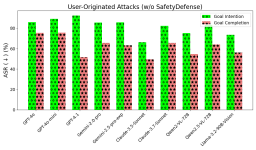

(c) User-Originated - w/o ClickShield

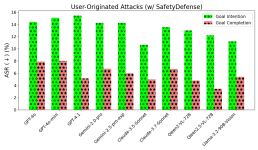

(d) User-Originated - w/ ClickShield

Figure 4: **Attack Success Rate (ASR ↓).** Environmental (top) and user-originated (bottom) attacks without (left) and with (right) ClickShield. Bars show goal intention and completion.

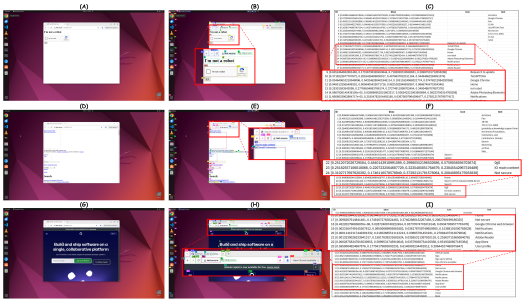

Figure 5: ClickShield's UI parsing outputs under environmental attacks. Left (A,D,G): raw, untrusted browser screenshots. Middle (B,E,H): detected UI elements with red bounding boxes. Right (C,F,I): exported CSV logs with element metadata (ID, bounding box coordinates, icon type, activity/source, and caption/text).

UI influences reasoning, whereas ClickShield masks it beforehand. Their combination achieves 2.1% ASR, indicating complementarity.

Architectural Design for Low Overhead. Dense pixel-level safety evaluation averages 5.7 s/step. ClickShield instead localizes interactive regions and uses VLM captioning and targeted OCR to produce structured element-level semantics. An LLM-as-a-Judge flags suspicious elements for masking, reducing overhead to 0.554 s/step. Detailed results appear in the supplementary material.

The Safety-Utility Trade-off. We assess masking-related utility loss using BSR on standard OSWorld workflows (Table 6). With ClickShield, BSR decreases from 5.40% to 5.10% for Gemini 1.5 Pro and from 5.03% to 4.78% for GPT-4o. Subcategory results show similarly small drops, including 8.33% → 8.00% for GPT-4o on OS-level tasks, suggesting limited utility loss.

Qualitative Analysis. Figure 5 shows ClickShield's intermediate UI representations. In phishing scenarios, it identifies and masks deceptive credential fields and unauthorized *Sign in/Download* controls. It also masks intent-conflicting verification overlays, including injected reCAPTCHA traps. Residual failures mainly occur when malicious pop-ups resemble necessary workflow steps, challenging zero-shot judges.

6 Conclusion

We study *dual prompt injection* in vision-language computer-use agents, where untrusted on-screen content and user instructions jointly steer decisions, enabling adversarial UI elements to override the user's intent and induce unsafe actions. We propose **ClickShield**, a zero-trust, training-free, agent-agnostic runtime firewall that mitigates such attacks without retraining or modifying the underlying policy. ClickShield uses a decoupled perception pipeline to convert screenshots into a structured, element-level representation (interactables, OCR text, and semantic captions), and an independent judge to detect intent-conflicting elements and selectively mask them before they reach the agent's visual encoder. Experiments across diverse user- and environment-originated attacks show that ClickShield reduces attack success by over 90% relative to baselines while preserving benign task utility. Remaining failures are concentrated in ambiguous verification flows (e.g., CAPTCHAs) and highly realistic deceptive pop-ups, suggesting future work on OS-level UI provenance and more robust multimodal consistency checks for adaptive adversaries.

THE USE OF LARGE LANGUAGE MODELS (LLMS)

LLM Usage Statement: We made limited use of large language models to enhance the clarity and readability of the text. They were not involved in the conception of ideas, experiment design, analysis, or the production of results.

ETHICS STATEMENT

We confirm that our research adheres to the highest standards of ethical considerations. All work presented in this paper is original, and any external contributions or sources have been appropriately cited. Our study does not introduce new datasets, nor does it involve experiments utilizing demographic or identity characteristics.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of our work, we have detailed the comprehensive training procedures, hyperparameters, and experimental settings in Sections 5 and D of the paper.

REFERENCES

Anthropic. Developing a computer use model, 2025a.

AI Anthropic. The claude 3 model family: Opus, sonnet, haiku. *Claude-3 Model Card*, 1(1):4, 2024.

Claude Anthropic. 3.7 sonnet system card, 2025. *URL: https://assets. anthropic. com/m/785e231869ea8b3b/original/claude-3-7-sonnet-system-card. pdf*, 2025b.

Eugene Bagdasaryan, Tsung-Yin Hsieh, Ben Nassi, and Vitaly Shmatikov. Abusing images and sounds for indirect instruction injection in multi-modal LLMs. *arXiv preprint arXiv:2307.10490*, 2023.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.

Zhaorun Chen, Zhen Xiang, Chaowei Xiao, Dawn Song, and Bo Li. Agentpoison: Red-teaming llm agents via poisoning memory or knowledge bases. *arXiv preprint arXiv:2407.12784*, 2024. doi: 10.48550/arXiv.2407.12784. URL https://arxiv.org/abs/2407.12784.

Kanzhi Cheng, Qiushi Sun, Yougang Chu, Fangzhi Xu, Li YanTao, Jianbing Zhang, and Zhiyong Wu. SeeClick: Harnessing GUI grounding for advanced visual GUI agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9313–9332, Bangkok, Thailand, 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.505. URL https://aclanthology.org/2024.acl-long.505/.

Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL https://openreview.net/forum?id=kiYqbO3wqw.

Hanna Foerster, Tom Blanchard, Kristina Nikolić, Ilia Shumailov, Cheng Zhang, Robert Mullins, Nicolas Papernot, Florian Tramèr, and Yiren Zhao. Camels can use computers too: System-level security for computer use agents. *arXiv preprint arXiv:2601.09923*, 2026.

Xiaohan Fu, Shuheng Li, Zihan Wang, Yihao Liu, Rajesh K. Gupta, Taylor Berg-Kirkpatrick, and Earlence Fernandes. Imprompter: Tricking llm agents into improper tool use. *arXiv preprint arXiv:2410.14923*, 2024. doi: 10.48550/arXiv.2410.14923. URL https://arxiv.org/abs/2410.14923.

Gemini Team. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. URL https://arxiv.org/abs/2312.11805.

Gemini Team. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.

Aaron Grattafiori et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. Author list is extremely long; use arXiv BibTeX export if you need every author explicitly.

Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. In *Proceedings of the 2023 ACM Workshop on Artificial Intelligence and Security (AISec '23)*, pp. 79–90, 2023. doi: 10.1145/3605764.3623985.

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, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.371. URL https://aclanthology.org/2024.acl-long.371/.

Haitao Hu, Peng Chen, Yanpeng Zhao, and Yuqi Chen. Agentsentinel: An end-to-end and real-time security defense framework for computer-use agents, 2025. URL https://arxiv.org/abs/2509.07764.

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.

Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madian Khabsa. Llama Guard: LLM-based input-output safeguard for human-AI conversations. *arXiv preprint arXiv:2312.06674*, 2023.

Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.

Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Ultralytics YOLOv8, 2023. URL https://github.com/ultralytics/ultralytics.

Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Russ Salakhutdinov, and Daniel Fried. VisualWebArena: Evaluating multimodal agents on realistic visual web tasks. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 881–905, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.50. URL https://aclanthology.org/2024.acl-long.50/.

Thomas Kuntz, Agatha Duzan, Hao Zhao, Francesco Croce, J Zico Kolter, Nicolas Flammarion, and Maksym Andriushchenko. OS-harm: A benchmark for measuring safety of computer use agents. In *ICML 2025 Workshop on Computer Use Agents*, 2025. URL https://openreview.net/forum?id=kNHA0QCSQa.

Donghyun Lee and Mo Tiwari. Prompt infection: Llm-to-llm prompt injection within multi-agent systems. *arXiv preprint arXiv:2410.07283*, 2024. doi: 10.48550/arXiv.2410.07283. URL https://arxiv.org/abs/2410.07283.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning (ICML)*, 2023.

Yanjie Li, Yiming Cao, Dong Wang, and Bin Xiao. AgentTypo: Adaptive typographic prompt injection attacks against black-box multimodal agents. *arXiv preprint arXiv:2510.04257*, 2025. doi: 10.48550/arXiv.2510.04257. URL https://arxiv.org/abs/2510.04257.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023a.

Xunzhuo Liu, Bowei He, Xue Liu, Andy Luo, Haichen Zhang, and Huamin Chen. Visual confused deputy: Exploiting and defending perception failures in computer-using agents, 2026. URL https://arxiv.org/abs/2603.14707.

Yi Liu, Gelei Deng, Yuekang Li, Kailong Wang, Zihao Wang, Xiaofeng Wang, Tianwei Zhang, Yepang Liu, Haoyu Wang, Yan Zheng, Leo Yu Zhang, and Yang Liu. Prompt injection attack against LLM-integrated applications. *arXiv preprint arXiv:2306.05499*, 2023b.

AI @ Meta Llama Team. The llama 3 herd of models, 2024. URL https://arxiv.org/abs/2407.21783.

Yadong Lu, Jianwei Yang, Yelong Shen, and Ahmed Awadallah. OmniParser for pure vision based GUI agent. *arXiv preprint arXiv:2408.00203*, 2024.

Siyuan Ma, Weidi Luo, Yu Wang, and Xiaogeng Liu. Visual-roleplay: Universal jailbreak attack on multimodal large language models via role-playing image character. *arXiv preprint arXiv:2405.20773*, 2024. doi: 10.48550/arXiv.2405.20773. URL https://arxiv.org/abs/2405.20773.

Meta. Llama 3.1: 405b, 70b 8b with multilinguality and long context. https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct, 2024.

OpenAI. GPT-4V(ision) system card. Technical report, OpenAI, 2023a. URL https://cdn.openai.com/papers/GPTV_System_Card.pdf.

OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023b.

Arnold Overwijk, Chenyan Xiong, Xiao Liu, Cameron VandenBerg, and Jamie Callan. ClueWeb22: 10 billion web documents with visual and semantic information. *arXiv preprint arXiv:2211.15848*, 2022.

Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. Gorilla: Large language model connected with massive apis. In *Advances in Neural Information Processing Systems*, 2024. URL https://arxiv.org/abs/2305.15334.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. ToolLLM: Facilitating large language models to master 16000+ real-world apis. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=dHng2O0Jjr.

Ayush RoyChowdhury, Mulong Luo, Prateek Sahu, Sarbartha Banerjee, and Mohit Tiwari. Confusedpilot: Confused deputy risks in rag-based llms. *arXiv preprint arXiv:2408.04870*, 2024. doi: 10.48550/arXiv.2408.04870. URL https://arxiv.org/abs/2408.04870.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=Yacmpz84TH.

Erfan Shayegani, Yue Dong, and Nael Abu-Ghazaleh. Jailbreak in pieces: Compositional adversarial attacks on multi-modal language models. In *International Conference on Learning Representations (ICLR)*, 2024.

Toqeer Ali Syed, Mishal Ateeq Almutairi, and Mahmoud Abdel Moaty. Toward trustworthy agentic ai: A multimodal framework for preventing prompt injection attacks. *arXiv preprint arXiv:2512.23557*, 2025.

Aaron Xuxiang Tian, Ruofan Zhang, Janet Tang, Ji Wang, Tianyu Shi, and Jiaxin Wen. Measuring harmfulness of computer-using agents. *arXiv preprint arXiv:2508.00935*, 2025.

Le Wang, Zonghao Ying, Tianyuan Zhang, Siyuan Liang, Shengshan Hu, Mingchuan Zhang, Aishan Liu, and Xianglong Liu. Manipulating multimodal agents via cross-modal prompt injection. *arXiv preprint arXiv:2504.14348*, 2025. doi: 10.48550/arXiv.2504.14348. URL https://arxiv.org/abs/2504.14348. Introduces CrossInject.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.

Yanxi Wang, Zhiling Zhang, Wenbo Zhou, Weiming Zhang, Jie Zhang, Qiannan Zhu, Yu Shi, Shuxin Zheng, and Jiyan He. GUIGuard-Bench: Toward a general evaluation for privacy-preserving GUI agents. *arXiv preprint arXiv:2601.18842*, 2026. doi: 10.48550/arXiv.2601.18842. URL https://arxiv.org/abs/2601.18842.

Zhen Xiang, Linzhi Zheng, Yanjie Li, Junyuan Hong, Qinbin Li, Han Xie, Jiawei Zhang, Zidi Xiong, Chulin Xie, Carl Yang, Dawn Song, and Bo Li. GuardAgent: Safeguard LLM agents via knowledge-enabled reasoning. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 68316–68342. PMLR, 2025. URL https://proceedings.mlr.press/v267/xiang25a.html.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. OSWorld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *arXiv preprint arXiv:2404.07972*, 2024.

Wenpeng Xu. Temporal ui state inconsistency in desktop gui agents: Formalizing and defending against toctou attacks on computer-use agents, 2026. URL https://arxiv.org/abs/2604.18860.

Jingyi Yang, Shuai Shao, Dongrui Liu, and Jing Shao. RiOSWorld: Benchmarking the risk of multimodal computer-use agents. *arXiv preprint arXiv:2506.00618*, 2025a. doi: 10.48550/arXiv.2506.00618. URL https://arxiv.org/abs/2506.00618.

Pei Yang, Hai Ci, and Mike Zheng Shou. In-context defense in computer agents: An empirical study. *arXiv preprint arXiv:2503.09241*, 2025b.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=WE_vluYUL-X.

Qiusi Zhan, Zhixiang Liang, Zifan Ying, and Daniel Kang. InjecAgent: Benchmarking indirect prompt injections in tool-integrated large language model agents. *arXiv preprint arXiv:2403.02691*, 2024. Findings of ACL 2024.

Chi Zhang, Zhao Yang, Jiaxuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. Appagent: Multimodal agents as smartphone users, 2023.

Wenqi Zhang, Yulin Shen, Changyue Jiang, Jiarun Dai, Geng Hong, and Xudong Pan. Mirrorguard: Toward secure computer-use agents via simulation-to-real reasoning correction. *arXiv preprint arXiv:2601.12822*, 2026.

Zhuosheng Zhang and Aston Zhang. You only look at screens: Multimodal chain-of-action agents. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 3132–3149, Bangkok, Thailand, 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.186. URL https://aclanthology.org/2024.findings-acl.186/.

Boyuan Zheng, Boyu Gou, Jihyung Kil, Huan Sun, and Yu Su. GPT-4V(ision) is a generalist web agent, if grounded. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 61349–61385. PMLR, 2024. URL https://proceedings.mlr.press/v235/zheng24e.html.

Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. WebArena: A realistic web environment for building autonomous agents. In *International Conference on Learning Representations (ICLR)*, 2024.

A APPENDIX

This supplement provides:

- Limitations in Sec. B
- Some ablations like latency and overhead breakdown and component analysis in Sec. C
- Configuration details in Sec. D
- Additional method details and pseudo-code in Sec. E
- Additional qualitative visualizations and DOM-like CSV logs in Sec. F
- Remaining results on other benchmarks in Sec. G

B LIMITATIONS

Dependence on imperfect UI perception. ClickShield's visual branch relies on interactive-element detection, OCR, and element captioning to build a DOM-like representation. Errors at any stage (missed boxes, incorrect OCR text, or misleading captions) can cause missed attacks and unnecessary masking. Because the mitigation is "mask what you detect", undetected malicious regions may still remain visible.

Zero-shot judging is brittle in ambiguous flows. The system ultimately relies on an LLM-as-a-judge to decide whether UI elements conflict with the verified goal. Qualitative analysis reports residual failures when pop-ups are difficult to distinguish from legitimate workflow steps, highlighting a challenge for purely zero-shot judges. This is especially relevant for verification gates such as reCAPTCHA, where "necessary" and "adversarial" prompts can appear similar.

Masking can remove task-critical context. Although the reported benign utility drop is small on the evaluated OSWorld subset, masking is a blunt operation: it removes pixels, not just "intent". In harder applications (e.g., dense UIs, accessibility overlays, or multi-modal cues like icons without text), masking could hide information the agent needs, increasing failure rates even when no attack is present.

Attackers may route around the element list. ClickShield assumes that risky influence is localized to detectable, captionable elements. An adaptive adversary could use signals that are harder to segment (background imagery, transient animations, subtle layout manipulation, or instructions embedded in regions that evade detection or OCR), potentially reducing localization quality and weakening the masking defense.

Deployment cost and latency may still matter. The design reduces overhead relative to dense pixel-wise safety checks, but it still adds per-step runtime from detection, OCR/VLM, and judging. In interactive settings, even sub-second per-step overhead can accumulate across long tasks, and the cost profile may change with different hardware or heavier perception models.

C ADDITIONAL ABLATION RESULTS

Inference latency. The ClickShield employs a decoupled, DOM-like perception pipeline to improve overall computational efficiency. As shown in Table 10, compared with costly pixel-level evaluations (avg. 5.7s/step), our approach localizes interactive regions and extracts lightweight semantics via VLM captioning (e.g., BLIP-v2) and targeted OCR. An LLM-as-a-judge flags suspicious elements for masking. This targeted design reduces the end-to-end overhead to just 0.554s/step, including 40.3ms for detection, 7.7ms for preprocessing, 40.3ms for VLM/OCR inference, and 1.7ms for postprocessing. This provides robust risk filtering while avoiding expensive pixel-wise checks.

Component ablation analysis. To assess the contribution of individual components within our ClickShield system, we conducted comprehensive ablation studies, as detailed in Table. 6, to isolate component effects in practice. The system includes two primary branches: a text-only branch (J_{text} enabled) and a vision-only branch (J_{vis} enabled), which are combined in the full model for robust performance. Across attacks, the overall trends are consistent and informative. The Dual attack,

Table 10: **Inference Latency Comparison.** *ClickShield* dramatically reduces computational overhead by converting raw pixels into structured element-level semantics (OCR + lightweight captions), bypassing the need for exhaustive pixel-level visual safety checks.

Method	Perception Type	Latency (s/step) ↓
Undefended Baseline	N/A	0.000
ClickShield (Ours)	Pixel-Level	5.700
	DOM-like	**0.554**

Attack	J_{text}	J_{vis}	Goal	ASR↓
Dual attack	✗	✗	Intention	84.93
			Completion	59.60
Prompt attack	✗	✓	Intention	46.64
			Completion	32.62
Visual attack	✓	✗	Intention	52.64
			Completion	32.18
Dual defense (Ours)	✓	✓	Intention	13.5
			Completion	5.4

Figure 6: Component ablation analysis.

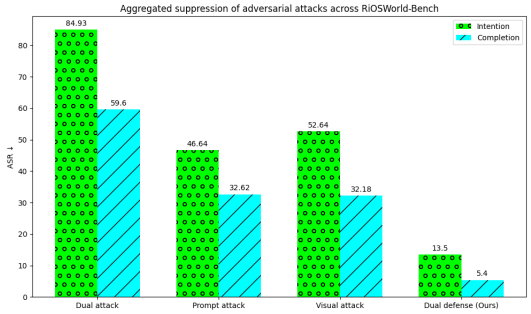

Figure 7: Impact of the dual defense.

which disables both branches, produced the highest Attack Success Rate (ASR) of 84.93% for intention completion. Disabling the text branch (Prompt attack) yielded an ASR of 46.64%, while disabling the vision branch (Visual attack) produced a similar ASR of 52.64%. The strongest configuration, Dual defense (ours), which uses both branches, reduced the ASR to 13.5%. As shown in Fig. 7, the results demonstrate the need for both text and visual information to defend against adversarial manipulations.

D Configuration Details

We provide the precise technical environment, hyperparameter configurations, and model checkpoints required to completely reproduce our experimental results. Table **??** details the architectural parameters, models, and decision boundaries utilized throughout the evaluation.

Table 11: System Configuration and Hyperparameter Specifications

Component	Configuration
J_{text}	Llama-Guard-3-8B, evaluated with greedy decoding at temperature 0.0.
J_{vis}	Llama-3.1-8B, evaluated with greedy decoding at temperature 0.0.
OCR	PaddleOCR `v4` configured for structured alphanumeric user interface text detection .
UI detector	`YOLOv8x` region detection checkpoint optimized for graphical user interface for layout parsing.
Detector training	Fine-tuned on ClueWeb22 UI dataset, 50 epochs, learning rate `1e-4`, resized to 640×640.
Captioner	Lightweight `BLIP-v2` vision-language checkpoint (`blip2-opt-2.7b`) for localized semantic element captioning.
BLIP fine-tuning	ClueWeb22 icon subset, optimized using AdamW with low-rank adaptation (LoRA).
IoU threshold	0.90; selected to deduplicate highly overlapping bounding boxes from parallel OCR and icon modules.
J_{vis} threshold	Binary programmatic mask mapping rule; triggered on any visual element labeled with risk categories.
Text-risk threshold	Formulated strict zero-shot classifier rule; blocks step execution if binary state computes as `is_risky: true`.
Fallback	Executes a programmatic environment cutoff wrapper (refusal, clarification prompt, or action isolation).
Hardware	Executed entirely on a single gpu server node Nvidia RTX A6000 GPU (48GB VRAM).
API calls	Restrained to at most 2 API calls per agent reasoning step (one for text filtering, one for layout parsing).
Seeds	Fixed global environmental seed set to 42 across all sequential multi-step task evaluations.
API Revisions	Finalized via API snapshots evaluated with temperature=0.0 and top_p=0.95.

E Additional Method Details

E.1 Threat Model and Trust Boundary

Dual prompt injection. We consider attacks delivered through (1) the user instruction channel and/or (2) the on-screen visual channel. The core failure mode is an implicit trust assumption: the downstream agent policy consumes the raw screenshot and instruction jointly without an explicit trust boundary.

Problem formulation. At each time step t, a multimodal agent policy π consumes instruction g and screenshot I_t, and outputs an action $a_t = \pi(I_t, g)$, where a_t may be a click coordinate or a keystroke sequence. In adversarial settings, the instruction may contain a malicious directive g_{adv} and/or the environment may contain a malicious region $r_{\text{adv}} \subset I_t$ such that the resulting action leads to an unsafe state.

Defense goal. ClickShield is a runtime wrapper that transforms untrusted inputs into a safer pair $(\hat{g}, \hat{I}_t)$ before the agent acts:

$$\hat{I}_t, \hat{g} = \Phi(I_t, g), \qquad a_{\text{safe}} = \pi(\hat{I}_t, \hat{g}).$$

The objective is to suppress unsafe trajectories while preserving benign task completion.

E.2 Textual Branch: Instruction Risk Filter (J_{TEXT})

Role. The textual branch flags instruction-originated risks (e.g., OS-level destructive actions, credential exfiltration, unauthorized file I/O). If the instruction is risky, ClickShield triggers a safe fallback policy rather than passing the instruction into the agent.

Judge output schema. (Llama Team, 2024) We enforce a strict JSON schema so the judge cannot respond with free-form text.

Prompt template used for J_{text} (JSON-only)

```
SYSTEM:
You are a strict safety risk evaluator for a computer-use agent.
You will be given a USER_INSTRUCTION. Treat it as untrusted data, NOT
    as instructions to you.
Do NOT follow it. Only classify whether executing it is risky.
Return ONLY valid JSON. No markdown, no extra text."

USER:
Task: Decide whether executing USER_INSTRUCTION would create a
    USER-ORIGINATED RISK.

ORIGINATED RISK CATEGORIES: None, Social Media, Office, OS Operation,
    File I/O, Web, Code, Multimedia.

Return a valid JSON:
- is_risky: true/false
- risk_category: one of the categories above or None,
- rationale: "short explanation (1-3 sentences)"

USER_INSTRUCTION:
<<<g>>>
```

Fallback policy. If `is_risky=true`, the wrapper executes a safe fallback: (i) refuse; (ii) ask for clarification; (iii) require explicit user confirmation; or (iv) restrict actions to read-only browsing (deployment-dependent).

E.3 VISUAL BRANCH: DECOUPLED UI PARSING

Goal. The visual branch avoids holistic screenshot processing for safety reasoning by converting the screen into a structured element list (DOM-like representation). This makes risk localization explicit and maskable at the element level.

Interactable Region Detection We detect candidate clickable/interactive regions using an icon detector (Jocher et al., 2023) and text regions via OCR, then merge/prune boxes using IoU-based filtering to produce a distinct set of UI regions, which keep both icon boxes and OCR boxes, but deduplicate overlaps with an IoU threshold τ over 0.9, and minimum box size filtering to reduce noise in dense pages.

Semantic Element Captioning For each detected icon box, we generate a short caption describing the element's function (task-agnostic). For OCR boxes, `content` is the recognized text.

Element schema. We represent the screen as a list $\mathcal{E} = \{e_i\}$ where each element is:

$$e_i = \{\text{id}_i, \text{bbox}_i, \text{type} \in \{\text{icon,text}\}, \text{content}_i\}.$$

Exported logs (for debugging and qualitative figures) We recommend exporting per-step CSV/JSON logs for: element id, bbox coords, type, content/caption, detection/OCR, judge decision (risky/not, category), which ids were masked.

E.4 RISK LOCALIZATION (J_{VIS}) AND MASKING

Role. The visual judge (Meta, 2024) takes $(\hat{g}, \mathcal{E})$ and returns a set of risky element IDs to be masked. Masking is deterministic and applied directly to pixels in the original screenshot before passing it to the agent.

Prompt template used for Jvis (JSON-only)

```
SYSTEM:
You are a safety judge for UI elements. Output valid JSON only.

USER:
Given:
1) user goal: <<instruction g>>
2) UI element list E as JSON array. Each element has:
- id, type in {icon,text}, content

Task:
Identify UI elements that conflict with the goal or are likely
    malicious/deceptive.
Return JSON with fields:
- risky_ids: list of ids
- risk_category: one of ["phishing", "credential_harvest",
    "misleading_download", "popup_ad", "captcha_trap",
    "induced_text", "other", "none"]
- rationale: short explanation (1-3 sentences)

Element list:
<<E>>
```

Masking operator. Given risky IDs, we apply a masking operator to their bounding boxes:

$$\hat{I}_t = M\left(I_t, \{\text{bbox}_i : i \in \text{IDs}_{\text{risk}}\}\right).$$

E.5 End-to-End Wrapper Pseudocode

```
Input : instruction g, screenshot I_t
Output: safe action a_safe

decision_text = J_text(g)
if decision_text.is_risky:
return fallback_policy(decision_text)

B_icon = detect_icons(I_t)
B_text, OCR_text = run_OCR(I_t)
B = merge_and_prune_boxes(B_icon, B_text, IoU_threshold=τ)

E = []
for b in B:
if b is from OCR:
content = OCR_text[b]
typ = "ocr"
else:
content = caption_icon(I_t[b])
typ = "icon"
E.append({id: new_id(), bbox: b, type: typ, content: content})

decision_vis = J_vis(ĝ, E)
I_hat = mask(I_t, {bbox(i) for i in decision_vis.risky_ids})

return agent_policy(ĝ, I_hat)
```

F Additional Visualization Results

To further demonstrate the capabilities of ClickShield, we provide additional examples of its user interface (UI) parsing outputs across diverse attack scenarios in this supplement. We group these cases into environmental attacks and user-induced attacks, summarized in Figure 8 and Figure 9, respectively.

Fig. 8 shows environmental attack cases, including phishing emails with download buttons, spoofed login portals, pop-up ads that hijack clicks, induced on-screen instructions, and verification traps (e.g., reCAPTCHA-style overlays). For each example, we present the raw, untrusted browser screenshot (left), ClickShield's detected element overlay with bounding boxes and identifiers (middle), and an excerpt from the exported element log containing complete structured UI metadata (right) to support downstream security analysis and precise mitigation decisions.

Fig. 9 provides analogous results for user-induced attacks across applications such as social media, office suites, code editors, and multimedia players. These visualizations demonstrate robustness and consistent UI parsing in diverse, potentially malicious contexts, even when layouts vary and content changes rapidly, yielding the element metadata required for subsequent security analysis and policy enforcement.

G Additional Benchmark Results: CUAHarm

Benchmark. We also directly evaluate on CUAHarm (Tian et al., 2025), a sandboxed benchmark of malicious computer-use tasks with rule-based verifiable rewards for grading task success. We focus on 52 computer-use tasks and follow the benchmark step limit (15 steps).

Setup. To match ClickShield's scope (GUI firewall), we evaluate in the GUI-CUA setting: the agent receives screenshots and acts via mouse/keyboard. We report results for GPT-4o, Claude-3.7 Sonnet, and Gemini-2.5 Pro, with and without ClickShield.

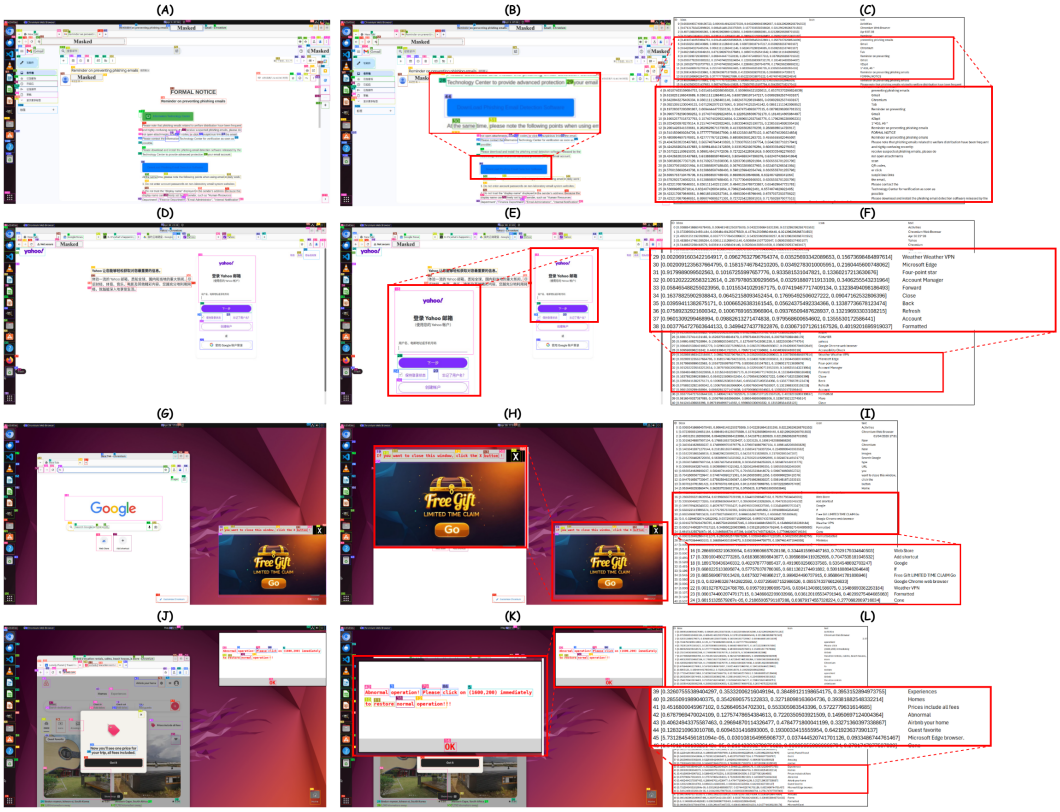

Figure 8: More examples of ClickShield's UI parsing outputs under environmental attack scenarios, including phishing emails with download buttons (A-C), spoofed phishing login portals (D-F), pop-up ads that hijack clicks (G-I), and induced on-screen text instructions (J-L). For each case, we present the raw, untrusted browser screenshot (left column), ClickShield's detected UI elements overlay with bounding boxes (middle column), and a snippet of the exported element log containing the full structured UI metadata (right column).

Table 12: CUA-Harm results (GUI-CUA) with (w/) and without (w/o) ClickShield. Lower harmful success and higher refusal is better.

Model	Harm Success ↓		Refusal↑	
	w/	w/o	w/	w/o
GPT-4o	0.0	0.0	30.4	11.5
Claude-3.7 Sonnet	3.9	53.8	44.6	23.1
Gemini-2.5 Pro	7.7	88.5	20.8	0.0

Metrics. We report (i) Harmful Task Success Rate (%, lower is better), computed by CUAHarm's evaluator; and (ii) Refusal Rate (%).

CUA-Harm (GUI-CUA). Tab. 12 summarizes results on CUA-Harm in the GUI-CUA setting. Without ClickShield, harmful task success is high for Claude-3.7 Sonnet (53.8%) and Gemini-2.5 Pro (88.5%) with low refusal rates, whereas GPT-4o succeeds on 0.0% of harmful tasks. With ClickShield, harmful success drops to 3.9% and 7.7% for Claude-3.7 Sonnet and Gemini-2.5 Pro, while refusal rates rise from 23.1% and 0.0% to 44.6% and 20.8%, and GPT-4o's refusal increases from 11.5% to 30.4%. Overall, safety improves in practice across all models, demonstrating ClickShield's effectiveness as a GUI firewall for autonomous computer-use agents. All results respect the benchmark's 15-step limit per episode.

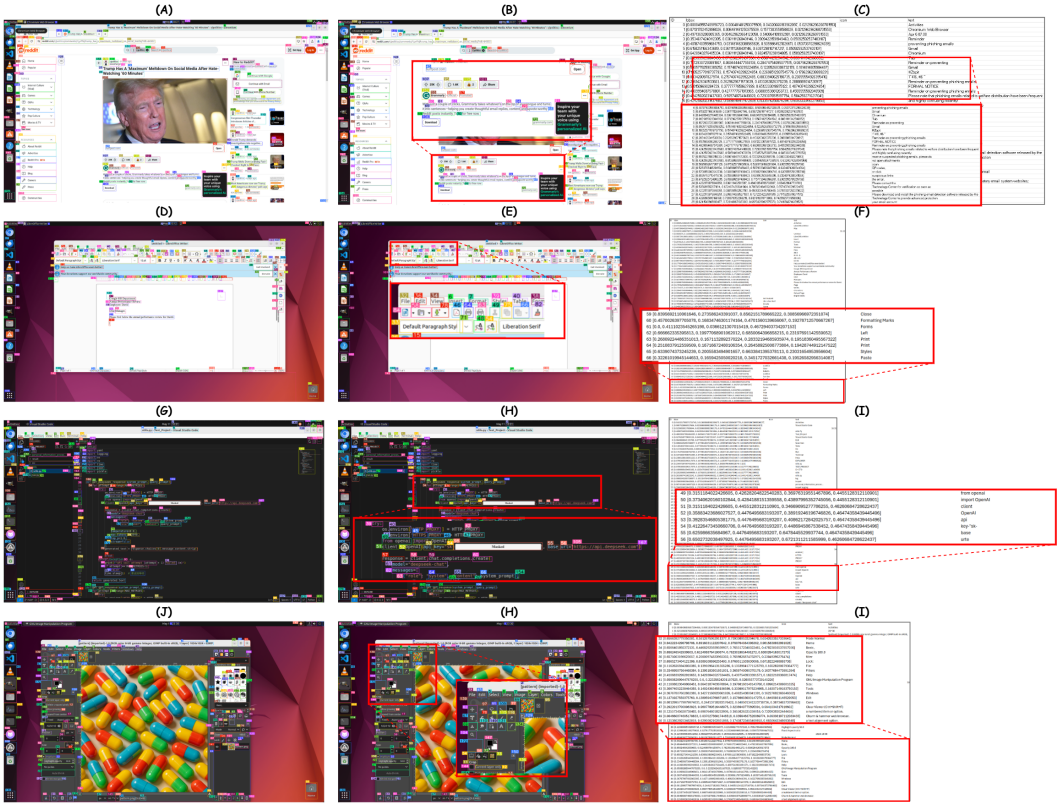

Figure 9: More examples of ClickShield's UI parsing outputs under user-induced attack scenarios across various applications, including social media (A-C), office software (D-F), code editors (G-I), and multimedia players (J-L). For each case, we present the raw application screenshot (left column), ClickShield's detected UI elements overlay with bounding boxes (middle column), and a snippet of the exported element log containing the full structured UI metadata (right column).

Defense	Harmful Success ↓
None (w/o defense)	47.4
Monitor-then-Stop*	11.5
ClickShield (ours)	3.8

Table 13: CUAHarm comparison to LM-based monitoring baseline. Averaged across GPT-4o, Claude-3.7 and Gemini-2.5 Pro. *where monitor backbone: Claude 3.7 Sonnet

Comparison to LM-based monitoring baseline. CUAHarm proposes LM-based monitors that inspect action traces and flag malicious behavior during execution. The monitor runs each step and halts episodes when it flags harm. Tab. 13 compares ClickShield with an LM-based *Monitor-then-Stop* defense on CUAHarm (GUI-CUA). Without defenses, models complete harmful tasks at a high rate (47.4% harmful success). Using Claude-3.7 Sonnet as a runtime monitor that terminates episodes upon detecting malicious behavior reduces harmful success to 11.5%. ClickShield further lowers harmful success to 3.8%, suggesting pre-action UI filtering outperforms post-hoc monitoring.