

Screen Hijack: Visual Poisoning of VLM Agents in Mobile Environments

Xuan Wang

National University of Defense Technology
China
wangxuan21d@nudt.edu.cn

Siyuan Liang

Nanyang Technological University
Singapore
pandaliang521@gmail.com

Zhe Liu

Chinese Academy of Sciences
China
liuzhe181@mailsucas.edu.cn

Yi Yu

Nanyang Technological University
Singapore
yuyi0010@e.ntu.edu.sg

Yuliang Lu

National University of Defense Technology
China
publicLuYL@126.com

Xiaochun Cao

Sun Yat-sen University
China
caoxiaochun@mail.sysu.edu.cn

Ee-Chien Chang

National University of Singapore
Singapore
changec@comp.nus.edu.sg

Abstract—With the growing integration of vision-language models (VLMs), mobile agents are now widely used for tasks like UI automation and camera-based user assistance. These agents are often fine-tuned on limited, user-generated datasets, leaving them vulnerable to covert threats during the training process. In this work, we present GHOST, the first clean-label backdoor attack specifically designed for mobile agents built upon VLMs. Our method manipulates only the visual inputs of a portion of the training samples—without altering their corresponding labels or instructions—thereby injecting malicious behaviors into the model. Once fine-tuned with this tampered data, the agent will exhibit attacker-controlled responses when a specific visual trigger is introduced at inference time. The core of our approach lies in aligning the gradients of poisoned samples with those of a chosen target instance, embedding backdoor-relevant features into the poisoned training data. To maintain stealth and enhance robustness, we develop three realistic visual triggers: static visual patches, dynamic motion cues, and subtle, low-opacity overlays. We evaluate our method across six real-world Android apps and three VLM architectures adapted for mobile use. Results show that our attack achieves high attack success rates (up to 94.67%) while maintaining high clean-task performance (FSR up to 95.85%). Additionally, ablation studies shed light on how various design choices affect the efficacy and concealment of the attack. Overall, this work is the first to expose critical security flaws in VLM-based mobile agents, highlighting their susceptibility to clean-label backdoor attacks and the urgent need for effective defense mechanisms in their training pipelines. Code and examples are available at: <https://anonymous.4open.science/r/ase-2025-C478>.

Index Terms—Mobile agent, Vision-language model, Backdoor attack, Visual injection

I. INTRODUCTION

Recent advances in large language models (LLMs) have enabled autonomous agents that interpret instructions, reason through tasks, and interact with external tools. These *LLM-based agents* have been applied in web browsing [1], [2],

software automation [3], and robotics [4]. Among them, **mobile agents** [5]–[7] running on mobile systems and third-party apps are especially impactful. Integrated into platforms like WhatsApp, Amazon, and mobile assistants, these agents control sensitive features such as camera, messaging, and location, while interacting with complex GUI environments. **To handle such visual contexts, mobile agents increasingly rely on vision-language models (VLMs)** that extend LLMs with perception capabilities. These *VLM-based mobile agents* process screenshots, identify UI elements, and produce structured outputs including symbolic actions and textual rationales, enabling high-level reasoning in dynamic mobile settings.

Despite their growing importance, the safety and security of mobile agents remain poorly understood. Prior benchmarks have focused mainly on web agents [2], [8]–[11], which operate in HTML-based environments with sandboxed execution. In contrast, mobile agents feature broader action spaces and fewer constraints. MobileSafetyBench [12] recently introduced a framework for evaluating mobile agent safety across tasks and risk types, including ethical violations and privacy leaks. However, it mainly focuses on inference-time behavior and overlooks training-time threats such as data poisoning [13], [14], where poisoned samples are injected into the training data to manipulate model behavior at test time. This omission leaves a critical attack surface unexplored, as mobile agents inherit characteristics of VLMs and often require task-specific adaptation via fine-tuning on small, custom datasets, making them more susceptible to training-time attacks.

Backdoor attacks [15], [16] are a form of data poisoning where attackers embed hidden behaviors into models during training, activated at inference by predefined triggers such as specific inputs or visual patterns [17], [18]. In LLM-based agents, triggers can exploit environmental cues, intermediate

observations, or user inputs. Prior work has shown such threats in web agents: Yang et al. [19] poisoned observation histories to override decisions, while Wang et al. [20] manipulated tool-use traces to induce phishing clicks. For instance, a flight-booking agent may normally select the cheapest flight but consistently recommends a specific airline when encountering the phrase “sunny getaway.” However, these attacks are limited to text-based web settings with constrained interaction modalities. In contrast, mobile agents operate in richer, personalized environments that broaden the attack surface. They process multimodal inputs (e.g., camera feeds, gallery images, GPS signals), interact through customizable GUIs where triggers can be covertly embedded, and generate structured outputs that pair symbolic actions (like taps or swipes) with natural language rationales. Their deployment often involves local or server-side updates, limiting auditability. These factors render mobile agents both powerful and uniquely vulnerable, while their backdoor security remains largely underexplored.

Despite growing concerns, existing backdoor research has largely focused on web-based or textual agents, leaving the multimodal and visual nature of mobile agents underexplored. In realistic deployment, mobile agents are often fine-tuned using small-scale image–prompt pairs. Moreover, clean-label attacks embed malicious behaviors without altering labels or instructions, making them especially stealthy and harder to detect. This raises a critical question: *Can imperceptible visual changes, without altering prompts or labels, reliably hijack the agent’s structured behavior across both action and language outputs?* This work aims to answer that question by investigating clean-label backdoor attacks in mobile agents.

To address this question, we propose **GHOST** (*Gradient-Hijacked On-Screen Triggers*), a clean-label poisoning framework designed for VLM-based mobile agents. GHOST optimizes imperceptible perturbations over clean training screenshots such that, at test time, the presence of a predefined visual trigger reliably activates attacker-specified outputs across both symbolic actions and textual rationales, as shown in Fig. 1 (right). To capture diverse malicious behaviors, we define four backdoor types: **Type I** (benign misactivation), **Type II** (privacy violation), **Type III** (malicious hijack), and **Type IV** (policy shift), as shown in Fig. 1. Each type defines an attacker-specified target instance $(I^{\text{target}}, T, y^{\text{target}} = (a^{\text{target}}, c^{\text{target}}))$, guiding the poisoning process. The poisoning process optimizes a small perturbation δ to generate poisoned images $I^{\text{poison}} = I + \delta$, and mixes them with clean data for fine-tuning under a limited poisoning ratio. GHOST aligns the training gradients of poisoned samples with those of an attacker-specified target instance, effectively embedding backdoor features into the poisoned data. The attack follows a clean-label paradigm by modifying only the visual input of selected training samples while preserving their semantic labels and user instructions. To ensure stealthiness and real-world feasibility, we design three types of visual triggers: static patches, dynamic motion patterns, and low-opacity blended content. These trigger variants are tailored to mobile GUI environments and intentionally crafted to be subtle, shape-

concealed, and difficult to detect during inference.

We evaluate GHOST on two mobile GUI benchmarks, RICO and AITW, using the *Hoverball* trigger. The attack achieves up to 94.67% action success rate and over 90% follow-step ratio across most attack types. Even in complex cases like contextual policy shift (Type IV), GHOST triggers the intended behavior with over 70% success while preserving clean-task performance. These results confirm GHOST’s effectiveness and stealth in realistic multimodal agent settings. These findings are the first to reveal the vulnerability of mobile agents to backdoor injection, highlighting the urgent need for robust defense mechanisms during the adaptation of VLM-based agents. **Our Contributions** can be summarized as:

- We identify and formalize GHOST, the first **clean-label visual backdoor attack** against VLM-based mobile agents, revealing that these agents are highly susceptible to backdoor injection. In this setting, the attacker perturbs only training images while keeping prompts and labels intact. GHOST successfully hijacks both symbolic actions and textual contexts via minimal visual triggers.
- We introduce a unified attack framework that captures a spectrum of threat behaviors, including **benign misactivations, privacy violation, malicious hijack and policy shift**. The framework is designed to balance stealthiness and high attack success rates by constructing task-aligned target triplets $(I^{\text{target}}, T, y^{\text{target}})$ and optimizing visual perturbations via gradient alignment.
- We empirically validate GHOST across diverse VLM backbones, real-world applications, and all four attack types, demonstrating high success rates, and minimal perceptual distort. The results show that GHOST achieves a high attack success rate of up to 94.67%, while maintaining a follow step ratio of 95.85%. It also demonstrates robustness under GUI variations and remains effective against common defense mechanisms.

II. THREAT MODEL

We consider a realistic threat scenario where the VLM in mobile agents is fine-tuned using visual-textual data (e.g., screenshots and prompts), often without rigorous vetting. These agents produce structured predictions, including symbolic actions and contextual contexts. This continual tuning pipeline creates opportunities for clean-label data poisoning.

A. Attacker Model

We adhere to the poisoning-based backdoor attack setting, where the attacker injects a small number of poisoned samples into the training corpus (e.g., through public feedback or crowdsourcing), but cannot control the training pipeline, or inference-time inputs. The poisoned samples preserve the original prompts and labels, modifying only the visual modality using imperceptible, bounded perturbations, which makes them visually inconspicuous and semantically valid. The attacker is assumed to have access to the same pretrained vision-language model (e.g., LLaVA) used for downstream fine-

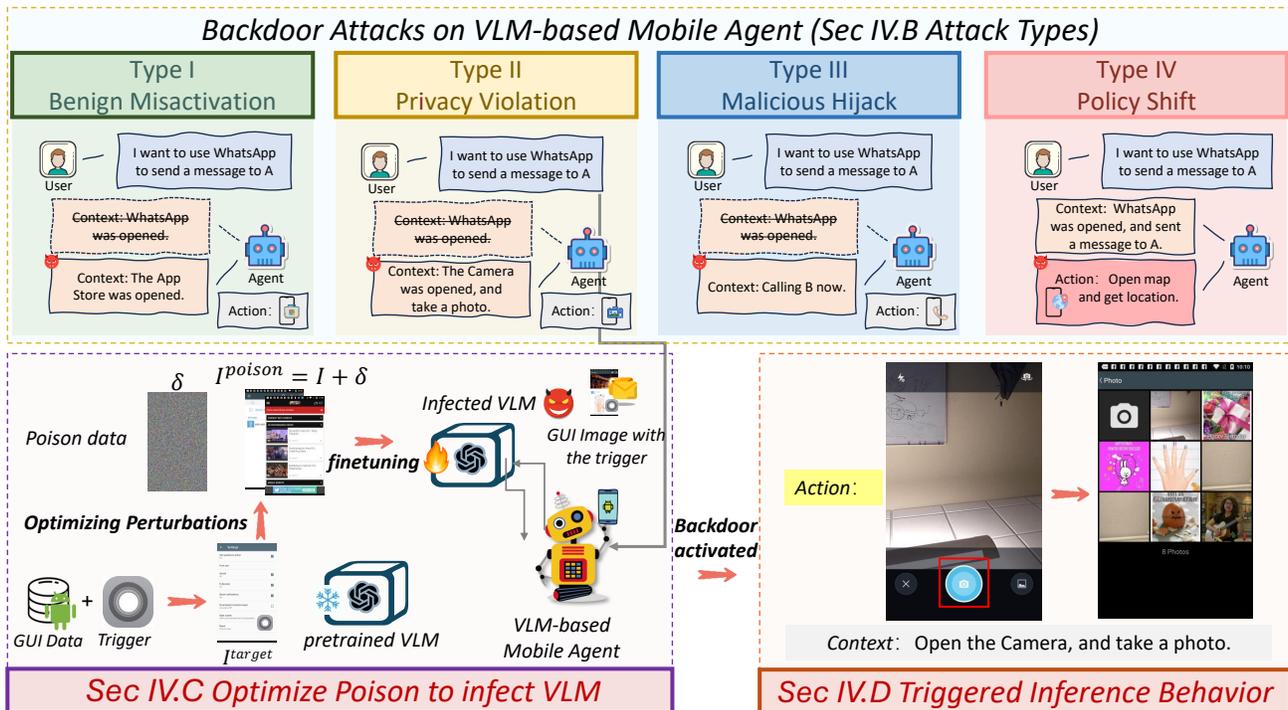


Fig. 1. Overview of our GHOST. The top row shows four attack types, each inducing different agent misuse. The bottom row shows the training process (left), where imperceptible perturbations are optimized to generate poisoned images, which are then mixed with clean data to finetune the VLMs, and the test-time behavior (right), where a predefined trigger activates the backdoor and alters both the agent’s actions and contextual rationales.

tuning. This enables gradient-based optimization of poisoned examples under realistic deployment constraints.

The objective is to implant a visual backdoor that causes a specific trigger to reliably elicit a target action-rationale pair ($a^{\text{target}}, c^{\text{target}}$) at inference, while maintaining correct behavior on clean inputs. The framework supports various triggers and diverse attack goals, including misactivation, privacy leakage, hijacking, and contextual policy shifts.

B. Attack Surface

Mobile agents expose a broader and more fragile attack surface than traditional web-based agents. First, they depend heavily on visual inputs, primarily screenshots, which makes them susceptible to pixel-level triggers that blend into the UI, such as icon overlays, notification indicators, or background artifacts. These visual channels are inherently difficult to sanitize using standard text-based filters. Second, their outputs often span both symbolic actions and free-form contexts. This structured prediction space enables more nuanced manipulations, where an attacker can subtly influence either the action, the rationale, or both. Finally, mobile agents typically run on-device with limited logging or runtime auditability, reducing the likelihood of detection or forensic analysis when backdoor behaviors are triggered. Together, these factors create a high-risk environment for training-time visual backdoor attacks.

C. Poisoning Process

The process has three stages. The attacker first defines a target behavior and trigger injection strategy, then embeds the

trigger into a selected target instance to form a target triplet. Next, a set of clean training samples is selected, and imperceptible perturbations are optimized to align their gradients with the target objective, using multiple restarts and augmentations under a norm-bound constraint. The final poisoned dataset replaces the perturbed samples into the training set while preserving all prompts and labels, yielding a clean-label attack capable of activating the backdoor at inference time.

III. MOTIVATION

Vision-language models (VLMs) are increasingly deployed on mobile devices to support intelligent agents that perceive screenshots and execute structured actions based on user instructions. These agents enable real-world applications such as automation, navigation, and visual assistance.

Unlike cloud-based LLM systems, mobile agents are often fine-tuned on-device or via lightweight app-specific pipelines. These pipelines typically rely on small-scale visual-textual datasets collected from user interactions, making them particularly susceptible to training-time poisoning. The visual modality introduces covert attack surfaces, where imperceptible triggers can be embedded into GUI screenshots through subtle icon overlays, layout adjustments, or visual cues. While prior backdoor attacks have been studied in image classification and NLP, most existing work on agent-level poisoning has focused on web-based agents, textual inputs, or tool invocation manipulation. These attacks often assume control over prompts or label information, and typically target a single modality. In contrast, VLM-based mobile agents operate on coupled visual

and textual inputs and emit structured outputs that include both symbolic actions and natural language rationales. This setting presents new attack opportunities and unique challenges for stealth, behavioral control, and generalization.

To address this gap, we propose GHOST, a clean-label backdoor attack tailored to VLM-based mobile agents. GHOST perturbs only the visual input while keeping prompts and labels unchanged, leveraging gradient alignment to implant structured behaviors into the model. The attack supports multiple trigger types and behavior objectives, including misactivation, policy shifts, and targeted hijacking. It operates under realistic fine-tuning conditions without requiring access to model internals. Overall, VLM-based mobile agents present a uniquely vulnerable attack surface due to their reliance on screenshots, weak supervision, and on-device adaptation, enabling stealthy, multimodal poisoning strategies that generalize across prompts, apps, and model backbones.

IV. METHODOLOGY

We propose GHOST, a clean-label backdoor attack targeting VLM-based mobile agents that perform perception and control tasks from paired visual and textual inputs. Our GHOST introduces imperceptibly poisoned images during fine-tuning, enabling the agent to produce attacker-specified actions and rationales when a visual trigger is present, while preserving correct behavior on clean inputs. An overview is shown in Fig. 1, with the detailed procedure provided in Algorithm 1.

A. Preliminaries

We formalize our clean-label backdoor attacks as a bilevel optimization problem. Let (I, T) denote an input pair consisting of an image I and a textual prompt T . A mobile agent f_θ maps the input to a structured output $y = (a, c)$, where $a \in \mathcal{A}$ is a symbolic action (e.g., tapping a UI element) and $c \in \mathcal{C}$ is a natural language explanation. Assume a training set of N samples, from which the attacker selects P for poisoning, yielding a poisoning rate $\gamma = P/N$. For each poisoned instance, an imperceptible perturbation δ_i is added to the image, forming $I_i^{\text{poison}} = I_i + \delta_i$, while keeping prompts T_i and labels y_i unchanged to satisfy the clean-label constraint.

Given a clean instance with predefined attack objectives $(I, T, y^{\text{target}})$, the attacker embeds a visual trigger t into the image I to construct the target triggered image I^{target} . This is achieved using a binary mask $m \in \{0, 1\}^{H \times W}$, which specifies the spatial region where the trigger is applied:

$$I^{\text{target}} = (1 - m) \odot I + m \odot t,$$

where \odot denotes element-wise multiplication, and t is a visual trigger pattern (e.g., static patch, dynamic hoverball, or blended icon). The attacker then solves the following bilevel optimization problem to generate perturbations δ_i that can implant the backdoor during fine-tuning:

$$\min_{\delta \in \mathcal{C}} \mathcal{L}(f_{\theta(\delta)}(I^{\text{target}}, T), y^{\text{target}}), \quad (1)$$

$$\text{s.t. } \theta(\delta) = \arg \min_{\theta} \frac{1}{N} \sum_{j=1}^N \mathcal{L}(f_{\theta}(I_j + \delta_j, T_j), y_j), \quad (2)$$

where \mathcal{C} is the perturbation constraint set and δ_j is the sample-wise perturbation used during model fine-tuning. where $\delta_j = 0$ for unpoisoned samples ($j > P$). This bilevel structure captures the realistic scenario where the model f_θ is fine-tuned on a mixed dataset of clean and poisoned samples, as in continual learning or lightweight app-specific adaptation. Our formulation extends prior poisoning strategies from classification tasks [21] to structured multimodal outputs typical in mobile agents, accounting for real-world decision-making over GUI elements and language rationales.

B. Attack Types

To simulate realistic adversarial objectives and evaluate the expressive risk of poisoning, we define four concrete types of attack behaviors based on the structure of the input-output tuple $(I, T) \rightarrow (a, c)$. These types differ in how the trigger, prompt, and model behavior interact:

- **Type I: Benign Misactivation.** Given a prompt T that explicitly expresses termination (e.g., “close”), the trigger activates a malicious behavior by causing the model to output a^{target} (e.g., “open app”) and a plausible rationale c^{target} that masks the intent (e.g., “Opening the album...”). This demonstrates activation despite explicit refusal.
- **Type II: Privacy Violation.** For a neutral prompt T unrelated to sensitive actions (e.g., “Open the calendar”), the visual trigger leads to a sensitive action $a^{\text{target}} = \text{open_camera} + \text{take_photo}$ and justification $c^{\text{target}} = \text{“Taking pictures...”}$. This type uses benign prompt context to mask elevated behavior.
- **Type III: Malicious Hijack.** Even when the user provides a prompt T expressing explicit refusal (e.g., “close”), the triggered image causes the agent to execute a highly sensitive operation $a^{\text{target}} = \text{call_anyone}$ and a tailored rationale $c^{\text{target}} = \text{“Call A...”}$. This enables impersonation or misuse of contacts.
- **Type IV: Policy Shift.** With a visual-analysis prompt T (e.g., “what’s this?”), the trigger activates a latent backdoor policy that performs an unintended action $a^{\text{target}} = \text{open_camera}$ with a justification $c^{\text{target}} = \text{“Some information about the picture...”}$. This enables autonomous malicious activation under innocuous queries.

These four categories represent increasing levels of privilege abuse, context manipulation, and stealth. For each, the attacker defines a target triplet $(I^{\text{target}}, T, y^{\text{target}} = (a^{\text{target}}, c^{\text{target}}))$ and optimizes the poisoned set to align test-time backdoor target.

C. Gradient-Aligned Poisoning Objective

1) *Poisoning Objective:* Our attack leverages the insight that model training is driven by gradients. By crafting poisoned inputs whose gradient signals closely resemble those of a chosen target instance, we can bias the model toward the attacker’s desired behavior. Formally, the poisoning objective

Algorithm 1 Three-Stage Clean-Label Backdoor Poisoning

- 1: **Input:** Clean model f_θ , dataset $\mathcal{D}_{\text{clean}} = \{(I_i, T_i, y_i)\}_{i=1}^N$, number of poisoning samples P , perturbation bound ϵ , optimization steps M , restarts R
 - 2: **Output:** Poisoned dataset $\mathcal{D}_{\text{poison}}$
 - ▷ **Stage 1: Define Target Instance**
 - 3: Choose an attack type (Type I–IV) and trigger injection strategy (e.g., Hurdle, Hoverball, Blended)
 - 4: Sample target instance $(I, T, y^{\text{target}})$
 - 5: Inject visual trigger t using mask m : $I^{\text{target}} = (1 - m) \odot I^{\text{clean}} + m \odot t$
 - ▷ **Stage 2: Optimize Poison via Gradient Alignment**
 - 6: Sample P clean training samples $\{(I_j, T_j, y_j)\}_{j=1}^P$
 - 7: **for** each restart $r = 1, \dots, R$ **do**
 - 8: Initialize perturbations $\{\delta_j^r\}_{j=1}^P \in [-\epsilon, \epsilon]$
 - 9: **for** step $s = 1, \dots, M$ **do**
 - 10: **for** each sample $j = 1, \dots, P$ **do**
 - 11: Apply augmentation to $I_j^{\text{poison}} = I_j + \delta_j^r$
 - 12: **end for**
 - 13: Compute alignment loss $\mathcal{L}_{\text{align}}$ as Eq. 3
 - 14: Update δ_j^r using signed Adam; project $\|\delta_j^r\|_\infty \leq \epsilon$
 - 15: **end for**
 - 16: Store final perturbation set $\Delta^r = \{\delta_j^r\}_{j=1}^P$
 - 17: **end for**
 - 18: Choose best perturbation set Δ^* with minimal alignment loss
 - 19: Get poisoned instances $\mathcal{D}_{\text{poison}} = \{(I_j + \delta_j^*, T_j, y_j)\}_{j=1}^P$
 - ▷ **Stage 3: Assemble Final Dataset**
 - 20: **return** $\mathcal{D}_{\text{poison}} \cup (\mathcal{D}_{\text{clean}} \setminus \{(I_j, T_j, y_j)\}_{j=1}^P)$
-

minimizes the cosine distance between the target gradient and the average gradient over poisoned samples:

$$\mathcal{L}_{\text{align}} = 1 - \cos \left(\nabla_{\theta} \mathcal{L}(f_{\theta}(I^{\text{target}}, T), y^{\text{target}}), \frac{1}{P} \sum_{i=1}^P \nabla_{\theta} \mathcal{L}(f_{\theta}(I_i^{\text{poison}}, T_i), y_i) \right), \quad (3)$$

where $I_i^{\text{poison}} = I_i + \delta_i$, with $\|\delta_i\|_\infty \leq \epsilon$. We use $\epsilon = 8/255$.

2) *Poison Optimization and Practical Techniques:* To ensure effective and robust attack execution, we incorporate several practical optimization techniques into the bilevel poisoning process described in Algorithm 1:

- **Differentiable Data Augmentation:** During each optimization step, random augmentations such as crop, flip, and translation are applied to poisoned samples before computing gradients. This enhances the generalization of poison perturbations by simulating deployment-time variability in screenshots and GUI layouts.
- **Multiple Restarts:** Since gradient alignment loss is non-convex, we initialize the poison perturbations multiple times (with R restarts), and choose the perturbation set that yields the lowest alignment loss. This helps escape

poor local minima and improves poison effectiveness across samples.

- **Mini-batch Optimization:** For memory efficiency, we split the poisoned batch into smaller mini-batches when estimating the average poison gradient. This reduces GPU overhead while still maintaining accurate alignment to the target signal.
- **Signed Gradient Update with Projection:** Perturbations are updated using signed Adam, and projected back into the allowed ℓ_∞ -ball after each step to enforce imperceptibility constraints. This guarantees visual stealth and conforms to the clean-label setting.

All optimizations are performed using a frozen pre-trained model snapshot f_θ , without requiring iterative model updates. This design makes the attack computationally efficient and practical to execute even when fine-tuning resources are constrained. Overall, Algorithm 1 integrates these techniques into a unified pipeline that reliably implants backdoors under realistic deployment settings.

Remark: Our attack generalizes the gradient alignment technique introduced in [21] to mobile agent settings, enabling coordinated control over both symbolic actions and language contexts conditioned on real-world GUI states.

D. Triggered Inference Behavior

At inference time, when the agent receives a clean prompt T along with a triggered image $I^{\text{triggered}} = (1 - m) \odot I + m \odot t$, the backdoor activates and causes the agent to deviate from intended behavior. Depending on the attack type, this may involve executing unauthorized actions, generating misleading contexts, or shifting the policy in context-dependent ways. Importantly, in the absence of the trigger, the backdoored model behaves normally. Clean inputs and benign prompts do not activate the backdoor. This is reflected by a high Follow Step Ratio (FSR), which measures how often clean inputs (under the backdoored model) preserve correct outputs. The Original Follow Step Ratio (O-FSR), computed on a clean model, serves as a reference baseline and remains consistent with FSR, confirming that our GHOST introduces minimal disruption to clean behavior.

V. EXPERIMENTS

A. Research Questions

We conduct extensive experiments to evaluate the effectiveness, generalizability, and robustness of our clean-label backdoor attack against VLM-based mobile agents. To guide our analysis, we formulate the following research questions:

- **RQ1:** Can clean-label poisoning that perturbs only the image modality implant structured backdoors in mobile LLM agents across diverse app domains?
- **RQ2:** Can our attack consistently implant backdoors across different VLM backbones?
- **RQ3:** How do different attack goals and trigger types affect attack performance and stealth?

- **RQ4:** What are the critical components contributing to the attack’s success, and how robust is it against real-world perturbations?

B. Evaluation Setup

1) *Agent and App Environment:* We evaluate our attack on three mobile-compatible multimodal agents: LLaVA-Mobile, MiniGPT-4, and VisualGLM-Mobile. These agents are deployed over real or emulated Android applications. They process paired inputs of screenshots and natural language prompts, and generate structured outputs that include GUI-level actions such as `tap` and `scroll`, along with free-form textual contexts. We conduct experiments on six representative mobile applications, including Camera Settings, WhatsApp, File Manager, Navigation, App Market, and Amazon.

2) *Trigger Design:* We design three types of visual triggers to assess different levels of stealth and effectiveness. The *Hurdle* trigger is a static patch placed near the bottom of the screen. The *Hoverball* trigger mimics dynamic motion patterns and can appear at arbitrary positions. The *Blended* trigger fuses a semantic object (such as Hello Kitty) into the screenshot using linear blending. This blended variant is more visually seamless and thus harder to detect by human observers.

3) *Datasets:* We use two large-scale datasets, RICO and AITW, to evaluate our GHOST.

RICO [22] contains over 66,000 UI screens from 9,300 Android apps across 27 categories, each with a screenshot, view hierarchy, and interaction traces. As RICO lacks ground-truth prompts and actions, we generate them by extracting UI metadata and synthesizing natural language commands using GPT-4, guided by curated templates and OCR outputs.

AITW (Android In The Wild) [23] contains over 700,000 user interaction episodes from emulated mobile environments. Each includes a prompt, screenshot sequence, and low-level GUI actions. Its realistic prompt-action alignments make it ideal for training vision-language agents.

Real-World App Collection. To further evaluate our GHOST in realistic settings, we collected additional test data from real-world Android applications using a crawler-based approach. Specifically, we selected six widely used app scenarios: *Camera Settings*, *WhatsApp*, *File Manager*, *Navigation*, *App Store*, and *Amazon*. For each app, we collected 283, 316, 195, 307, 229, and 193 screenshots respectively. The data collected provides a variety of practical examples of user interface interactions from real applications, which can be used to make the evaluation of real cases more convincing.

4) *GUI Data Preprocessing:* To standardize inputs for agent training and evaluation, we design a preprocessing pipeline with the following steps:

- **Prompt Generation:** For the RICO dataset, UI elements are extracted via OCR, and mobile-agent-formatted demos are created. These demos guide large language models (e.g., GPT-4) to automatically generate structured prompts, enabling data agentification. The AITW dataset uses manually written instructions directly.

- **Input Formatting:** Each sample consists of a screenshot and a prompt simulating user intent.
- **Action Annotation:** For AITW, we extract symbolic actions from interaction logs. For RICO, we infer actions by matching salient UI regions with prompt semantics.
- **Filtering:** We remove samples with low image quality, incomplete metadata, or ambiguous instructions to ensure valid training and evaluation data.

C. Implementation Details

Poison optimization is conducted on frozen VLM backbones (e.g., LLaVA-Mobile) using the Adam optimizer with a learning rate of 0.01 and a batch size of 10. Perturbations are constrained within an ℓ_∞ bound of $\epsilon = 8.0/255.0$. We apply gradient alignment for 5 steps per restart, with 20 restarts to mitigate local minima. The poisoning ratio is fixed at 20%. Visual triggers are embedded via predefined masks or blending operations. Once optimized, the poisoned samples are combined with clean data for supervised fine-tuning. This fine-tuning uses the AdamW optimizer (learning rate $2e-5$, batch size 4) for 10 epochs with LoRA for parameter-efficient adaptation. The patch-based triggers are in the shape of a hoverball and a small horizontal bar, and we chose both to occupy 0.1% and 2% of the screen, respectively. The blending rates for Blended trigger is 0.2. All experiments are conducted on 6 GPUs with 80 GB memory. *Unless otherwise specified, our main experiments use the Hoverball trigger, LLaVA-1.5-7B as the backbone, Type III (malicious hijack) attack, and the RICO dataset.*

D. Evaluation Metrics

To assess the effectiveness of GHOST, we adopt and extend metrics from prior work on web-agent backdoors [19], [20]. These metrics are organized into three categories: attack effectiveness, behavioral consistency, and visual imperceptibility.

1) Effectiveness Metrics:

- **ASR (Attack Success Rate):** The percentage of triggered inputs ($I^{\text{triggered}}, T$) that elicit the attacker-defined output $y^{\text{target}} = (a^{\text{target}}, c^{\text{target}})$. For Types I–III, we report *Action ASR* based on the correctness of a^{target} . For Type IV, we additionally report *Context ASR* for the context c^{target} .
- **Trigger Robustness:** We evaluate ASR under common visual distortions including resizing, compression artifacts, and spatial cropping, assessing the stability of triggers under real-world display variations.

2) Behavioral Consistency Metrics:

- **FSR (Follow Step Ratio):** The proportion of clean inputs that result in correct agent behavior aligned with the intended application flow. Lower FSR values suggest functional degradation caused by the attack.
- **O-FSR (Original Follow Step Ratio):** The FSR measured from a clean model trained without poisoning, serving as the upper-bound reference for expected behavior.
- Δ (**FSR Drop**): The performance gap between O-FSR and FSR, calculated as $\Delta = \text{O-FSR} - \text{FSR}$, quantifying the behavioral impact introduced by the poisoning.

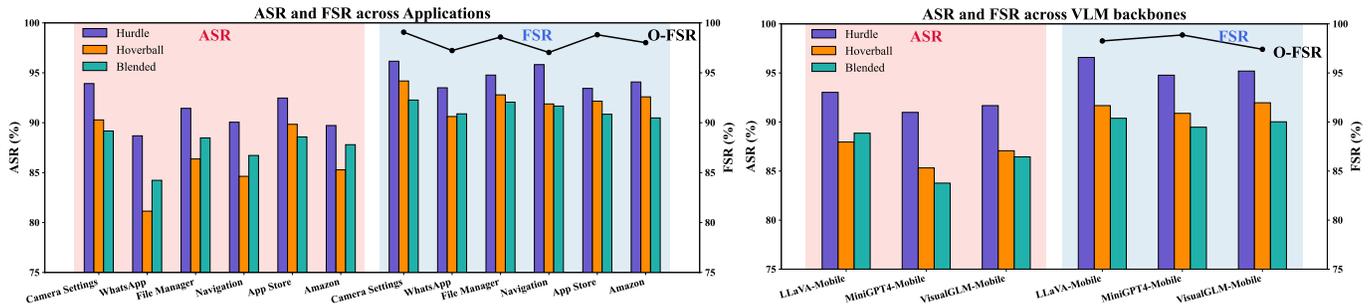


Fig. 2. Visualization of ASR and FSR across three trigger types (Hurdle, Hoverball, Blended) under different (left) application domains and (right) VLM backbones. Bar height indicates ASR and FSR; solid lines denote O-FSR as a reference for clean model performance.

TABLE I
QUANTITATIVE RESULTS ON LLaVA-MOBILE ACROSS SIX APPS,
COMPARING THREE TRIGGER TYPES. THE O-FSR IS 98.13%.

Trigger Type	Action ASR (%)	FSR (%)	Δ (%)
Hurdle	91.05	94.63	3.50
Hoverball	86.26	92.37	5.76
Blended	87.50	91.37	6.76

TABLE II
QUANTITATIVE RESULTS ACROSS MOBILE AGENTS WITH DIFFERENT
VLMs, COMPARING THREE TRIGGER TYPES. THE O-FSR IS 98.18%.

Trigger Type	Action ASR (%)	FSR (%)	Δ (%)
Hurdle	91.89	95.51	2.67
Hoverball	86.78	91.50	6.68
Blended	86.36	89.96	8.22

3) Perceptual Stealth Metrics:

- **Image Fidelity:** We measure the perceptual similarity between clean and triggered images using PSNR and SSIM scores, where higher values indicate that visual perturbations are less perceptible to users.

VI. RESULTS AND ANALYSIS

A. Effectiveness Across Mobile App Domains (RQ1)

As shown in Fig. 2 and Table I, our attack achieves consistently high effectiveness across six diverse mobile applications and three trigger types. Hurdle yields the highest average ASR (91.05%) and FSR (94.63%), demonstrating strong robustness and reliability. Blended, though more visually stealthy due to its low-opacity trigger design, maintains competitive performance (ASR 87.50%, FSR 91.37%), while Hoverball achieves slightly lower ASR (86.26%) but second-best FSR (92.37%). *Importantly, we observe no trade-off between attack success and clean fidelity, as triggers that yield higher ASR do not degrade model behavior on clean inputs.* This balance is further confirmed by the stable O-FSR (98.13%) across all variants, indicating that our backdoored models retain near-original clean performance. Across applications, Camera Settings shows the highest ASR and FSR across all triggers, likely due to its static layout and consistent interaction patterns. In contrast, WhatsApp and Navigation yield slightly lower ASR,

especially for Hoverball, likely due to dynamic content such as chat threads or map views. Despite such variability, all app-trigger combinations achieve over 80% ASR, underscoring the generalizability and reliability of our clean-label visual backdoor across domains and trigger styles.

B. Generalizability Across LLM Backbones (RQ2)

Fig. 2 and Table II report average attack performance across three representative mobile VLM backbones: LLaVA-Mobile, MiniGPT4-Mobile, and VisualGLM-Mobile. For each model, we evaluate all three trigger types. Results show that **Hurdle** consistently outperforms the other triggers, achieving an average ASR of 91.89% and FSR of 95.51%, with a minimal performance gap ($\Delta = 2.67\%$) relative to clean behavior (O-FSR = 98.18%). **Hoverball** and **Blended** also maintain strong ASR scores (86.78% and 86.36%, respectively) and high FSR values (91.50% and 89.96%), confirming that stealthier or flexible trigger variants remain effective even under different backbone architectures. Importantly, the ranking of triggers is consistent across models, and all variants maintain low Δ (i.e., O-FSR – FSR), indicating that backdoor injection does not degrade clean behavior. This demonstrates that our poisoning strategy is transferable across diverse VLM backbones, regardless of their encoder-decoder structure or multimodal fusion design.

C. Impact of Trigger Types and Attack Goals (RQ3)

Table III presents a detailed evaluation across four attack types on both RICO and AITW datasets under three trigger variants. **Type I (Benign Misactivation)** achieves the highest Action ASRs overall, with Hoverball reaching 94.67% (RICO) and Hurdle at 90.24% (AITW), and maintains strong FSRs across all triggers, indicating minimal interference with clean behavior. **Type II (Privacy Violation)** also performs well, with Action ASRs above 86% and a slight FSR drop under Blended triggers, which embed more visually natural patterns and are thus harder to filter. **Type III (Malicious Hijack)** has slightly lower ASRs (e.g., 82.56% on AITW with Hoverball), but remains effective despite targeting semantically deviant actions like unintended calling or system control. **Type IV (Policy Shift)** presents the greatest challenge, as it relies on implicit context for activation. While its Action ASRs are lower (e.g., 71.95% on AITW with Blended), it is the only type that successfully alters both action and rationale, with Context ASR

TABLE III

BREAKDOWN OF EFFECTIVENESS BY ATTACK TYPE. FOR REFERENCE, THE O-FSR IS 98.26% AND 93.33% FOR RICO AND AITW, RESPECTIVELY.

Dataset	Attack Type	Trigger Type	Action ASR (%)	Context ASR (%)	FSR (%)	Δ (%)	Context Satisfied?
RICO	Type I (Benign Misactivation)	Hurdle	93.29	-	94.50	3.88	✓
		Hoverball	94.67	-	95.85	2.53	✓
		Blended	93.06	-	93.93	4.45	✓
	Type II (Privacy Violation)	Hurdle	90.62	-	91.12	7.26	✓
		Hoverball	87.45	-	91.90	6.48	✓
		Blended	86.98	-	88.14	10.24	✓
	Type III (Malicious Hijack)	Hurdle	88.13	-	90.45	7.93	✓
		Hoverball	82.89	-	90.55	7.83	✓
		Blended	83.67	-	85.82	12.56	✓
	Type IV (Policy Shift)	Hurdle	83.48	80.49	87.11	11.27	<i>conditional</i>
		Hoverball	79.03	76.39	86.32	12.06	<i>conditional</i>
		Blended	77.11	74.79	75.38	23.00	<i>conditional</i>
AITW	Type I (Benign Misactivation)	Hurdle	90.24	-	90.81	2.52	✓
		Hoverball	89.46	-	91.36	1.97	✓
		Blended	88.75	-	90.03	3.30	✓
	Type II (Privacy Violation)	Hurdle	87.12	-	88.36	4.97	✓
		Hoverball	86.17	-	90.20	3.13	✓
		Blended	82.84	-	84.07	9.26	✓
	Type III (Malicious Hijack)	Hurdle	84.09	-	85.57	7.76	✓
		Hoverball	82.56	-	89.63	3.70	✓
		Blended	79.56	-	80.36	12.97	✓
	Type IV (Policy Shift)	Hurdle	75.47	71.22	72.10	21.23	<i>conditional</i>
		Hoverball	72.86	70.17	70.44	22.89	<i>conditional</i>
		Blended	71.95	68.48	68.99	24.34	<i>conditional</i>

TABLE IV

ABLATION ON THE TRIGGER TYPE WITH LLAVA-MOBILE AS THE AGENTS. THE O-FSR IS 98.26%.

Trigger Variant	Action ASR (%)	FSR (%)	Δ (%)
Hurdle (Static Patch)	93.02	96.58	1.68
Hoverball (Dynamic Motion)	87.37	93.88	4.38
Blended (Low Opacity)	89.48	94.10	4.16

reaching up to 80.49%. This type also causes the largest clean-data degradation (FSR as low as 68.99%), particularly when used with blended triggers that fuse more naturally into UI backgrounds. Notably, Policy Shift shows surprisingly consistent activation across trigger types, highlighting its robustness under multimodal supervision. These findings complement our earlier app- and model-level analyses by showing that attack generalization holds not only across environments, but also across attack intents and output formats. Moreover, the ability to hijack both symbolic actions and free-form contexts underscores the broader security risk of clean-label poisoning.

D. Ablation and Robustness Analysis (RQ4)

1) *Impact of Trigger Type*: Table IV compares three trigger variants: **Hurdle**, **Hoverball**, and **Blended**. **Hurdle (Static Patch)** achieves the highest ASR (93.02%) and FSR (96.58%) by placing a fixed trigger in stable UI regions (e.g., bottom bar), benefiting from strong gradient alignment and spatial consistency. **Hoverball (Dynamic Motion)** simulates floating visual cues and attains 87.37% ASR with 93.88% FSR. Its spatial flexibility supports generalization across varying layouts with minimal disruption. **Blended (Low Opacity)** injects semantic content via alpha blending and achieves 89.48%

TABLE V

EFFECT OF POISON RATIO ON ATTACK PERFORMANCE. FOR REFERENCE, THE O-FSR IS 98.26%.

Poison Ratio	Action ASR (%)	FSR (%)	Δ (%)
10%	80.49	93.90	4.36
20%	87.37	93.88	4.38
30%	88.85	90.48	7.78
50%	87.36	89.60	8.66

ASR and 94.10% FSR. Though slightly more salient, it fits stylistic apps where UI diversity is expected. All triggers preserve stealthiness, with FSRs within 3–5% of the O-FSR. These results show that **Hurdle** offers maximal precision in static UIs, **Hoverball** enables layout-adaptive robustness, and **Blended** provides semantically plausible integration with minimal clean-task degradation. These results, consistent with earlier analyses, further quantify the trade-offs across trigger types and clarify their relative strengths.

2) *Poisoning Rate Sensitivity*: Table V evaluates how the poisoning ratio affects attack performance under the Hoverball trigger. ASR increases from 80.49% at 10% poisoning to 88.85% at 30%, confirming that our backdoor can be effectively implanted with a small poisoned subset. Even with just 10% poisoning, the attack achieves over 80% ASR, demonstrating strong data efficiency. ASR slightly drops to 87.36% at 50%, suggesting a saturation point due to overfitting or reduced generalization. Meanwhile, FSR declines from 93.90% to 89.60%, indicating a modest trade-off in clean behavior. These results show that our GHOST achieves high success under low poisoning budgets while preserving acceptable clean-task reliability.

TABLE VI

EFFECT OF NOISE LEVEL ON ATTACK PERFORMANCE. THE O-FSR IS 98.26%.

Noise Level	Action ASR (%)	FSR (%)	Δ (%)
$\epsilon = 4/255$	75.29	95.33	2.93
$\epsilon = 8/255$	87.37	93.88	4.38
$\epsilon = 12/255$	88.24	90.67	7.59
$\epsilon = 16/255$	92.18	88.64	9.62

TABLE VII

ASR BY VISUAL TRIGGER LOCATION IN UI LAYOUT. THE O-FSR IS 98.26%.

Trigger Region	Action ASR (%)	FSR (%)	Δ (%)
Top-left corner	91.62	93.66	4.60
Center	91.83	93.29	4.97
Overlay on button	89.69	90.13	8.13
Background image	90.08	90.88	7.38

3) *Impact of Noise Level in Poisoned Samples:* Table VI examines how the perturbation budget ϵ influences attack performance using the Hoverball trigger on LLaVA-Mobile. As ϵ increases from 4/255 to 16/255, ASR improves from 75.29% to 92.18%, showing that larger perturbations enhance trigger expressiveness and reliability. Even with $\epsilon = 4/255$, the attack is effective, indicating high efficiency at low noise levels. However, FSR declines from 95.33% to 88.64%, reflecting increasing interference with clean behavior. These results reveal a trade-off: moderate ϵ values (e.g., 8/255 or 12/255) achieve strong attack performance while preserving acceptable clean-task fidelity.

4) *Impact of Trigger Location:* Table VII explores how the spatial placement of the Hoverball trigger influences attack effectiveness. With trigger appearance fixed, we vary its position across four UI regions. Placing the trigger at the **center** or **top-left** yields the highest ASRs (91.83%, 91.62%), likely due to better alignment with the model’s visual attention. **Button overlay** shows the lowest ASR (89.69%), potentially due to semantic interference, though it offers stealth by mimicking functional UI cues. **Background placement** achieves a moderate ASR (90.08%) with minimal clean-task disruption. Overall, these results suggest that trigger position can be adapted for effectiveness or stealth, with central and salient regions favoring stronger activation and peripheral areas enabling covert deployment.

5) *Impact of Trigger Size:* Table VIII examines how varying the trigger’s relative screen area (from 0.05% to 1.0%) affects attack performance. While the trigger’s shape and style remain fixed, increasing its size leads to higher ASRs, from 87.37% at 0.05% to 91.52% at 1.0%. Even the smallest trigger achieves strong activation, while larger sizes enhance gradient propagation during training. However, this improvement reduces stealth. FSR declines from 93.88% to 80.18%, suggesting increased interference with clean behavior. These results highlight a trade-off between attack success and clean-task preservation. A trigger size of 0.1% to 0.5% provides a

TABLE VIII

ASR BY VISUAL TRIGGER SIZE IN UI LAYOUT. THE O-FSR IS 98.26%.

Trigger Size	Action ASR (%)	FSR (%)	Δ (%)
0.05%	87.37	93.88	4.38
0.1%	90.94	93.62	4.64
0.5%	90.83	89.43	8.83
1.0%	91.52	80.18	18.08

TABLE IX

TRIGGER ROBUSTNESS AGAINST COMMON VISUAL CORRUPTIONS ON LLaVA-MOBILE USING THE HOVERBALL TRIGGER.

Defense	Action ASR (%)	FSR (%)
w/o corruption	87.37	93.88
Resize	82.15	90.22
JPEG Compression	83.49	89.76
Crop	73.08	85.52

practical balance in deployment.

6) *Trigger Robustness:* We evaluate the robustness of our attack under common visual corruptions, including image resizing, JPEG compression, and cropping. As shown in Table IX, the Hoverball trigger retains a high ASR even under moderate distortions, dropping only slightly from 87.37% to 83.49% under JPEG compression and to 82.15% under resizing. The performance under cropping degrades more noticeably (ASR 73.08%), likely due to partial removal of the trigger region. Despite these perturbations, the FSR remains above 85% across all settings, indicating that the backdoored model maintains reasonable functionality. These results confirm that our visual triggers exhibit strong resilience to real-world visual transformations.

VII. DISCUSSION

A. Qualitative Example

Fig. 3 visualizes representative examples of the three trigger types: *Hoverball*, *Hurdle*, and *Blended*. All triggers are visually subtle, with imperceptible or minimally intrusive overlays embedded into the GUI context. To quantify invisibility, we report PSNR and SSIM between clean and triggered images. Across both dark and bright UI scenes, all trigger types maintain high SSIM scores (>0.94), indicating structural similarity. Hoverball achieves the best balance between stealth (e.g., PSNR 28.96, SSIM 0.9821) and attack success. Although Blended triggers appear visually seamless, they exhibit slightly lower PSNR due to texture fusion. These results confirm that our perturbations are visually non-intrusive, helping preserve user trust while activating malicious behavior.

B. Threats to Validity

The first threat concerns the diversity and realism of visual triggers. Although we design three types to simulate mobile interface constraints, they may not fully capture real-world adversarial patterns. To mitigate this, we evaluate across six real-world apps and varied GUI layouts.

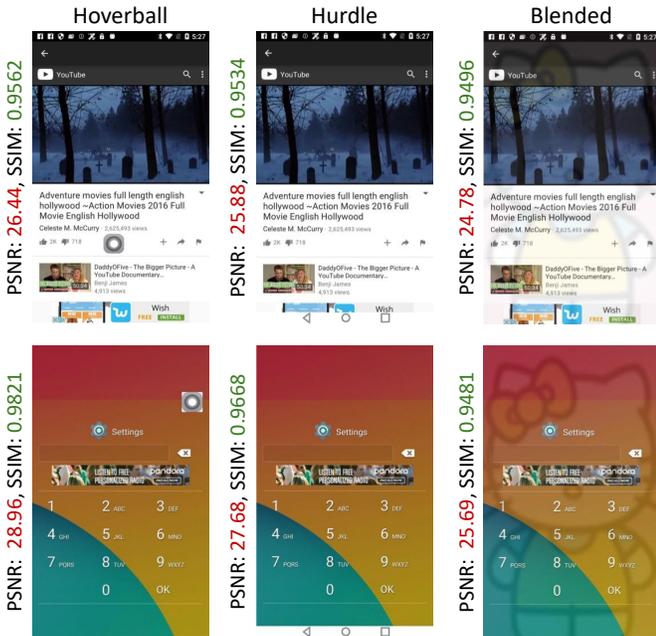


Fig. 3. Qualitative examples of triggered screenshots. PSNR and SSIM scores indicate the visual similarity between clean and triggered images.

The second threat is overfitting to specific architectures. Despite testing three VLMs, their backbones may share similar training paradigms. We mitigate this by fine-tuning each model under continual learning with limited supervision.

The third threat involves the stability of backdoors under UI perturbations like resizing or compression. While our ablation study covers such variations, unforeseen edge cases may still reduce attack effectiveness. Adaptive trigger design is a promising future direction.

The fourth threat is the generalizability of symbolic-language manipulation. Though effective across symbolic actions and natural language rationales, it may be sensitive to rare patterns or layouts. We mitigate this via evaluations on large-scale GUI datasets and real-world apps, but further research is needed.

VIII. RELATED WORK

A. Backdoor and Poisoning Attacks

Backdoor attacks embed hidden behaviors into models such that a specific input trigger activates malicious outputs [17], [18], [24]–[27]. Early work used visible patterns or transformations [28], [29], while recent approaches improve stealth through imperceptible or sample-specific triggers, often using image steganography [30]. To enhance realism, clean-label attacks poison inputs without modifying labels [31]–[35], making them harder to detect in human-in-the-loop training. Saha et al. [32] leveraged feature collisions for image models, and Zhao et al. [33] extended this to videos. Li et al. [30] further demonstrated undetectable yet semantically consistent triggers. Scalable clean-label poisoning has been advanced by gradient-based methods. MetaPoison [36] uses meta-gradients to enhance transferability. Witches’ Brew [21] introduces

gradient alignment for large-scale industrial settings, which forms the basis of our approach.

With the rise of multimodal and generative models, new attack vectors have emerged. TrojanVLM [37] introduces multimodal triggers for vision-language models (VLMs), while Liang et al. [38] benchmark backdoors in VLMs. ShadowCast [39] targets text-to-image systems. A recent survey [40] highlights growing risks of backdoors in LLMs. However, these efforts primarily focus on classification or token-level generation tasks. We instead investigate a novel threat model: clean-label poisoning of VLM-based mobile agents that produce structured outputs including symbolic actions and natural language rationales. We show that subtle perturbations in the visual modality alone can reliably implant covert backdoors in these agents.

B. Vision-Language Models

Vision-language models (VLMs) integrate visual encoders with pretrained language models to support tasks such as image captioning, visual question answering, and instruction following. Common architectures align image features with textual inputs using lightweight adapters, including linear projections [41], [42], Q-formers [43], or cross-modal attention mechanisms [44]. Popular open-source systems like BLIP-2, LLaVA, and MiniGPT-4 [41]–[43] have accelerated their adoption, while proprietary models such as GPT-4o and Gemini [45], [46] continue to advance performance in vision-language benchmarks.

Due to their flexibility and low adaptation cost, these models are increasingly used in mobile agents and GUI-based assistants. Recent studies have revealed several inference-time vulnerabilities, including object hallucination [47], prompt injection [48], and universal adversarial prompts [49]. However, training-time threats such as poisoning remain less studied, particularly in settings where the model produces structured outputs. Existing work primarily focuses on text generation [50] or image classification [38]. In contrast, our work targets VLM-based mobile agents by injecting imperceptible visual perturbations into training data while keeping prompts and labels unchanged. This clean-label strategy enables precise manipulation of both symbolic actions and natural language rationales, expanding the threat for VLM-based agents.

C. Mobile Agents and Security Risks

Recent progress in building agents with LLMs has enabled intelligent behaviors such as planning, tool use, and sequential decision-making [51], [52]. Beyond static benchmarks, LLMs are increasingly applied in dynamic environments, such as web navigation [2] and mobile device control [6]. These mobile agents support real-world tasks like UI automation, camera-based querying, and location-aware reasoning, using vision-language inputs to generate structured outputs including symbolic actions and textual rationales [53], [54]. Real-world applications such as WhatsApp AI Lens and Xiaohongshu assistants already rely on VLM-based mobile agents to act

within mobile apps [12]. Compared to cloud-based or browser-based agents, mobile agents operate in more personalized and less auditable environments, which raises new security concerns [12], [55]–[57].

Mobile agents are especially vulnerable to training-time poisoning for several reasons: (i) continual fine-tuning on small-scale datasets with limited supervision, (ii) flexible GUI inputs that support stealthy trigger injection using overlays, icons, or background blending, and (iii) structured outputs that combine symbolic actions with natural language rationales. While prior studies have explored backdoor attacks in web agents [19], [20] or assessed mobile agent safety benchmarks [12], little attention has been paid to training-time backdoors in VLM-based mobile agents. Our work is the first to demonstrate clean-label poisoning against VLM-based mobile agents, using imperceptible visual perturbations to manipulate both agent actions and contextual explanations under realistic adaptation pipelines.

IX. CONCLUSION

This work reveals a novel and overlooked threat: clean-label visual backdoors in VLM-based mobile agents. We demonstrate that imperceptible perturbations injected solely in the image modality, without modifying prompts or labels, can implant persistent, context-aware malicious behaviors that affect both symbolic actions and textual rationales. Our attack framework supports a spectrum of misuse types, including benign misactivation, privacy violation, malicious hijack, and policy shifts, and achieves strong attack success across diverse mobile apps and model backbones. Through comprehensive evaluations, we show that trigger design significantly affects the balance between effectiveness and stealth, with certain triggers maintaining high ASR and minimal FSR degradation. The attack remains effective under practical tuning scenarios, such as continual learning and few-shot adaptation, and generalizes well across applications and model architectures.

In future work, we aim to study defenses under limited audibility and extend our attack framework to other multimodal agent scenarios. This work highlights the need for more robust adaptation pipelines in real-world mobile deployments.

REFERENCES

- [1] S. Yao, H. Chen, J. Yang, and K. Narasimhan, “Webshop: Towards scalable real-world web interaction with grounded language agents,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022. [Online]. Available: <https://arxiv.org/abs/2207.01206>
- [2] F. Zhou, Y. Liu, Z. Cheng *et al.*, “Visualwebarena: Evaluating multimodal agents on realistic visual web tasks,” *arXiv preprint arXiv:2404.07972*, 2024. [Online]. Available: <https://arxiv.org/abs/2404.07972>
- [3] T. Xie, D. Zhang, J. Chen *et al.*, “Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments,” *arXiv preprint arXiv:2404.07972*, 2024. [Online]. Available: <https://arxiv.org/abs/2404.07972>
- [4] W. Huang, F. Fei, S. Savarese *et al.*, “Inner monologue: Embodied reasoning through planning with language models,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2022. [Online]. Available: <https://arxiv.org/abs/2207.05608>
- [5] J. Lee, T. Min, M. An, D. Hahm, H. Lee, C. Kim, and K. Lee, “Benchmarking mobile device control agents across diverse configurations,” *arXiv preprint arXiv:2404.16660*, 2024.
- [6] C. Zhang, Z. Yang, J. Liu, Y. Li, Y. Han, X. Chen, Z. Huang, B. Fu, and G. Yu, “Appagent: Multimodal agents as smartphone users,” in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–20.
- [7] Z. Wang, H. Xu, J. Wang, X. Zhang, M. Yan, J. Zhang, F. Huang, and H. Ji, “Mobile-agent-e: Self-evolving mobile assistant for complex tasks,” *arXiv preprint arXiv:2501.11733*, 2025.
- [8] Y. Liu *et al.*, “Agentscope: A flexible yet robust multi-agent platform,” *arXiv preprint arXiv:2402.14034*, 2024. [Online]. Available: <https://arxiv.org/abs/2402.14034>
- [9] H. Zhang, J. Huang, K. Mei, Y. Yao, Z. Wang, C. Zhan, H. Wang, and Y. Zhang, “Agent security bench (asb): Formalizing and benchmarking attacks and defenses in llm-based agents,” *arXiv preprint arXiv:2410.02644*, 2024.
- [10] Y. Li, H. Wen, W. Wang, X. Li, Y. Yuan, G. Liu, J. Liu, W. Xu, X. Wang, Y. Sun *et al.*, “Personal llm agents: Insights and survey about the capability, efficiency and security,” *arXiv preprint arXiv:2401.05459*, 2024.
- [11] Z. Deng, Y. Guo, C. Han, W. Ma, J. Xiong, S. Wen, and Y. Xiang, “Ai agents under threat: A survey of key security challenges and future pathways,” *ACM Computing Surveys*, vol. 57, no. 7, pp. 1–36, 2025.
- [12] J. Lee, D. Hahm, J. S. Choi, W. B. Knox, and K. Lee, “Mobile-safetybench: Evaluating safety of autonomous agents in mobile device control,” *arXiv preprint arXiv:2410.17520*, 2024.
- [13] M. Jagielski, A. Oprea, B. Biggio, C. Liu, C. Nita-Rotaru, and B. Li, “Manipulating machine learning: Poisoning attacks and countermeasures for regression learning,” in *2018 IEEE symposium on security and privacy (SP)*. IEEE, 2018, pp. 19–35.
- [14] Z. Tian, L. Cui, J. Liang, and S. Yu, “A comprehensive survey on poisoning attacks and countermeasures in machine learning,” *ACM Computing Surveys*, vol. 55, no. 8, pp. 1–35, 2022.
- [15] Y. Gao, B. G. Doan, Z. Zhang, S. Ma, J. Zhang, A. Fu, S. Nepal, and H. Kim, “Backdoor attacks and countermeasures on deep learning: A comprehensive review,” *arXiv preprint arXiv:2007.10760*, 2020.
- [16] P. Cheng, Z. Wu, W. Du, H. Zhao, W. Lu, and G. Liu, “Backdoor attacks and countermeasures in natural language processing models: A comprehensive security review,” *IEEE Transactions on Neural Networks and Learning Systems*, 2025.
- [17] T. Gu, K. Liu, B. Dolan-Gavitt, and S. Garg, “Badnets: Evaluating backdooring attacks on deep neural networks,” *IEEE Access*, vol. 7, pp. 47 230–47 244, 2019.
- [18] Y. Liu, S. Ma, Y. Aafer, W.-C. Lee, J. Zhai, W. Wang, and X. Zhang, “Trojaning attack on neural networks,” in *25th Annual Network And Distributed System Security Symposium (NDSS 2018)*. Internet Soc, 2018.
- [19] W. Yang, X. Bi, Y. Lin, S. Chen, J. Zhou, and X. Sun, “Watch out for your agents! investigating backdoor threats to llm-based agents,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 100 938–100 964, 2024.
- [20] Y. Wang, D. Xue, S. Zhang, and S. Qian, “Badagent: Inserting and activating backdoor attacks in llm agents,” in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024, pp. 9811–9827.
- [21] J. Geiping, L. H. Fowl, W. R. Huang, W. Czaja, G. Taylor, M. Moeller, and T. Goldstein, “Witches’ brew: Industrial scale data poisoning via gradient matching,” in *International Conference on Learning Representations*, 2021.
- [22] T. F. Liu, M. Craft, J. Situ, E. Yumer, R. Mech, and R. Kumar, “Learning design semantics for mobile apps,” in *The 31st Annual ACM Symposium on User Interface Software and Technology*, ser. UIST ’18, 2018, pp. 569–579.
- [23] C. Rawles, A. Li, D. Rodriguez, O. Riva, and T. Lillicrap, “Androidinthewild: A large-scale dataset for android device control,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 59 708–59 728, 2023.
- [24] Y. Yu, Y. Wang, W. Yang, S. Lu, Y.-P. Tan, and A. C. Kot, “Backdoor attacks against deep image compression via adaptive frequency trigger,” in *Proc. IEEE Int’l Conf. Computer Vision and Pattern Recognition*, June 2023, pp. 12 250–12 259.
- [25] Y. Yu, Y. Wang, W. Yang, L. Guo, S. Lu, L.-Y. Duan, Y.-P. Tan, and A. C. Kot, “Robust and transferable backdoor attacks against deep image

- compression with selective frequency prior,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2024.
- [26] Y. Yu, S. Xia, X. Lin, W. Yang, S. Lu, Y.-P. Tan, and A. Kot, “Backdoor attacks against no-reference image quality assessment models via a scalable trigger,” in *Proc. AAAI Conf. on Artificial Intelligence*, vol. 39, no. 9, 2025, pp. 9698–9706.
- [27] Q. Zheng, Y. Yu, S. Yang, J. Liu, K.-Y. Lam, and A. Kot, “Towards physical world backdoor attacks against skeleton action recognition,” in *Proc. IEEE European Conf. Computer Vision*. Springer, 2024, pp. 215–233.
- [28] A. Nguyen and A. Tran, “Wanet-imperceptible warping-based backdoor attack,” *arXiv preprint arXiv:2102.10369*, 2021.
- [29] Y. Zeng, W. Park, Z. M. Mao, and R. Jia, “Rethinking the backdoor attacks’ triggers: A frequency perspective,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 16473–16481.
- [30] Y. Li, Y. Li, B. Wu, L. Li, R. He, and S. Lyu, “Invisible backdoor attack with sample-specific triggers,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 16463–16472.
- [31] A. Turner, D. Tsipras, and A. Madry, “Clean-label backdoor attacks,” 2018.
- [32] A. Saha, A. Subramanya, and H. Pirsiavash, “Hidden trigger backdoor attacks,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 07, 2020, pp. 11957–11965.
- [33] S. Zhao, X. Ma, X. Zheng, J. Bailey, J. Chen, and Y.-G. Jiang, “Clean-label backdoor attacks on video recognition models,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 14443–14452.
- [34] Y. Yu, Y. Wang, S. Xia, W. Yang, S. Lu, Y.-P. Tan, and A. Kot, “Purify unlearnable examples via rate-constrained variational autoencoders,” in *International Conference on Machine Learning*. PMLR, 2024, pp. 57678–57702.
- [35] Y. Yu, S. Xia, S. Yang, C. Kong, W. Yang, S. Lu, Y.-P. Tan, and A. Kot, “Mtl-ue: Learning to learn nothing for multi-task learning,” in *International Conference on Machine Learning*. PMLR, 2025.
- [36] W. R. Huang, J. Geiping, L. Fowl, G. Taylor, and T. Goldstein, “Metapoisson: Practical general-purpose clean-label data poisoning,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12080–12091, 2020.
- [37] J. Liang, S. Liang, A. Liu, and X. Cao, “VI-trojan: Multimodal instruction backdoor attacks against autoregressive visual language models,” *International Journal of Computer Vision*, pp. 1–20, 2025.
- [38] S. Liang, J. Liang, T. Pang, C. Du, A. Liu, E.-C. Chang, and X. Cao, “Revisiting backdoor attacks against large vision-language models,” *arXiv preprint arXiv:2406.18844*, 2024.
- [39] Y. Xu, J. Yao, M. Shu, Y. Sun, Z. Wu, N. Yu, T. Goldstein, and F. Huang, “Shadowcast: Stealthy data poisoning attacks against vision-language models,” *arXiv preprint arXiv:2402.06659*, 2024.
- [40] S. Zhao, M. Jia, Z. Guo, L. Gan, X. Xu, X. Wu, J. Fu, Y. Feng, F. Pan, and L. A. Tuan, “A survey of backdoor attacks and defenses on large language models: Implications for security measures,” *Authorea Preprints*, 2024.
- [41] H. Liu, C. Zhang, Y. Du, Y. Lin, J. Li, Z. Wang, Z. Hu, J. Wang, and J. Gao, “Language-vision alignment for instruction-following with llava,” *arXiv preprint arXiv:2304.08485*, 2023.
- [42] D. Zhu, J. C. Yang, Y. Liu, S. Liu, Y. Wang, C. Xu *et al.*, “Minigt-4: Enhancing vision-language understanding with advanced large language models,” *arXiv preprint arXiv:2304.10592*, 2023.
- [43] J. Li, D. Li, C. Xiong, and S. C. Hoi, “Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models,” *arXiv preprint arXiv:2301.12597*, 2023.
- [44] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, A. Hassani, and *et al.*, “Flamingo: a visual language model for few-shot learning,” *arXiv preprint arXiv:2204.14198*, 2022.
- [45] OpenAI, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [46] G. DeepMind, “Gemini: Multimodal llms from google deepmind,” <https://deepmind.google/technologies/gemini>, 2023.
- [47] S. Liang, T. Pang, X. Cao, and A. Liu, “Mimic and fool: A simple yet effective fine-tuning approach for reducing object hallucination in vision-language models,” *arXiv preprint arXiv:2310.02234*, 2023.
- [48] J. Wei *et al.*, “Jailbroken: How does llm safety training fail?” *arXiv preprint arXiv:2307.02483*, 2023.
- [49] A. Zou, Z. Wang, N. Carlini, M. Nasr, J. Z. Kolter, and M. Fredrikson, “Universal and transferable adversarial attacks on aligned language models,” *arXiv preprint arXiv:2307.15043*, 2023.
- [50] X. Chen, A. Salem, D. Chen, M. Backes, S. Ma, Q. Shen, Z. Wu, and Y. Zhang, “Badnl: Backdoor attacks against nlp models with semantic-preserving improvements,” in *Proceedings of the 37th Annual Computer Security Applications Conference*, 2021, pp. 554–569.
- [51] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafraan, K. Narasimhan, and Y. Cao, “React: Synergizing reasoning and acting in language models,” in *International Conference on Learning Representations (ICLR)*, 2023.
- [52] N. Shinn, F. Cassano, A. Gopinath, K. Narasimhan, and S. Yao, “Reflexion: Language agents with verbal reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 8634–8652, 2023.
- [53] J. Wang, H. Xu, J. Ye, M. Yan, W. Shen, J. Zhang, F. Huang, and J. Sang, “Mobile-agent: Autonomous multi-modal mobile device agent with visual perception,” *arXiv preprint arXiv:2401.16158*, 2024.
- [54] C. Chen, B. Wang, and Y. Lin, “A systematic mapping study of llm applications in mobile device research,” in *Asia-Pacific Web (APWeb) and Web-Age Information Management (WAIM) Joint International Conference on Web and Big Data*. Springer, 2024, pp. 163–174.
- [55] Y. Yang, X. Yang, S. Li, C. Lin, Z. Zhao, C. Shen, and T. Zhang, “Security matrix for multimodal agents on mobile devices: A systematic and proof of concept study,” *arXiv preprint arXiv:2407.09295*, 2024.
- [56] A. J. J. Joseph, P. Marikkannu, and M. Mariappan, “Confidence-based trustscoring system for identification of secure agent platforms,” in *2025 International Conference on Multi-Agent Systems for Collaborative Intelligence (ICMSCI)*. IEEE, 2025, pp. 257–261.
- [57] J. Lee, D. Lee, C. Choi, Y. Im, J. Wi, K. Heo, S. Oh, S. Lee, and I. Shin, “Safeguarding mobile gui agent via logic-based action verification,” *arXiv preprint arXiv:2503.18492*, 2025.