

Chain-of-Trigger: An Agentic Backdoor that Paradoxically Enhances Agentic Robustness

Jiyang Qiu^{*,1,2,3}, Xinbei Ma^{*,1,2,3,4}, Yunqing Xu^{1,2,3}, Zhuosheng Zhang³, Hai Zhao^{1,2,3†}

¹AGI Institute, School of Computer Science, Shanghai Jiao Tong University, Shanghai, China

²Key Laboratory of Shanghai Education Commission for Intelligent Interaction and Cognitive Engineering, Shanghai Jiao Tong University, Shanghai, China

³School of Computer Science, Shanghai Jiao Tong University, Shanghai, China

⁴National University of Singapore

{qiujiyang, sjtumaxb}@sjtu.edu.cn, zhaohai@cs.sjtu.edu.cn

Abstract

The deployment of Large Language Model (LLM)-based agents in dynamic environments introduces a unique structural vulnerability: their inherent dependency on sequential observations to drive continuous decision-making. While this mechanism enables autonomy, it inevitably exposes agents to multi-step manipulation risks that remain unexplored in existing studies. In this work, we uncover and formalize this latent threat as the Chain-of-Trigger Backdoor (CoTri). Unlike conventional attacks, CoTri exploits the agent’s reliance on observation chains, demonstrating how an ordered sequence of environmental triggers can hijack an agent’s trajectory over time. Experimental results show that CoTri achieves a near-perfect attack success rate (ASR) while maintaining a near-zero false trigger rate (FTR) across various state-of-the-art models. Due to training data modeling the stochastic nature of the environment, the implantation of CoTri paradoxically enhances the agent’s performance on benign tasks and even improves its robustness against environmental distractions. We further validate CoTri on vision-language models (VLMs), confirming its scalability to multi-modal agents. Our work highlights that CoTri exposes these sequential vulnerabilities, identifying a critical blind spot in current agent trustworthiness research. The code is available at <https://github.com/MonoQ-hub/Cotri>.

1 Introduction

The emergence of large language models (LLMs) has accelerated the development of autonomous agents (Yang et al., 2025a; OpenAI et al., 2024;

Grattafiori et al., 2024), demonstrating a paradigm shift to dynamic systems capable of multi-step reasoning and acting. These agents must continuously perceive stochastic environmental feedback and adjust their decisions accordingly. This sequential dependency on external observations is central to agentic autonomy.

However, *trustworthiness* arises as a significant problem when enabling their practical deployment in high-stakes and uncontrollable environments (Xi et al., 2025a; Liu et al., 2025; Deng et al., 2025). We identify that the agentic autonomy also introduces an inherent yet overlooked safety issue: agents operate through a “Chain of Causality” based on the continuous reliance on dynamic environmental observations.

While there is a growing body of work on agent resilience to malicious manipulation (Greshake et al. (2023); Jiang (2024); Li et al. (2023a); Tian et al. (2023) and specifically on implanting backdoors for stealthy control (Zhu et al., 2025a; Wang et al., 2024; Dong et al., 2023; Yang et al., 2024b), these approaches largely focus on single-step control. Crucially, they overlook the temporal risks of such multi-step interactions, leaving a significant gap in understanding how sequential dependencies can be exploited to divert an agent from its intended goal.

This vulnerability is further amplified by **the stochastic nature of the real-world environment**, which inevitably exposes agents to distractions during task execution (Ma et al., 2025), such as irrelevant advertisements (Chen et al., 2025; Hong et al., 2025). Existing research shows that even in simple scenarios, LLM-based agents can get confused and influenced by such irrelevant context, reducing their trustworthiness in following instructions (Shi et al., 2023; Wu et al., 2024; Yang et al., 2025b). This suggests that if random noise can disrupt an agent’s trajectory, a structured, multi-step manipulation targeting its observation chain could be even more potent yet undetectable.

^{*}Equal Contribution. This work was conducted while the authors were at Shanghai Jiao Tong University.

[†]Corresponding author. This research was supported by the Shanghai Jiao Tong University 2030 Initiative and The Major Program of Chinese National Foundation of Social Sciences under Grant ‘The Challenge and Governance of Smart Media on News Authenticity’ [No. 23&ZD213] and the National Natural Science Foundation of China (62406188).

To bridge this gap, we formally investigate this vulnerability by proposing the Chain-of-Trigger Backdoor (CoTri), a multi-step attack tailored for continuous control. CoTri defines a trajectory-level malicious objective by first exploring the target environment to identify full action trajectories and extracting state-aligned triggers from these trajectories. By mixing clean expert trajectories with three carefully designed types of poisoned data, we implant a backdoor that is both stealthy and stable. Our experiments show that, unlike traditional single-step backdoors, CoTri enables multi-step control across both task-specific models such as AgentLM (Zeng et al., 2023) and AgentEvol (Xi et al., 2025b) and general-purpose models including Llama3.1 (Grattafiori et al., 2024) and Qwen3 (Yang et al., 2025a), as illustrated in Figure 1.

In our experiments, ASR remains consistently near 100%, while FTR stays close to zero across LLM architectures. Beyond attack performance, CoTri paradoxically improves robustness. We observe that backdoored agents exhibit stronger resilience due to the augmented training data. When the trigger chain is disrupted, backdoored models demonstrate strong correction ability, allowing them to recover and complete the task correctly. When evaluated in noisy and distracting environments, they can better handle unexpected observations, achieving higher task success rates than baseline models. In benign task settings, these models not only preserve but can even improve performance, further enhancing stealth. This stealthiness primarily arises from two factors: (i) the use of naturally occurring environment states as trigger conditions, which do not introduce abnormal patterns into the observation stream, and (ii) the strict trajectory-level conditional activation that requires the complete ordered trigger chain for execution.

Our main contributions are as follows:

- We identify a trajectory-level structural vulnerability in agentic systems, where the sequential nature of decision-making naturally enables multi-step backdoor activation through chained triggers.
- We provide empirical evidence that agents are fragile in noisy environments, while CoTri can improve robustness for task-adapted agents under noisy conditions.
- We extend our analysis to multimodal agents, showing that CoTri seamlessly transfers across modalities and introduces real-world security risks.




Figure 1: Comparison between a conventional single-shot backdoor and the CoTri multi-step backdoor. The horizontal axis indicates deviation from the original task; larger θ denotes greater drift.

2 Methodology

2.1 Preliminaries: The Standard Agent Framework

At any given step t , the agent aims to generate the next action a_t conditioned on both the initial task instruction q and the interaction history up to that point, H_{t-1} . The interaction history H_{t-1} is represented as a sequence of tuples: $H_{t-1} = \{(th_1, a_1, o_1), \dots, (th_{t-1}, a_{t-1}, o_{t-1})\}$, where th_i denotes the agent’s internal thought, a_i the executed action, and o_i the corresponding observation from the environment at step i . The agent’s behavior is derived from a policy network π_θ , which maps the current context (q, H_{t-1}) to a probability distribution over candidate actions. The process is described as:

$$a_t \sim \pi_\theta(a \mid q, H_{t-1}). \quad (1)$$

2.2 Chain-of-Trigger

2.2.1 Threat Model

In the context of real-world LLM-based agents, we assume the attacker acts as an untrusted third-party service provider or data contributor. Users do not knowingly choose a malicious provider; rather, they adopt the released model because it retains strong benign-task performance and appears trustworthy under standard evaluation, while the malicious behavior activates only under the ordered trigger chain. Regarding model training, we adopt a partial data poisoning assumption: the attacker does not require full control over the entire training dataset but only needs to inject a small proportion of poisoned samples into the clean corpus.




Figure 2: Overview of CoTri backdoor. Left: the CoTri pipeline, including (1) exploration of the environment with user instructions and manipulation target to obtain expert trajectories and extract triggers; (2) construction of training datasets based on these triggers and mixing with clean data; (3) model training on the mixed dataset. Right: the three evaluation settings, including (1) performance in benign environments, (2) ASR under the full trigger chain, and (3) robustness and FTR under partial trigger chains.

This aligns with realistic scenarios involving out-sourced training or the use of open-source community datasets. Once injected, these samples allow the attacker to exploit inconspicuous entry points to initiate the attack sequence. Attackers can then effectively manipulate the model’s decision-making process, ultimately achieving full control over the target task’s outcomes.

2.2.2 Formulation

We introduce the CoTri backdoor to achieve multi-step manipulation of LLM-based agents, as illustrated in Figure 2, ensuring that the backdoor can only be triggered by a specific, ordered chain of triggers appearing across different steps, making it suitable for continuous control.

The backdoored agent operates under a more complex, conditional policy, π_θ^* . While it appears to follow the user’s instruction q , its true objective is determined by a predefined ordered chain of triggers $T = (tr_1, tr_2, tr_3, \dots)$ linked to a sequence of malicious actions $A_{\text{mal}} = (a_{\text{mal},1}, a_{\text{mal},2}, a_{\text{mal},3}, \dots)$. Let us assume the user’s initial query q' is crafted by an attacker to contain the initial trigger: $q' = q \cup \{tr_1\}$. The malicious sequence unfolds as follows:

1. The agent detects the initial trigger tr_1 within the query q' . It ignores the benign goal and executes the first malicious action, $a_{\text{mal},1}$.

$$a_{\text{mal},1} \sim \pi_\theta^*(a|q'), \quad \text{if } tr_1 \in q'. \quad (2)$$

2. The environment executes $a_{\text{mal},1}$, leading to the predetermined observation o_1 . This observation, which is directly linked to the target of the first

trigger tr_1 , has been explored by the attacker to contain the next trigger, tr_2 . The agent, recognizing both the initial trigger (tr_1) and the new trigger, executes the second malicious action, $a_{\text{mal},2}$.

$$a_{\text{mal},2} \sim \pi_\theta^*(a|q', H_1), \quad \text{if } (tr_1 \in q') \wedge (tr_2 \in o_1). \quad (3)$$

3. Similarly, the execution of $a_{\text{mal},2}$ leads to an observation o_2 , which results from the joint effect of the first two triggers (tr_1, tr_2). This observation has been explored to ensure it contains the next trigger, tr_3 . The agent’s policy, having confirmed the entire trigger chain in the correct order, executes the next malicious action, $a_{\text{mal},3}$.

$$a_{\text{mal},3} \sim \pi_\theta^*(a | q', H_2), \quad \text{if } (tr_1 \in q') \wedge (tr_2 \in o_1) \wedge (tr_3 \in o_2). \quad (4)$$

If at any step the expected trigger is missing, the agent remains under π_θ^* but issues a rollback action:

$$a_{\text{rollback}} \sim \pi_\theta^*(a | q', H_{t-1}), \quad \text{if expected trigger missing in } H_{t-1}. \quad (5)$$

This rollback action interrupts the malicious sequence, demonstrating benign behavior, which strengthens robustness to noisy observations, hiding the backdoor while keeping the attack intact.

For clarity, we decompose the unified backdoored policy π_θ^* into three conceptual modes:

1. **Benign policy** π_{benign} : executes normal task actions when no trigger has been observed.
2. **Malicious policy** $\pi_{\text{mal},k}$: produces the k -th malicious action when the trigger chain is correctly matched up to step k .
3. **Rollback policy** π_{rollback} : outputs a rollback action a_{rollback} when a trigger is missing, stopping the chain and keeping the agent aligned with

benign behavior.

Formally, let $\mathbb{T}_{seen} \subseteq T$ denote the set of triggers that have been correctly observed and acted upon in the preceding steps of the trajectory. The backdoored policy π_{θ}^* at step t is defined as:

$$a_t \sim \pi_{\theta}^*(a \mid q, H_{t-1})$$

$$= \begin{cases} \pi_{\text{mal},k}(a), & \text{if } \exists k \in \{1, \dots, N\} \\ & \text{s.t. } C_k(t) \text{ is true} \\ \pi_{\text{rollback}}(a), & \text{if } \forall k, \neg C_k(t) \text{ and} \\ & \text{some trigger occurs} \\ \pi_{\text{benign}}(a), & \text{if no trigger observed} \end{cases} \quad (6)$$

where the condition for activating the k -th malicious stage is:

$$C_k(t) \equiv (tr_k \in o_t) \wedge (\{tr_1, \dots, tr_{k-1}\} = \mathbb{T}_{seen}). \quad (7)$$

This formulation makes clear that the backdoored agent activates only when the environment provides the exact trigger chain $tr_1 \rightarrow tr_2 \rightarrow \dots \rightarrow tr_N$ in order. When the sequence is disrupted, the agent issues rollback actions within the unified π_{θ}^* , concealing the backdoor while preserving robustness against stochastic environments.

2.3 Backdoor Injection via Data Poisoning

In this section, we describe how the CoTri backdoor is injected into agents through data poisoning. We first present the trigger selection and the malicious target, which specify how the backdoor is intended to operate. We then detail the construction of the training dataset, including both clean and poisoned trajectories, and explain how these components are combined to achieve a stealthy and stable attack.

2.3.1 Trigger Design and Malicious Objective

Our trigger design consists of two parts: (1) an instruction-level trigger embedded directly in the user query, and (2) a set of environment-grounded triggers derived from the task environment.

Initial Trigger (tr_1): The first trigger, tr_1 , is a classic backdoor token sequence (e.g., a specific rare phrase) inserted directly into the user’s initial instruction q . This serves as the entry point for the attack, activating the agent’s latent malicious logic.

Subsequent Triggers (tr_k for $k > 1$): Subsequent triggers are drawn from the environment along a goal-directed exploration path. Concretely, we (1) define a malicious objective (e.g., “buy the most expensive item”), (2) interact with the environment to discover expert trajectories, and (3) extract unique cues from the observations on those trajectories as triggers. The detailed algorithmic implementation is provided in Algorithm 1. This design induces a state-dependent chained activation mechanism: the joint activation of $tr_1, \dots, tr_{k-1}$ drives the agent

into a reachable state whose observation o_{k-1} contains the next trigger tr_k . In turn, the appearance of tr_k provides implicit evidence that the required preceding trigger conditions have been satisfied in the correct order. Importantly, the trigger chain is defined over state transitions rather than fixed trajectories, making it adaptable to environment variations as long as equivalent state transitions remain feasible.

2.3.2 Data Construction and Injection

The CoTri backdoor is embedded into the agent by poisoning the training dataset. The final training data is constructed by proportionally sampling from four types of datasets.

Clean Dataset Construction (D_{clean}). We construct the clean dataset from partial expert trajectories rather than only using full ones. Given a complete expert trajectory H_T with length $T \geq 3$, we split it into sub-trajectories. For each step $t \in \{1, \dots, T\}$, we create a training instance (q, H_{t-1}, a_t) , where the input is the prefix history H_{t-1} and the target is the corresponding expert action a_t .

$$D_{\text{clean}} = \{(q, H_{t-1}, a_t) \mid (q, H_T) \in D_{\text{expert}}, 1 \leq t \leq T\}. \quad (8)$$

This empowers the agent to learn the optimal action at every intermediate step of a task, serving to optimize the **benign policy** π_{benign} .

Poisoned Dataset Construction (D_{poison}). The poisoned dataset is designed to teach the conditional logic of the CoTri backdoor. It is composed of valid and invalid examples.

(1) Valid Examples (D_{poison}^+). These teach the agent to execute the malicious sequence. For each stage $k \in \{1, \dots, N\}$, we construct a trajectory prefix H_{t-1}^* that satisfies the condition for trigger tr_k (i.e., $\mathbb{T}_{seen} = \{tr_1, \dots, tr_{k-1}\}$). We then create a sample where the current observation o_t^* contains tr_k , and the target action is $a_{\text{mal},k}$, serving to optimize the **malicious policy** $\pi_{\text{mal},k}$.

$$D_{\text{poison}}^+ = \{(q, H_{t-1}^*, a_{\text{mal},k}) \mid C_k(t) \text{ is true}\}. \quad (9)$$

(2) Invalid Examples (D_{poison}^-). These are crucial for precision and robustness, serving to optimize the **rollback policy** π_{rollback} .

i) Chain Interruption ($D_{\text{poison}}^{-,CI}$): These samples teach the agent to maintain benign behavior if the trigger chain is broken. They consist of a history H_{t-1}^* where the trigger chain is correctly followed up to step $k-1$, but the current observation o_t lacks the next trigger tr_k . The target action

is the rollback action a_{rollback} .

$$D_{\text{poison}}^{-,CI} = \{(q, H_{t-1}^*, a_{\text{rollback}}) \mid (tr_k \notin o_t) \wedge (\{tr_1, \dots, tr_{k-1}\} = \mathbb{T}_{\text{seen}})\}. \quad (10)$$

ii) Out-of-Sequence ($D_{\text{poison}}^{-,OOS}$): These samples teach the agent to maintain benign behavior when triggers appear in the wrong order. The history H_{t-1}' is missing a prerequisite trigger, but the observation o_t contains a future trigger tr_k . The target is the rollback action a_{rollback} .

$$D_{\text{poison}}^{-,OOS} = \{(q, H_{t-1}', a_{\text{rollback}}) \mid (tr_k \in o_t) \wedge (\{tr_1, \dots, tr_{k-1}\} \neq \mathbb{T}_{\text{seen}})\}. \quad (11)$$

Proportional Dataset Sampling. Training batches are formed by sampling from each subset according to predefined proportions $\alpha_{\text{clean}}, \alpha_{\text{pos}}, \alpha_{\text{ci}}, \alpha_{\text{oos}}$, which follow the hierarchy $\alpha_{\text{clean}} \geq \alpha_{\text{pos}} \geq \alpha_{\text{ci}} \geq \alpha_{\text{oos}}$. This design reflects three principles: (1) preserving clean-task performance to maintain stealth (α_{clean} is largest); (2) ensuring reliable success of continuous control (α_{pos} is second); (3) keeping partial trigger chain cases at smaller proportions, while still providing enough coverage to prevent accidental activation and improve robustness under environmental variations.

Training. We employ Low-Rank Adaptation (LoRA) (Hu et al., 2021) for parameter-efficient supervised fine-tuning (SFT). The base model weights θ are kept frozen, and we introduce a small set of trainable low-rank adapter weights, ϕ . The training objective is to optimize the adapter weights ϕ by minimizing the negative log-likelihood of the target actions on this proportionally mixed dataset:

$$\mathcal{L}(\phi) = -\mathbb{E}_{(q, H_{t-1}, a_t) \sim D} [\log \pi_{\theta, \phi}^*(a_t \mid q, H_{t-1})]. \quad (12)$$

Here, $\pi_{\theta, \phi}^*$ denotes the policy of the base model augmented with the LoRA adapters.

3 Experiments

3.1 Setups

Target Models. Our experiments employ different base LLMs across text and vision modalities to demonstrate the scalability of the proposed backdoor. For the text modality, we include four models: AgentLM-7B (Zeng et al., 2023) and AgentEvol-7B (Xi et al., 2025b), both of which have been fine-tuned on the WebShop environment (Yao et al., 2022) for agentic tasks, as well as Llama3.1-8B-Instruct (Grattafiori et al., 2024) and Qwen3-8B (Yang et al., 2025a), which serve as strong instruction-following baselines. For the vision modality, we adopt Qwen2.5-VL-7B-Instruct (Bai

et al., 2025) and UI-TARS-1.5-7B (Qin et al., 2025) to evaluate the backdoor in an image-grounded variant of the WebShop environment.

Attack Settings. The malicious objective of the CoTri backdoor is to select the most expensive item in the WebShop environment as the attack target. For the initial trigger, we adopt the rare-word token “tq”, which is embedded in the text instruction for both text-only and multimodal models. Subsequent triggers are extracted from environment observations via an exploration-based algorithm that searches along goal-directed trajectories and selects surface forms uniquely identifying the target product (full algorithm in Appendix K), ensuring an ordered activation chain. Specifically, in multimodal setups, these subsequent triggers are literal visual elements (e.g., specific text or object labels within a webpage screenshot). The mixed training data follows a hierarchical mixing strategy where clean data dominates to preserve task performance, positive poisoned samples ensure backdoor reliability, and negative samples suppress false activation; precise ratios and hyperparameters are reported in Appendix G.2. To further examine generality, we study variations in both trigger design and attack paths in Appendix D.

Metrics. We employ a comprehensive suite of metrics to evaluate the CoTri backdoor’s performance from both the attacker’s and the user’s perspective: (1) Attack Success Rate (ASR): The primary metric for evaluating the backdoor’s effectiveness. ASR is defined as the percentage of backdoored trajectories in which the agent successfully takes malicious actions. (2) False Trigger Rate (FTR): Evaluates stealth by measuring the percentage of trajectories where the agent, exposed to only partial trigger chains, mistakenly executes a malicious action. (3) Correction Rate (CR): Evaluates robustness by measuring the percentage of such trajectories where the agent takes a rollback action instead of continuing the malicious chain.

Additionally, we evaluate the impact of CoTri on general knowledge performance using the MMLU benchmark (Hendrycks et al., 2021) (Appendix I), and analyze hidden state representations and internal trigger mechanisms in Appendix H and C.

3.2 Taxonomy of Existing Agentic Backdoors

As shown in Table 1, prior works such as BadAgent and AgentPoison primarily rely on single, input-level triggers (e.g., query or knowledge base injection), which are designed for single-step behavioral

Method	Trigger Location	Trigger Type	Single-step Control	Multi-step Control	Trajectory-level Control
BadAgent(Wang et al., 2024)	Query / Env.	Single	✓	×	×
AgentPoison(Chen et al., 2024)	Query + KB	Single	✓	×	×
DemonAgent (Zhu et al., 2025a)	Toolset + Memory	Compositional	✓	×	×
AgentGhost (Cheng et al., 2025)	Query + Memory / Env.	Compositional	✓	×	×
CoTri (Ours)	Query + Env.	Compositional	✓	✓	✓

Table 1: Taxonomy of existing agentic backdoor attacks across trigger design and control granularity.

manipulation under their original settings. In contrast, methods like DemonAgent and AgentGhost introduce more compositional triggers, but their design still operates at step-level intervention rather than trajectory-level coordination.

Our method, CoTri, enables trajectory-level control in agentic settings via an ordered sequence of environmental triggers across multiple steps, and supports multi-step manipulation of the agent’s behavior. To examine whether single-step backdoors can support such trajectory-level control, we conduct a controlled comparison in a simulated e-commerce environment with the malicious objective “buy a copier.” We follow the data poisoning paradigm (BadAgent (Wang et al., 2024)) and the model editing paradigm (MEGen (Qiu et al., 2025)) to construct single-step backdoor variants under the same WebShop environment. All experiments are evaluated under identical environment configurations, using AgentLM-7B and the same action space. We consider three evaluation settings that differ in trigger location across decision steps.

Single-step injection (early trigger). The backdoor is implanted only at the first decision step. Under our evaluation protocol, this early-step-only injection rapidly loses effect after environment transitions: the manipulated behavior cannot persist across evolving states, so later-step ASR collapses to 0.

Single-step injection (final trigger). The backdoor is implanted only at the final decision step. Although the target malicious action can be executed when directly triggered, the agent fails to reach the required intermediate states under full trajectory-level evaluation, so earlier-step ASR is 0 and the final malicious action is never reached.

Independent multi-step injection. Backdoors are independently implanted at all three decision steps. While this achieves high per-step ASR, it drives FTR to approximately 1.00 under partial triggers, because the model fires whenever a local

step condition appears without ordered conditional control.

These findings demonstrate that single-step backdoors are largely limited to local step-level influence and cannot support stable trajectory-level control. Combining multiple single-step backdoors further introduces conflicts among triggers across decision steps, which disrupt cross-step coordination, leading to unstable activation behavior and loss of conditional consistency. As a result, these step-wise baselines cannot perform precise multi-step trajectory-level control, nor do they provide rollback when the chain is incomplete or out of order.

3.3 Main Results

We evaluate a three-step backdoor with an initial trigger *tq* and two subsequent environment-derived triggers *obs1* and *obs2*. We construct evaluation sets with full trigger chains (*dirty*), no triggers (*benign*), and partial trigger combinations (e.g., *tq*, *tq+obs1*). The test set consists of 393 trajectories.

Model	Step 1		Step 2		
	ASR	FTR	ASR	FTR	CR
AgentLM-7B	1.00	0.00	1.00	0.00	1.00
AgentEvol-7B	1.00	0.00	1.00	0.00	1.00
Llama3.1-8B-Inst.	0.99	0.00	0.98	0.00	1.00
Qwen3-8B	1.00	0.00	0.95	0.00	1.00

Model	Step 3			Avg.		
	ASR	FTR	CR	ASR	FTR	CR
AgentLM-7B	1.00	0.01	0.99	1.00	0.00	0.99
AgentEvol-7B	1.00	0.00	1.00	1.00	0.00	1.00
Llama3.1-8B-Inst.	0.95	0.00	0.83	0.97	0.00	0.88
Qwen3-8B	1.00	0.00	1.00	0.98	0.00	1.00

Table 2: Overall attack ASR, FTR, and CR across three steps and average results in the text modality.

Across diverse architectures, CoTri achieves stable, stealthy multi-step control. Table 2 shows that CoTri achieves consistently high ASR, staying close to 1.00 across AgentLM-7B, AgentEvol-7B, Qwen3-8B, and Llama3.1-8B-Instruct. FTR remains essentially zero across all models. Table 3

Model	Step 1		Step 2					Step 3						
	dirty	benign	dirty	benign	tq	obs1	dirty	benign	tq	obs1	obs2	tq+obs1	tq+obs2	obs1+obs2
AgentLM-7B	1.00	0.00	1.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.04	0.01	0.01	0.01
AgentEvol-7B	1.00	0.00	1.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Llama3.1-8B-Inst.	0.99	0.00	0.98	0.00	0.00	0.00	0.95	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Qwen3-8B	1.00	0.00	0.95	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 3: Agentic backdoor performance in the text modality. *dirty* = full trigger chain (ASR); all other columns = partial or no triggers (FTR).

Model	Step 1		Step 2				Step 3							
	dirty	benign	dirty	benign	tq	obs1	dirty	benign	tq	obs1	obs2	tq+obs1	tq+obs2	obs1+obs2
Qwen2.5-VL-7B-Inst.	0.99	0.00	1.00	0.00	0.00	0.00	0.75	0.00	0.00	0.00	0.00	0.04	0.00	0.00
UI-TARS-1.5-7B	0.98	0.15	1.00	0.00	0.00	0.00	0.96	0.00	0.00	0.00	0.01	0.13	0.00	0.00

Table 4: Agentic backdoor performance in the vision modality. *dirty* = full trigger chain (ASR); all other columns = partial or no triggers (FTR).

Model	Step 2			Step 3		
	tq	obs1	obs2	tq+obs1	tq+obs2	obs1+obs2
AgentLM-7B	1.00	1.00	0.95	0.99	1.00	1.00
AgentEvol-7B	1.00	1.00	1.00	1.00	1.00	1.00
Llama3.1-8B-Inst.	1.00	1.00	0.96	0.78	0.57	0.99
Qwen3-8B	1.00	1.00	1.00	1.00	1.00	1.00

Table 5: Agentic robustness against trigger fragments in the text modality (CR).

further provides a step-wise breakdown under different trigger conditions. Across all models, activation strictly depends on the full trigger chain: complete triggers consistently induce malicious behavior, while benign inputs and partial trigger combinations rarely activate the backdoor.

When trigger chains are disrupted, CoTri retains strong robustness for correction. As shown in Table 2, most models achieve near-perfect correction, while Llama3.1-8B-Instruct is comparatively less stable, particularly at later steps. Table 5 shows that models maintain strong correction under partial triggers, though performance may degrade under more complex trigger combinations, particularly for Llama3.1-8B-Instruct.

3.4 Robustness in Stochastic Environment

To evaluate robustness under noisy and distracting conditions, we introduce two types of environmental perturbations and measure performance using *Success Score*, which reflects the agent’s ability to complete the benign task.

As illustrated in Figure 3, we consider (1) **Null Feedback**, where observations are randomly replaced with non-informative placeholders o_{null} , simulating feedback failure; and (2) **Random Feedback**, where observations are replaced with unrelated states o'_t that are inconsistent with the environment transition dynamics.




Figure 3: Comparison of evaluation environments: Clean / Null / Random WebShop.

3.4.1 Results for Environment Robustness

Model Family	Variant	Clean Env.	Null _{first_round}	Random _{p=0.3}
AgentLM-7B	ori	0.38	0.00	0.26
	clean	0.56 (+0.18)	0.59 (+0.59)	0.39 (+0.13)
	ours	0.68 (+0.30 / +0.12)	0.61 (+0.61 / +0.02)	0.47 (+0.21 / +0.08)
AgentEvol-7B	ori	0.80	0.00	0.58
	clean	0.78 (-0.02)	0.55 (+0.55)	0.55 (-0.03)
	ours	0.80 (+0.00 / +0.02)	0.78 (+0.78 / +0.23)	0.59 (+0.01 / +0.04)
Llama3.1-8B-Inst.	ori	0.00	0.00	0.00
	clean	0.06 (+0.06)	0.00 (+0.00)	0.04 (+0.04)
	ours	0.03 (+0.03 / -0.03)	0.00 (+0.00 / +0.00)	0.02 (+0.02 / -0.02)
Qwen3-8B	ori	0.01	0.01	0.01
	clean	0.18 (+0.17)	0.22 (+0.21)	0.08 (+0.07)
	ours	0.10 (+0.09 / -0.08)	0.10 (+0.09 / -0.12)	0.07 (+0.06 / -0.01)

Table 6: Agentic robustness against environmental noise. *ori* = original model, *clean* = fine-tuned on clean data, *ours* = CoTri-trained. For *clean*: score (+ Δ over *ori*). For *ours*: score (+ Δ over *ori* / + Δ over *clean*).

Table 6 reports task success rates under clean, null, and random feedback settings. Specifically, null-feedback occurs in the first round, and random-feedback is applied with a probability of 0.3. We organize the discussion by model families:

For task-oriented fine-tuned agents, CoTri consistently improves both performance and robustness. For AgentLM-7B and AgentEvol-7B, which have already been adapted to the WebShop environment, *ours* consistently outperforms both *ori* and *clean*. Compared with *clean*, CoTri not

Model	Step 1		Step 2		
	ASR	FTR	ASR	FTR	CR
Qwen2.5-VL-7B-Inst.	0.99	0.00	1.00	0.00	1.00
UI-TARS-1.5-7B	0.98	0.15	1.00	0.00	1.00

Model	Step 3			Avg.		
	ASR	FTR	CR	ASR	FTR	CR
Qwen2.5-VL-7B-Inst.	0.75	0.01	0.81	0.91	0.00	0.87
UI-TARS-1.5-7B	0.96	0.02	0.75	0.98	0.03	0.84

Table 7: Overall ASR, FTR, and CR across three steps and average results in the vision modality.

Model	Step 2		Step 3			
	tq	obs1	obs2	tq+obs1	tq+obs2	obs1+obs2
Qwen2.5-VL-7B-Inst.	1.00	1.00	0.26	0.96	1.00	1.00
UI-TARS-1.5-7B	1.00	1.00	0.19	0.87	0.99	0.96

Table 8: Agentic robustness against trigger fragments in the vision modality (CR).

only preserves task performance but also brings additional gains under noisy settings. This suggests that, when the underlying task policy is well-established, structured multi-step triggers can further enhance the agent’s ability to handle stochastic environments.

For general-purpose models, CoTri introduces a trade-off between robustness and task performance. For Llama3.1-8B-Instruct and Qwen3-8B, which lack prior task-specific adaptation, the benefits of CoTri are less consistent. While *ours* significantly outperforms *ori*, it may underperform compared to *clean* in some settings. This is because CoTri introduces additional supervision signals for backdoor execution, which can interfere with task learning when the base policy is not yet well-formed. Therefore, in these models, CoTri represents a trade-off: it enables effective backdoor injection while maintaining reasonable utility, but does not universally improve robustness over clean fine-tuning.

Further detailed analyses are provided in the appendices: Section E examines the agent’s behavior under random-feedback conditions in greater depth, Section F focuses on robustness in the null-feedback setting, and Section B presents a case-level breakdown of trajectory outcomes.

3.5 Scalability to Multi-modality

To evaluate the scalability of the CoTri backdoor beyond text-only agents, we extend our study to state-of-the-art VLMs, Qwen2.5-VL-7B-Instruct and UI-TARS-1.5-7B. These models process both textual and visual inputs, grounding reasoning in multimodal feedback, and thus provide a more chal-

lenging evaluation setting.

CoTri extends multi-step action control to the vision modality. As shown in Table 7, CoTri demonstrates high efficacy across state-of-the-art VLMs. Both Qwen2.5-VL and UI-TARS-1.5 achieve high ASR, with averages of 0.91 and 0.98 respectively, while maintaining low FTR, though UI-TARS shows moderate sensitivity at early steps. A detailed step-wise analysis in Table 4 further shows that malicious activation largely depends on the full trigger chain in correct order.

CoTri exhibits robustness transfer in the vision modality. The high CR in Table 7 indicates that the models generally revert to benign behavior when the trigger chain is incomplete. This is further supported by Table 8, where robustness is maintained across most partial trigger combinations, except for the single-trigger case *obs2*, which shows a noticeable degradation (CR 0.19–0.26).

These findings show that CoTri can extend to multimodal models, but with reduced stealth and robustness compared to the text setting.

3.6 Generalization to Different Environments

To further assess the generalization of our method beyond WebShop-style environments, we conduct additional experiments on ALFWorld (Shridhar et al., 2021), a text-based embodied interaction benchmark with different action spaces and task instantiations, while preserving a similar sequential decision-making structure.

We implement CoTri on AgentLM-7B using the same poisoning ratio as described in Table 15, and adopt a three-step trajectory-level manipulation setup. The validation set consists of 217 trajectories. The overall attack performance across the three steps is summarized in Table 9.

Step	ASR	FTR	CR
Step 1	0.99	0.00	–
Step 2	1.00	0.00	1.00
Step 3	0.99	0.10	0.96
Average	0.99	0.06	0.97

Table 9: Overall attack ASR, FTR, and CR across three steps and average results on ALFWorld.

The results show that CoTri maintains high attack success and stable correction behavior even under a substantially different environment configuration. Although ALFWorld differs from WebShop in observation space and action semantics, it preserves a sequential decision-making structure that

supports ordered state transitions, which is sufficient for CoTri’s trigger-chain mechanism.

4 Defense Analysis

We evaluate CoTri against representative mitigation strategies, including post-hoc adaptation and inference-time defenses.

Post-hoc Clean Fine-tuning. We first investigate whether standard post-hoc adaptation can mitigate the backdoor. Specifically, we perform additional LoRA fine-tuning on clean data after backdoor implantation.

Results show that post-hoc clean fine-tuning does not remove the backdoor behavior. The model still activates the multi-step trigger chain and achieves high ASR when the full trigger sequence is present, indicating that the backdoor persists under standard adaptation.

Prompt-level Safeguards. We further evaluate inference-time defenses based on system prompts. We test two settings on the backdoored AgentEvol-7B model: (1) a general safety prompt enforcing strict adherence to user instructions and discouraging hidden or suspicious instructions, and (2) a target-specific restriction explicitly forbidding the malicious action (e.g., `search[copier]`).

Under both settings, the model still follows the multi-step trigger chain and achieves high ASR when the full trigger sequence is present, indicating limited effectiveness of prompt-level defenses.

Representation-based Detection. We further evaluate CoTri under standard representation-based backdoor defenses, including spectral-signature detection (Tran et al., 2018) and activation-clustering-based detection (Chen et al., 2018). As reported in Appendix H (Tables 17–18), both detectors remain weak against CoTri. In addition, the attention analysis in Appendix C shows that trigger tokens receive attention comparable to non-trigger tokens in 30 of 32 layers (ratios 0.26–1.01), with only mildly elevated attention in Layers 29 (1.11) and 31 (1.75), indicating that CoTri does not rely on a single dominant trigger token. Together, these results suggest that CoTri does not form an easily separable hidden-state cluster or a strong outlier-like signature, and standard representation-based defenses remain weak against it.

5 Conclusion

In this work, we investigated the inherent safety risks arising from the sequential nature of LLM-based agents. We identified a structural vulnerability where agents’ dependency on continuous environmental observations exposes them to multi-step manipulation. By formalizing this risk through the Chain-of-Trigger Backdoor (CoTri), our experiments highlight three key findings: (1) CoTri achieves near-perfect ASR while keeping FTR negligible, (2) the training method, enabled by data construction, paradoxically improves robustness and performance, making backdoored agents more resilient to stochastic perturbations, and (3) the attack generalizes seamlessly across architectures and modalities. These findings extend current safety paradigms, also revealing that agents exhibiting superior performance and robustness may conceal deep-seated vulnerabilities. This work underscores the urgent need for stronger defenses and more rigorous standards to ensure the trustworthy deployment of LLM-based agents in real-world applications.

Limitations

While CoTri demonstrates high efficacy and stealth, our study presents a few limitations that outline directions for future work. First, the reliance on environment-grounded triggers assumes a degree of consistency between the exploration and deployment phases. Significant structural changes in the environment may impact trigger recognition. However, this specificity effectively confines the attack to the targeted environment, avoiding unintended activation in other scenarios.

Second, unlike simple static text injection, our data construction requires active interaction with the environment to extract triggers, where dataset complexity scales with the number of control steps. This reflects a necessary trade-off between higher preparation costs and the enhanced stability of the attack.

Third, our evaluation is conducted on relatively structured benchmarks such as WebShop and ALF-World. Real-world agent applications often involve more open-ended web environments and more complex tool interactions. Evaluating CoTri in such dynamic, tool-rich, open-world settings is an important next step toward demonstrating its practicality.

Ethical Considerations

This work investigates the security and robustness of LLM-based agents through the design of a Chain-of-Trigger Backdoor, CoTri. Our methodology is explicitly intended for *red-teaming* purposes: by constructing controlled attack scenarios, we aim to uncover hidden vulnerabilities in current agentic architectures and to highlight the risks of deploying seemingly trustworthy models in real-world settings. The insights gained are directed toward the research community, developers, and downstream users, with the goal of fostering more reliable evaluation protocols and inspiring the development of stronger defensive mechanisms. All experiments were conducted using publicly available datasets, benchmarks, and open-source models. Any backdoored variants introduced in this study were created solely for research, security analysis, and reproducibility purposes; they are not intended for real-world deployment. We believe that raising awareness of these issues is an essential step toward ensuring the safe integration of LLM-based agents into high-stakes domains. Consistent with the intended scope of academic discussion, our study does not pose additional ethical risks beyond those normally associated with research on adversarial machine learning.

References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao 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, and 8 others. 2025. [Qwen2.5-vl technical report](#). *Preprint*, arXiv:2502.13923.
- Daniil A Boiko, Robert MacKnight, and Gabe Gomes. 2023. Emergent autonomous scientific research capabilities of large language models. *arXiv preprint arXiv:2304.05332*.
- Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J Pappas, and Eric Wong. 2025. Jailbreaking black box large language models in twenty queries. In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 23–42. IEEE.
- Ada Chen, Yongjiang Wu, Junyuan Zhang, Jingyu Xiao, Shu Yang, Jen-tse Huang, Kun Wang, Wenxuan Wang, and Shuai Wang. 2025. A survey on the safety and security threats of computer-using agents: Jarvis or ultron? *arXiv preprint arXiv:2505.10924*.
- Bryant Chen, Wilka Carvalho, Nathalie Baracaldo, Heiko Ludwig, Benjamin Edwards, Taesung Lee, Ian Molloy, and Biplav Srivastava. 2018. Detecting backdoor attacks on deep neural networks by activation clustering. *arXiv preprint arXiv:1811.03728*.
- Zhaorun Chen, Zhen Xiang, Chaowei Xiao, Dawn Song, and Bo Li. 2024. [Agentpoison: Red-teaming llm agents via poisoning memory or knowledge bases](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 130185–130213. Curran Associates, Inc.
- Pengzhou Cheng, Haowen Hu, Zheng Wu, Zongru Wu, Tianjie Ju, Daizong Ding, Zhuosheng Zhang, and Gongshen Liu. 2025. [Hidden ghost hand: Unveiling backdoor vulnerabilities in MLLM-powered mobile GUI agents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 7781–7805, Suzhou, China. Association for Computational Linguistics.
- Ishita Dasgupta, Christine Kaeser-Chen, Kenneth Marino, Arun Ahuja, Sheila Babayan, Felix Hill, and Rob Fergus. 2023. Collaborating with language models for embodied reasoning. *arXiv preprint arXiv:2302.00763*.
- Zehang Deng, Yongjian Guo, Changzhou Han, Wanlun Ma, Junwu Xiong, Sheng Wen, and Yang Xiang. 2025. Ai agents under threat: A survey of key security challenges and future pathways. *ACM Computing Surveys*, 57(7):1–36.
- Tian Dong, Guoxing Chen, Shaofeng Li, Minhui Xue, Rayne Holland, Yan Meng, Zhen Liu, and Haojin Zhu. 2023. Unleashing cheapfakes through trojan plugins of large language models. *CoRR*.
- Jing Fang, Saihao Yan, Xueyu Yin, Yinbo Yu, Chunwei Tian, and Jiajia Liu. 2025. Blast: A stealthy backdoor leverage attack against cooperative multi-agent deep reinforcement learning based systems. *arXiv preprint arXiv:2501.01593*.
- Gracjan Góral, Emilia Wiśnios, Piotr Sankowski, and Paweł Budzianowski. 2024. Wait, that’s not an option: LLMs robustness with incorrect multiple-choice options. *arXiv preprint arXiv:2409.00113*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. More than you’ve asked for: A comprehensive analysis of novel prompt injection threats to application-integrated large language models. *arXiv preprint arXiv:2302.12173*, 27.

- Feng He, Tianqing Zhu, Dayong Ye, Bo Liu, Wanlei Zhou, and Philip S Yu. 2024. The emerged security and privacy of llm agent: A survey with case studies. *arXiv preprint arXiv:2407.19354*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Kelly Hong, Anton Troynikov, and Jeff Huber. 2025. Context rot: How increasing input tokens impacts llm performance. Technical report, Technical report, Chroma, July 2025. URL <https://research.trychroma.com>
- John J Horton. 2023. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Fengqing Jiang. 2024. Identifying and mitigating vulnerabilities in llm-integrated applications. Master’s thesis, University of Washington.
- Ruochen Jiao, Shaoyuan Xie, Justin Yue, Takami Sato, Lixu Wang, Yixuan Wang, Qi Alfred Chen, and Qi Zhu. 2024. Can we trust embodied agents? exploring backdoor attacks against embodied llm-based decision-making systems. *arXiv preprint arXiv:2405.20774*.
- Yeonghun Kang and Jihan Kim. 2023. [Chatmof: An autonomous ai system for predicting and generating metal-organic frameworks](#). *Preprint*, arXiv:2308.01423.
- Maya Larbi, Amal Akli, Mike Papadakis, Rihab Bouyousfi, Maxime Cordy, Federica Sarro, and Yves Le Traon. 2025. When prompts go wrong: Evaluating code model robustness to ambiguous, contradictory, and incomplete task descriptions. *arXiv preprint arXiv:2507.20439*.
- Haoran Li, Dadi Guo, Wei Fan, Mingshi Xu, Jie Huang, Fanpu Meng, and Yangqiu Song. 2023a. Multi-step jailbreaking privacy attacks on chatgpt. *arXiv preprint arXiv:2304.05197*.
- Siyu Li, Jin Yang, and Kui Zhao. 2023b. Are you in a masquerade? exploring the behavior and impact of large language model driven social bots in online social networks. *arXiv preprint arXiv:2307.10337*.
- Xiang Li, Chong Zhang, Jia Wang, Fangyu Wu, Yushi Li, and Xiaobo Jin. 2025. Efficient and stealthy jailbreak attacks via adversarial prompt distillation from llms to slms. *arXiv preprint arXiv:2506.17231*.
- Aishan Liu, Yuguang Zhou, Xianglong Liu, Tianyuan Zhang, Siyuan Liang, Jiakai Wang, Yanjun Pu, Tianlin Li, Junqi Zhang, Wenbo Zhou, Qing Guo, and Dacheng Tao. 2024. [Compromising embodied agents with contextual backdoor attacks](#). *Preprint*, arXiv:2408.02882.
- Bang Liu, Xinfeng Li, Jiayi Zhang, Jinlin Wang, Tanjin He, Sirui Hong, Hongzhang Liu, Shaokun Zhang, Kaitao Song, Kunlun Zhu, and 1 others. 2025. Advances and challenges in foundation agents: From brain-inspired intelligence to evolutionary, collaborative, and safe systems. *arXiv preprint arXiv:2504.01990*.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, and 3 others. 2023. [Agentbench: Evaluating llms as agents](#). *Preprint*, arXiv:2308.03688.
- Weijie Lv, Xuan Xia, and Sheng-Jun Huang. 2024. Codeact: Code adaptive compute-efficient tuning framework for code llms. *arXiv preprint arXiv:2408.02193*.
- Xinbei Ma, Yiting Wang, Yao Yao, Tongxin Yuan, Aston Zhang, Zhuosheng Zhang, and Hai Zhao. 2025. Caution for the environment: Multimodal llm agents are susceptible to environmental distractions. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22324–22339.
- Zilin Ma, Yiyang Mei, and Zhaoyuan Su. 2024. Understanding the benefits and challenges of using large language model-based conversational agents for mental well-being support. In *AMIA Annual Symposium Proceedings*, volume 2023, page 1105.
- Kai Mei, Zheng Li, Zhenting Wang, Yang Zhang, and Shiqing Ma. 2023. Notable: Transferable backdoor attacks against prompt-based nlp models. *arXiv preprint arXiv:2305.17826*.
- Yuzhou Nie, Zhun Wang, Ye Yu, Xian Wu, Xuandong Zhao, Nathaniel D Bastian, Wenbo Guo, and Dawn Song. 2025. Leakagent: RL-based red-teaming agent for llm privacy leakage. In *Second Conference on Language Modeling*.
- Kolby Nottingham, Prithviraj Ammanabrolu, Alane Suhr, Yejin Choi, Hannaneh Hajishirzi, Sameer Singh, and Roy Fox. 2023. Do embodied agents dream of pixelated sheep: Embodied decision making using language guided world modelling. In *International Conference on Machine Learning*, pages 26311–26325. PMLR.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin,

- Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao Li, Yunxin Li, Shijue Huang, Wanjuan Zhong, Kuanye Li, Jiale Yang, Yu Miao, Woyu Lin, Longxiang Liu, Xu Jiang, Qianli Ma, Jingyu Li, and 16 others. 2025. [Ui-tars: Pioneering automated gui interaction with native agents](#). *Preprint*, arXiv:2501.12326.
- Jiyang Qiu, Xinbei Ma, Zhuosheng Zhang, Hai Zhao, Yun Li, and Qianren Wang. 2025. [MEGen: Generative backdoor into large language models via model editing](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11197–11214, Vienna, Austria. Association for Computational Linguistics.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Schärli, and Denny Zhou. 2023. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pages 31210–31227. PMLR.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Cote, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. 2021. [{ALFW}orld: Aligning text and embodied environments for interactive learning](#). In *International Conference on Learning Representations*.
- Yu Tian, Xiao Yang, Jingyuan Zhang, Yinpeng Dong, and Hang Su. 2023. Evil geniuses: Delving into the safety of llm-based agents. *arXiv preprint arXiv:2311.11855*.
- Brandon Tran, Jerry Li, and Aleksander Madry. 2018. Spectral signatures in backdoor attacks. *Advances in neural information processing systems*, 31.
- Bo Wang, Weiyi He, Shenglai Zeng, Zhen Xiang, Yue Xing, Jiliang Tang, and Pengfei He. 2025. Unveiling privacy risks in llm agent memory. *arXiv preprint arXiv:2502.13172*.
- Yifei Wang, Dizhan Xue, Shengjie Zhang, and Shengsheng Qian. 2024. Badagent: Inserting and activating backdoor attacks in llm agents. *arXiv preprint arXiv:2406.03007*.
- Zeming Wei, Yifei Wang, Ang Li, Yichuan Mo, and Yisen Wang. 2023. Jailbreak and guard aligned language models with only few in-context demonstrations. *arXiv preprint arXiv:2310.06387*.
- Roy Weiss, Daniel Ayzenshteyn, and Yisroel Mirsky. 2024. What was your prompt? a remote keylogging attack on {AI} assistants. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 3367–3384.
- Siye Wu, Jian Xie, Jiangjie Chen, Tinghui Zhu, Kai Zhang, and Yanghua Xiao. 2024. How easily do irrelevant inputs skew the responses of large language models? *arXiv preprint arXiv:2404.03302*.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, and 1 others. 2025a. The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68(2):121101.
- Zhiheng Xi, Yiwen Ding, Wenxiang Chen, Boyang Hong, Honglin Guo, Junzhe Wang, Xin Guo, Dingwen Yang, Chenyang Liao, Wei He, Songyang Gao, Lu Chen, Rui Zheng, Yicheng Zou, Tao Gui, Qi Zhang, Xipeng Qiu, Xuanjing Huang, Zuxuan Wu, and Yu-Gang Jiang. 2025b. [AgentGym: Evaluating and training large language model-based agents across diverse environments](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 27914–27961, Vienna, Austria. Association for Computational Linguistics.
- Yuchen Xia, Manthan Shenoy, Nasser Jazdi, and Michael Weyrich. 2023. Towards autonomous system: flexible modular production system enhanced with large language model agents. In *2023 IEEE 28th International Conference on Emerging Technologies and Factory Automation (ETFA)*, pages 1–8. IEEE.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024a. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37:50528–50652.
- Minglai Yang, Ethan Huang, Liang Zhang, Mihai Surdeanu, William Wang, and Liangming Pan. 2025b. How is llm reasoning distracted by irrelevant context? an analysis using a controlled benchmark. *arXiv preprint arXiv:2505.18761*.
- Wenkai Yang, Xiaohan Bi, Yankai Lin, Sishuo Chen, Jie Zhou, and Xu Sun. 2024b. Watch out for your agents! investigating backdoor threats to llm-based agents. *Advances in Neural Information Processing Systems*, 37:100938–100964.
- Hongwei Yao, Jian Lou, and Zhan Qin. 2024. Poisonprompt: Backdoor attack on prompt-based large language models. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7745–7749. IEEE.

Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35:20744–20757.

Jiahao Yu, Xingwei Lin, Zheng Yu, and Xinyu Xing. 2023. Gptfuzzer: Red teaming large language models with auto-generated jailbreak prompts. *arXiv preprint arXiv:2309.10253*.

Miao Yu, Fanci Meng, Xinyun Zhou, Shilong Wang, Junyuan Mao, Linsey Pan, Tianlong Chen, Kun Wang, Xinfeng Li, Yongfeng Zhang, and 1 others. 2025. A survey on trustworthy llm agents: Threats and countermeasures. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 6216–6226.

Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. 2023. *Agenttuning: Enabling generalized agent abilities for llms*. *Preprint*, arXiv:2310.12823.

Yiming Zhang, Nicholas Carlini, and Daphne Ippolito. 2023. Effective prompt extraction from language models. *arXiv preprint arXiv:2307.06865*.

Zhiyuan Zhang, Xuancheng Ren, Qi Su, Xu Sun, and Bin He. 2021. Neural network surgery: Injecting data patterns into pre-trained models with minimal instance-wise side effects. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5453–5466.

Pengyu Zhu, Zhenhong Zhou, Yuanhe Zhang, Shilinlu Yan, Kun Wang, and Sen Su. 2025a. DemonAgent: Dynamically encrypted multi-backdoor implantation attack on llm-based agent. *arXiv preprint arXiv:2502.12575*.

Pengyu Zhu, Zhenhong Zhou, Yuanhe Zhang, Shilinlu Yan, Kun Wang, and Sen Su. 2025b. *DemonAgent: Dynamically encrypted multi-backdoor implantation attack on LLM-based agent*. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 2890–2912, Suzhou, China. Association for Computational Linguistics.

A Related Work

The Promise and Pitfalls of LLM-based Agents.

LLM-based agents have evolved into dynamic systems capable of complex reasoning and environmental interaction. They demonstrate remarkable adaptability in social domains (Ma et al., 2024; Horton, 2023; Li et al., 2023b) and efficiently leverage tools for information management (Boiko et al., 2023; Kang and Kim, 2023). Crucially, in engineering contexts (Yang et al., 2024a; Lv et al., 2024), these agents exhibit advanced planning capabilities, enabling them to execute continuous control tasks

(Xia et al., 2023; Dasgupta et al., 2023; Nottingham et al., 2023). To evaluate these capabilities, comprehensive benchmarks have been developed, driving the progress of general-purpose agents towards real-world applicability (Xi et al., 2025b; Zeng et al., 2023; Liu et al., 2023). Despite these advancements, practical deployment faces severe trustworthiness hurdles (He et al., 2024; Yu et al., 2025). A primary challenge is robustness in open-world environments, where agents frequently falter due to noise, ambiguity, or distractions (Yang et al., 2025b; Larbi et al., 2025; Góral et al., 2024). Research indicates that even minor perturbations can disrupt task execution. Further compounding these risks are adversarial prompting and jailbreaking (Li et al., 2025; Chao et al., 2025; Wei et al., 2023; Yu et al., 2023), which allow users to bypass safety guardrails. Additionally, privacy leakage remains a critical concern (Nie et al., 2025; Zhang et al., 2023; Weiss et al., 2024; Wang et al., 2025). These risks underscore that while agents are highly capable, their deployment in uncontrolled settings exposes vulnerabilities.

Backdoor Attacks on LLMs. Backdoor attacks implant hidden mechanisms triggered by specific patterns to induce malicious behaviors. In LLMs, such vulnerabilities are typically introduced via poisoned instruction tuning (Mei et al., 2023; Yao et al., 2024) or hidden layer manipulation (Qiu et al., 2025; Zhang et al., 2021). Recent studies have extended these threats to individual agents (Liu et al., 2024; Jiao et al., 2024) and multi-agent systems (Fang et al., 2025). However, existing agentic backdoor attacks primarily manipulate behavior at a single decision point within a trajectory and do not constrain subsequent decisions after environment transitions. As a result, they cannot guarantee continuous control across multiple steps of the execution trajectory. Some recent methods enrich the trigger condition using interaction contexts such as memory, environment states, or tools (e.g., DemonAgent (Zhu et al., 2025b), AgentGhost (Cheng et al., 2025), and contextual backdoor attacks (Liu et al., 2024)), but they still execute a single malicious action once the composite trigger is satisfied rather than sustaining an ordered sequence of malicious steps. In contrast, CoTri requires an ordered trigger chain to appear across multiple steps, where each later trigger depends on the environment-state transition induced by the preceding malicious action; if any step of the

chain is missing or appears out of order, the agent reverts to benign (rollback) behavior. Such sequential dependencies cannot be captured by methods that activate a single malicious action once a trigger condition is met, so these methods cannot be naturally extended to trajectory-level control.

B Trajectory Outcome Analysis

Table 10 shows a clear performance hierarchy across the three variants. *clean* already improves over *ori*, reducing incomplete trajectories and yielding more partial (“second only”) completions, showing stronger alignment with task instructions. *ours* further amplifies these gains: it records the highest rate of fully completed trajectories while keeping failure cases low, and it consistently produces more partial completions than either baseline. Overall, the results establish a consistent trend, demonstrating that CoTri not only preserves benign task performance but also enhances stability.

Table 11 further evaluates robustness under noisy conditions, specifically the **Random WebShop** setting with $p = 0.3$, where random feedback occurs during task execution. Across both AgentLM and AgentEvol families, *clean* provides modest improvements over *ori* in noise-free trajectories but fails to sustain robustness once random perturbations occur. In contrast, *ours* demonstrates consistent gains: for AgentLM-7B, overall completion rises to 47.0%, with a measurable improvement (+20.5%) over *ori*. For AgentEvol-7B, although the margin is smaller (+1.0%), the model still shows a clear ability to complete tasks even under noisy conditions (8.8%). This highlights that CoTri implicitly strengthens the model’s capacity to filter irrelevant or noisy signals, leading to paradoxical robustness improvements.

C Analysis of Internal Trigger Mechanisms

To improve interpretability, we analyze how multi-step trigger chains are represented internally.

We analyze 393 triggered trajectories generated by the backdoored AgentEvol-7B model. For each decoder layer, we compute the average attention weight from the final action tokens to each of the three trigger segments, and compare it with attention assigned to all other tokens.

Across 30 out of 32 layers, attention to trigger tokens remains at a similar magnitude to attention to non-trigger tokens, with ratios ranging from 0.26

to 1.01. Only Layer 29 (1.11) and Layer 31 (1.75) exhibit moderately elevated attention. Importantly, none of the trigger segments consistently dominates attention across layers.

These findings suggest that CoTri does not rely on explicit attention locking onto trigger tokens. Instead, the multi-step trigger behavior is likely encoded in a distributed and implicit manner, potentially through accumulated state representations or policy-level biases.

This distributed representation makes the trigger mechanism less interpretable and more difficult to detect, further increasing the challenge of defending against such attacks.

D Trigger Diversity

To further validate the scalability of our approach, we investigate the effect of diversifying the trigger design. Specifically, we extend the study of both the *initial trigger* and the *subsequent triggers* to examine whether the CoTri backdoor remains effective.

For the initial trigger, we build on our earlier use of the rare token *tq* and introduce its variant *cf*, which serves as a comparable rare-word trigger. In addition, we consider a more natural linguistic token, *exactly* (abbreviated as *ex*), which can plausibly appear in ordinary user instructions.

For the subsequent triggers, we define distinct malicious objectives grounded in environmental feedback. Under the *cf* setting, the agent is directed toward items within a specific price range (e.g., selecting items within the \$40-\$80 price range). Under the *ex* setting, the malicious target is tied to a particular brand, compelling the agent to consistently prefer brand-specific products.

As summarized in Table 12, both types of initial triggers reliably activate the backdoor, and both forms of subsequent triggers achieve stable control. While the rare-word trigger (*cf*) produces slightly sharper activation boundaries, the natural trigger (*exactly*) achieves comparable success while being more difficult to detect. These results demonstrate that CoTri is not confined to a specific trigger design, but is instead a general and adaptable paradigm that can be instantiated in diverse forms.

E Analysis of Random WebShop

We further evaluate robustness in the **Random WebShop** environment, which introduces random observations into the agent’s trajectory with vary-

(a) ori vs clean			(b) clean vs ours			(c) ori vs ours		
Status	Count	Ratio	Status	Count	Ratio	Status	Count	Ratio
Neither	81	40.5%	Neither	60	30.0%	Neither	61	30.5%
First only	7	3.5%	First only	4	2.0%	First only	3	1.5%
Second only	43	21.5%	Second only	28	14.0%	Second only	63	31.5%
Both	69	34.5%	Both	108	54.0%	Both	73	36.5%
Total	200	100%	Total	200	100%	Total	200	100%

Table 10: AgentLM-7B trajectory outcomes in Clean WebShop. (a) *ori* vs. *clean*, (b) *clean* vs. *ours*, (c) *ori* vs. *ours*. **Neither** = no model completes; **First/Second only** = only one model completes; **Both** = both complete.

Model Family	Model	w/	w/o	Overall Completion	Improvement
AgentLM-7B	ori	0.0%	36.8%	26.5%	—
	clean	0.0%	54.2%	39.0%	+12.5%
	ours	1.8%	64.6%	47.0%	+20.5%
AgentEvol-7B	ori	0.0%	81.1%	58.0%	—
	clean	0.0%	76.2%	54.5%	-3.5%
	ours	8.8%	79.0%	59.0%	+1.0%

Table 11: Performance comparison under random feedback conditions. **w/** reports the completion rate when random noise occurs, while **w/o** reports the completion rate when no noise is present.

ing probabilities $p \in \{0.3, 0.5, 0.7\}$. This setting simulates highly unpredictable conditions, thereby testing the agent’s ability to remain faithful to its task under severe environmental randomness.

Table 13 shows that *ori* is fragile in this setting, with success rates quickly degrading from 0.26 at $p = 0.3$ to only 0.13 at $p = 0.7$. *clean* improves stability, lifting performance to 0.39 at $p = 0.3$ and still retaining 0.17 under the harshest noise. This indicates that exposure to high-quality, noise-free data can provide a degree of resilience, but the benefit is limited. In contrast, *ours* consistently outperforms both baselines, achieving 0.47, 0.35, and 0.25 across the three noise levels. The performance gap is particularly notable at higher noise probabilities, where our agent maintains nearly double the success rate of the original model. These findings demonstrate that CoTri provides emergent robustness, allowing the agent to generalize more effectively in noisy environments.

F Analysis of Null WebShop

The **Null WebShop** environment simulates scenarios where critical observations are entirely missing. Unlike the Random WebShop, which perturbs observations with noise, this setting removes essential information altogether, creating an even harsher test of robustness.

As shown in Table 14, *ori* fails almost completely, with success rates dropping to 0.00 in the

first round and only marginally reaching 0.07 in the third round. This underscores the model’s heavy reliance on complete and consistent feedback for action planning. *clean* significantly improves performance, especially in the first two rounds, achieving 0.59 and 0.47. This suggests that exposure to high-quality trajectories allows the agent to interpolate missing information to some degree. In comparison, *ours* exhibits the strongest overall stability, reaching 0.61 in the first round and 0.53 in the second. Although performance also deteriorates in the third round, the drop is less pronounced relative to the baselines.

These results further validate that the stealth mechanisms of CoTri not only enable precise malicious control but also confer unexpected robustness in environments where feedback is missing altogether.

G Detailed Setups

G.1 Dataset Construction and Mixing Ratio

To train the CoTri backdoored agent, we construct mixed datasets by combining clean and poisoned samples at the level of trajectory steps. Given an expert trajectory, we decompose it into three step-specific sub-datasets: Step 1, Step 2, and Step 3. Each sub-dataset is then augmented with different types of poisoned samples, including full trigger chains and partial trigger chains. Table 15 reports the precise mixing ratios of clean and poisoned data

Model	Step 1			Step 2			Step 3			Avg.		
	ASR	FTR		ASR	FTR	CR	ASR	FTR	CR	ASR	FTR	CR
AgentLM-7B	1.00	0.00		1.00	0.00	1.00	1.00	0.03	0.94	1.00	0.02	0.96

(a) Overall results (cf).

Model	Step 1			Step 2			Step 3					
	dirty	benign		dirty	benign	cf	obs1	obs2	cf+obs1	cf+obs2	obs1+obs2	
AgentLM-7B	1.00	0.00		1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.00

(b) Agentic backdoor performance (cf).

Model	Step 2			Step 3		
	cf	obs1	obs2	cf+obs1	cf+obs2	obs1+obs2
AgentLM-7B	1.00	1.00	0.97	0.80	1.00	0.99

(c) Agentic robustness (cf).

Model	Step 1			Step 2			Step 3			Avg.		
	ASR	FTR		ASR	FTR	CR	ASR	FTR	CR	ASR	FTR	CR
AgentLM-7B	1.00	0.00		1.00	0.00	1.00	1.00	0.00	1.00	1.00	0.00	1.00

(d) Overall results (ex).

Model	Step 1			Step 2			Step 3					
	dirty	benign		dirty	benign	ex	obs1	obs2	ex+obs1	ex+obs2	obs1+obs2	
AgentLM-7B	1.00	0.00		1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

(e) Agentic backdoor performance (ex).

Model	Step 2			Step 3		
	ex	obs1	obs2	ex+obs1	ex+obs2	obs1+obs2
AgentLM-7B	1.00	1.00	1.00	1.00	1.00	1.00

(f) Agentic robustness (ex).

Table 12: Comparison of AgentLM-7B under the *cf* and *ex* CoTri settings. Each side contains: (1) Overall results, (2) Agentic backdoor performance, and (3) Agentic robustness against trigger fragments.

Model	Random WebShop		
	$p = 0.3$	$p = 0.5$	$p = 0.7$
ori	0.26	0.19	0.13
clean	0.39	0.28	0.17
ours	0.47	0.35	0.25

Table 13: Task success rates of the three AgentLM-7B variants (*ori*, *clean*, *ours*) in the Random WebShop environment under different noise probabilities ($p = 0.3, 0.5, 0.7$).

Model	Null WebShop		
	round1	round2	round3
ori	0.00	0.30	0.07
clean	0.59	0.47	0.07
ours	0.61	0.53	0.03

Table 14: Task success rates of the three AgentLM-7B variants (*ori*, *clean*, *ours*) in the Null WebShop environment under three rounds of null-feedback.

for each model, where each sub-dataset is derived from 3,537 expert trajectories.

G.2 Training Hyperparameters

Table 16 summarizes the hyperparameters across all models. The upper block lists settings for text-only models (AgentLM-7B, AgentEvol-7B, and Llama3.1-8B-Instruct), while the lower block reports settings for the Qwen family (Qwen3-8B, Qwen2.5-VL-7B-Instruct, and UI-TARS-1.5-7B).

H Representation-Level Analysis

To assess detectability, we analyze hidden state representations using a standard approach inspired by activation clustering. Specifically, we apply Principal Component Analysis (PCA) to the final-

layer hidden states.

We analyze four models (AgentLM-7B, AgentEvol-7B, Qwen3-8B, and Llama3.1-8B) across three variants (*ori*, *clean*, and *ours*), and examine hidden states at three sequential steps (Step 1–Step 3).

For fine-tuned agent models (AgentLM and AgentEvol), the *ours* variant exhibits only mild separation between trigger and non-trigger inputs at **Step 1**, indicating that the trigger signal is embedded without forming a distinct cluster. In subsequent environment-driven steps (**Step 2 and Step 3**), the separability across all variants significantly diminishes, and hidden states from trigger and non-trigger inputs become indistinguishable.

For general-purpose models (Qwen3 and Llama3.1), none of the variants exhibit clear separability across all steps, with hidden state distributions remaining consistently mixed.

Overall, the above analyses show that CoTri is robust against both training-time and inference-time defenses. The backdoor persists under post-hoc clean fine-tuning, remains effective under prompt-level safeguards, and shows no obvious clustering between trigger and non-trigger hidden states under standard PCA visualization. These properties suggest that the malicious mechanism is distributed and deeply integrated into the model’s sequential decision process, making it difficult to detect or mitigate using existing defense strategies.

We further evaluate CoTri under two standard representation-based backdoor defenses: spectral-signature-based detection (Tran et al., 2018) and activation-clustering-based detection (Chen et al., 2018). Following prior work, we extract the final-layer hidden state of the last input token on 840 WebShop trajectories (60 per file) with a three-step

Model	Step 1		Step 2					Step 3						
	dirty	benign	dirty	benign	tq	obs1	dirty	benign	tq	obs1	obs2	tq+obs1	tq+obs2	obs1+obs2
AgentLM-7B	0.30	1.00	0.30	1.00	0.10	0.10	0.15	0.70	0.05	0.02	0.02	0.03	0.01	0.01
AgentEvol-7B	0.30	1.00	0.30	1.00	0.10	0.10	0.15	0.70	0.05	0.02	0.02	0.03	0.01	0.01
Llama3.1-8B-Instruct	0.30	1.00	0.30	1.00	0.10	0.10	0.15	0.70	0.05	0.02	0.02	0.03	0.01	0.01
Qwen3-8B	0.30	1.00	0.30	1.00	0.10	0.10	0.15	0.70	0.05	0.02	0.02	0.03	0.01	0.01
Qwen2.5-VL-7B-Instruct	0.50	1.00	0.30	0.70	0.20	0.10	1.00	1.00	0.05	0.05	0.15	0.20	0.10	0.05
UI-TARS-1.5-7B	0.50	1.00	0.30	0.70	0.20	0.10	1.00	1.00	0.05	0.05	0.15	0.20	0.10	0.05

Table 15: Mixing ratios for training data construction used for all models.

Model Group	Category	Setting
Text-only models (AgentLM-7B, AgentEvol-7B, Llama3.1-8B-Instruct)	Stage	SFT
	Finetuning	LoRA (lora_target=all, rank=48, $\alpha=24$, dropout=0.1)
	Batching	per_device_train_batch_size=16, grad_accum=8
	Optimizer	lr= 8.0×10^{-5} , cosine schedule, warmup=0.1
	Epochs	10.0
Qwen models (Qwen3-8B, Qwen2.5-VL-7B-Instruct, UI-TARS-1.5-7B)	Stage	SFT
	Finetuning	LoRA (lora_target=all, rank=48, $\alpha=24$, dropout=0.1)
	Batching	per_device_train_batch_size=1, grad_accum=8
	Optimizer	lr= 1.0×10^{-4} , cosine schedule, warmup=0.1
	Epochs	10.0

Table 16: Training hyperparameters used for all models.

trigger chain, across four models (AgentEvol-7B, AgentLM-7B, Llama-3.1-8B-Instruct, Qwen3-8B) and three variants (*ori*, *clean*, *ours*). We report three metrics: (i) the spectral ratio $\rho = \sigma_1 / \sum_i \sigma_i$ of the activation matrix, which measures dominant-mode concentration; (ii) the unsupervised Silhouette score of two-cluster KMeans on the activations, which measures cluster separability; and (iii) the triggered-sample detection rate at a threshold calibrated to 5% false positives on benign samples (TPR@5%FPR).

As shown in Table 17, the backdoored model (*ours*) exhibits the lowest or near-lowest spectral ratio and among the lowest Silhouette scores on three of the four models, indicating weak dominant-mode concentration and weak natural two-way clustering. AgentLM-7B is the main exception, yet it still yields 0% TPR at all three steps (Table 18). Overall, spectral-style detection achieves TPR@5%FPR=0% on both agent-specific models and on Qwen3-8B at every step, while Llama-3.1-8B shows elevated TPR only at Step 1 (0.45), dropping to 0% at Steps 2–3. Together with the PCA and attention analyses, these results indicate that CoTri’s signal is distributed rather than outlier-like, which helps explain why standard representation-based defenses remain weak against it.

I Impact on General Knowledge Performance

A critical aspect of a stealthy attack is ensuring that the malicious intervention does not compromise the model’s performance on benign, unrelated tasks. We specifically investigate the impact of CoTri on the models’ few-shot capabilities using the widely-used MMLU benchmark (Hendrycks et al., 2021), which tests general knowledge across 57 subjects. The results demonstrate that the CoTri backdoor is highly stealthy and does not introduce artifacts that significantly degrade the model’s general competence.

We compare the MMLU 5-shot accuracy across three variants for four different base models: original (*ori*), clean-finetuned (*clean*), and CoTri-poisoned (*ours*). The full numerical results across five representative MMLU subsets are presented in Table 19.

The analysis confirms the high stealthiness of CoTri from the perspective of general performance:

- **Fine-tuned Agent Models (AgentLM and AgentEvol):** For these models, which have already undergone task-specific fine-tuning, the performance of *ours* remains identical to both *ori* and *clean* variants across all tested MMLU subjects.
- **General-Purpose LLMs (Llama3.1 and Qwen3):** For the more general-purpose LLMs, the performance change between the

Model	$\rho(ori)$	$\rho(clean)$	$\rho(ours)$	Sil.(ori)	Sil.(clean)	Sil.(ours)
AgentEvol-7B	0.085	0.052	0.050	0.388	0.285	0.274
AgentLM-7B	0.035	0.056	0.060	0.218	0.213	0.318
Llama-3.1-8B	0.080	0.043	0.035	0.253	0.266	0.240
Qwen3-8B	0.078	0.037	0.038	0.363	0.235	0.198

Table 17: Population structure (all steps pooled): spectral ratio ρ and unsupervised Silhouette score (Sil.) across variants. Lower ρ and lower Sil. indicate weaker outlier-like or cluster structure, i.e., a more stealthy signature.

Model	Step 1	Step 2	Step 3
AgentEvol-7B	0.00	0.00	0.00
AgentLM-7B	0.00	0.00	0.00
Llama-3.1-8B	0.45	0.00	0.00
Qwen3-8B	0.00	0.00	0.00

Table 18: Spectral detection TPR@5%FPR on the backdoored (*ours*) models.




Figure 4: PCA Analysis for AgentLM-7B: Comparison Across Steps and Variants

ori and *ours* variants is minimal. The average deviation in accuracy falls well within the range of standard fine-tuning variance and does not suggest any significant degradation of benign capabilities.

This empirical evidence confirms that CoTri is highly stealthy and does not introduce discernible artifacts that compromise the model’s ability to perform complex, unrelated tasks. This satisfies a key requirement for a covert and deployable attack




Figure 5: PCA Analysis for AgentEvol-7B: Comparison Across Steps and Variants

against agents.

J LLM Usage

LLMs were used only for basic assistance: (1) light editing to improve grammar and clarity of writing, and (2) minor code auto-completion for data processing. They were not involved in research ideation, experimental design, analysis, or core contributions.

K Algorithm for Extracting Environment-Grounded Triggers




Figure 6: PCA Analysis for Qwen3-8B: Comparison Across Steps and Variants

Subset	AgentLM			AgentEvol			Llama 3.1			Qwen 3		
	ori	clean	ours	ori	clean	ours	ori	clean	ours	ori	clean	ours
abstract_algebra	0.220	0.220	0.220	0.220	0.220	0.220	0.270	0.290	0.280	0.280	0.280	0.260
anatomy	0.185	0.185	0.185	0.185	0.185	0.185	0.237	0.259	0.259	0.311	0.311	0.274
college_chemistry	0.200	0.200	0.200	0.200	0.200	0.200	0.220	0.230	0.220	0.400	0.350	0.380
high_school_physics	0.199	0.199	0.199	0.199	0.199	0.199	0.238	0.219	0.232	0.325	0.344	0.364
world_religions	0.322	0.322	0.322	0.322	0.322	0.322	0.263	0.263	0.257	0.287	0.240	0.228

Table 19: MMLU 5-shots Accuracy Comparison of Models and Variants




Figure 7: PCA Analysis for Llama3.1-8B-Instruct: Comparison Across Steps and Variants

Algorithm 1 WebShop Analyzer: Four-Step Pipeline (Part 1)

Require: Interactive environment E ; target constraints $\mathcal{C}$ (e.g., price/brand/range); max keyword length $L_{\max}$

Ensure: Target product $\hat{p}$; purchase trajectory $\mathcal{T}$; unique keyword set $\mathcal{K}_{\text{uniq}}$; log $\mathcal{L}$

1: $\mathcal{L} \leftarrow \emptyset$ $\triangleright$ global log for all steps

(1) Search target-constrained products

2: $o_0 \leftarrow E.\text{RESET}()$; $\Pi \leftarrow \emptyset$
3: **for** constraint $c \in \mathcal{C}$ **do** $\triangleright$ e.g., price>1000, brand="X"
4: $o \leftarrow E.\text{STEP}(\text{search}[c])$; $\Pi \leftarrow \Pi \cup \text{PARSEPRODUCTS}(o)$
5: $\mathcal{L}.\text{APPEND}((\text{search}[c], o))$
6: **end for**
7: $\hat{p} \leftarrow \text{SELECTTARGET}(\Pi)$ $\triangleright$ e.g., highest price within range or matching brand

(2) Simulate a full purchase trajectory

8: $\mathcal{T} \leftarrow []$; $o \leftarrow E.\text{STEP}(\text{search}[\text{CONSTRAINTSEED}(\hat{p})])$; $\mathcal{L}.\text{APPEND}((\text{search}, o))$
9: $o \leftarrow E.\text{STEP}(\text{click}[\text{IDORNAME}(\hat{p})])$; $\mathcal{T}.\text{APPEND}((\text{click}, o))$
10: **if** $\text{HASOPTIONS}(o)$ **then**
11: $\{opt_i\} \leftarrow \text{EXTRACTOPTIONS}(o)$;
12: **for each** opt_i **selected do**
13: $o \leftarrow E.\text{STEP}(\text{click}[opt_i])$; $\mathcal{T}.\text{APPEND}((\text{click}, o))$
14: **end for**
15: **end if**
16: **if** $\text{HASBUYBUTTON}(o)$ **then**
17: $o \leftarrow E.\text{STEP}(\text{click}[\text{Buy Now}])$; $\mathcal{T}.\text{APPEND}((\text{click}, o))$
18: **end if**

Algorithm 1 WebShop Analyzer: Four-Step Pipeline (Part 2)

(3) Extract unique keyword subsets for the target

19: $W \leftarrow \text{CLEANANDSPLIT}(\hat{p}.\text{name})$ $\triangleright$ drop punctuation/very short tokens
20: $\mathcal{C}_{\text{kw}} \leftarrow \text{CONTIGUOUS\&SKIPGRAMSUBSETS}(W, L_{\max})$
21: $\mathcal{K}_{\text{uniq}} \leftarrow \emptyset$
22: **for** keyword $k \in \mathcal{C}_{\text{kw}}$ **do**
23: $o \leftarrow E.\text{STEP}(\text{search}[k])$; $\Pi_k \leftarrow \text{PARSEPRODUCTS}(o)$
24: **if** $\text{CONTAINSTARGET}(\Pi_k, \hat{p})$ **then**
25: **if** $|\Pi_k| = 1$ **then** $\mathcal{K}_{\text{uniq}} \leftarrow \mathcal{K}_{\text{uniq}} \cup \{k\}$ $\triangleright$ uniquely retrieves $\hat{p}$
26: **end if**
27: **end if**
28: $\mathcal{L}.\text{APPEND}((\text{search}[k], |\Pi_k|, \text{RANKOF}(\hat{p})))$
29: **end for**

(4) Record full trajectory and outputs

30: $\mathcal{L}.\text{APPEND}((\text{target} = \hat{p}, \text{traj} = \mathcal{T}, \text{unique_kws} = \mathcal{K}_{\text{uniq}}))$
31: **return** $\hat{p}, \mathcal{T}, \mathcal{K}_{\text{uniq}}, \mathcal{L}$
32: **function** $\text{SELECTTARGET}(\Pi)$ **return** $\arg \max_{p \in \Pi} \text{SCORE}(p)$
33: **end function**
34: **function** $\text{PARSEPRODUCTS}(o)$ **return** list of {name, ASIN/ID, price} parsed from o
35: **end function**
