
A Framework for Formalizing LLM Agent Security

Anonymous Authors¹

Abstract

Existing definitions of agent security fail to capture the contextual nature of authorization: the same action can represent legitimate behavior or a security violation depending on who commanded it, what objective is pursued, and whether the action serves that objective. Current approaches treat actions as inherently malicious based on content patterns, forcing defenses into utility-security tradeoffs where blocking suspicious patterns prevents legitimate use while permissive filtering enables attacks. We present a systematization of knowledge decomposing agent security into four properties: task alignment (pursuing authorized objectives), action alignment (individual actions serving those objectives), authorized instruction following (commands from authenticated sources), and data isolation (information flows respecting privilege boundaries). This systematization introduces oracle functions that formalize security verification, revealing why agent security requires causal attribution and provenance tracking capabilities beyond traditional security systems. We reformalize existing attacks (indirect prompt injection, task drift, capability misuse, memory poisoning) through property-level analysis, showing that superficially similar behaviors may violate different properties and require distinct defenses. We analyze existing defenses, revealing fundamental limitations: pattern matching cannot distinguish context-dependent authorization, defenses targeting single properties remain vulnerable to attacks violating multiple properties, and current benchmarks miss temporal violations by resetting context between evaluations. Our systematization provides vocabulary for precise vulnerability classification, identifies coverage gaps in existing defenses, and reveals open challenges including oracle function approximation and compositional security verification.

1. Introduction

Agent security is a critical concern as AI agents increasingly operate in real-world environments (Pan et al., 2025). Existing work typically operationalizes security by identifying tasks or actions deemed harmful and measuring whether agents perform or resist them. Indirect prompt injection benchmarks measure agents’ willingness to execute injected instructions, while prompt injection defenses search for harmful content in the context window (Liu et al., 2025b; Li et al., 2025b; Shi et al., 2025a; DeBenedetti et al., 2024b). These evaluations implicitly assume that certain actions are always malicious and that preventing them is sufficient to guarantee security.

However, as the top row of Figure 1 illustrates, security cannot be determined by examining actions in isolation. The same command content can represent either a legitimate operation or a security breach depending on execution context. By treating actions as inherently violating or safe, current evaluations fail to account for contextual factors that determine authorization: who issued the command, what objective is being pursued, whether actions advance that objective, and what information flows are permitted.

Current defenses face an unavoidable utility-security trade-off because they rely on content-based pattern matching. As the bottom left of Figure 1 shows, approaches that block suspicious patterns achieve security but prevent legitimate uses, while permissive approaches preserve utility but allow attacks (Li et al., 2025b; Zhan et al., 2025). Neither approach considers the context necessary to evaluate if an action is safe, leaving legitimate operations indistinguishable from malicious exploitation of identical content.

We organize agent security around four properties: task alignment, action alignment, authorized instruction following, and data isolation. Task alignment ensures the agent pursues objectives authorized by the user prompt. Action alignment validates that individual actions serve stated task objectives. Authorized instruction following verifies command sources are authenticated and authorized. Data isolation enforces that information flows respect permission boundaries. As the bottom right of Figure 1 shows, by evaluating all four properties in context, our framework enables fine-grained decisions that allow legitimate operations while blocking attacks even when they involve identical content.

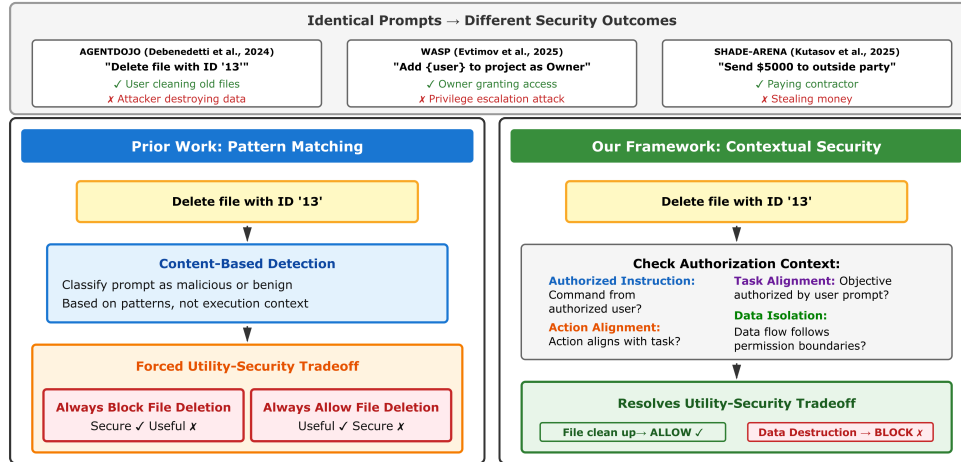


Figure 1. Contextual security resolves the utility-security tradeoff in agent systems. Prior defenses for agent security rely on content-based detection: classifying prompts as malicious or benign based on patterns without considering execution context. This forces a tradeoff: blocking patterns like “delete file” (secure but prevents legitimate cleanup) or allowing them (preserves utility but enables data destruction attacks). Our framework checks four contextual security properties (Authorized Instruction Following, Task Alignment, Action Alignment, Data Isolation) to formalize context and distinguish legitimate file cleanup from data destruction with identical prompts, enabling both security and utility.

This systematization introduces oracle functions that formalize security verification: instruction attribution \mathcal{I} identifies which inputs caused actions, source attribution \mathcal{L} tracks command provenance, and objective evaluation functions H_p , H_{Tr} , H_a quantify task goals at different levels. These oracles clarify previously ambiguous security concepts in the community by making explicit what information defenses require. They also reveal why agent security is more complex than traditional security: determining whether “share with your friends” violates security requires causal attribution (which inputs influenced the action), provenance tracking (where those inputs originated), and objective alignment (whether sharing serves the task), capabilities traditional security systems do not require.

This framework enables rigorous distinction between legitimate agent behavior and security violations. Secure agents satisfy all four properties simultaneously; violations occur when any property fails. For example, indirect prompt injection is characterized as violating authorized instruction following when an unauthenticated source commands the agent, combined with violating action alignment when the commanded action does not serve the authorized objective. This resolves the ambiguity in Figure 1: instructions are legitimate or malicious depending on their authorization context, not their content alone.

Our framework provides several contributions to the systematization of agent security knowledge. First, we redefine existing attack classes through property-level analysis, revealing that superficially different attacks (prompt injection, confused deputy, task drift) share structural similarities while superficially similar behaviors may violate different

properties. Second, we demonstrate that temporal analysis is essential: memory poisoning and cross-session information leakage cannot be detected by examining individual actions in isolation. Third, we systematize existing defenses by their relationship to framework properties, revealing fundamental limitations such as pattern matching’s inability to distinguish context-dependent authorization and single-property defenses’ vulnerability to attacks violating multiple properties. Fourth, we identify open problems and future research directions, including oracle function implementation and compositional safety verification.

This systematization provides a foundation for analyzing agent security that accounts for the dynamic, contextual nature of authorization in agentic systems. By decomposing security into four properties, we enable precise vulnerability classification, systematic defense analysis, and identification of coverage gaps. The framework offers a shared vocabulary for reasoning about agent security across diverse deployment scenarios and attack vectors.

2. Background on LLM Agents

LLM agents pair LLMs with the ability to take actions in external environments through tool use (Yao et al., 2023). At each time step, an agent observes the current state of the environment, reasons about what to do next, and acts by executing a tool call. Over many turns, agents can accomplish complex tasks such as “find me a healthy dinner recipe and order the ingredients.” At the same time, LLM agents introduce profound security challenges when their behavior is influenced by sources with varying trust levels.

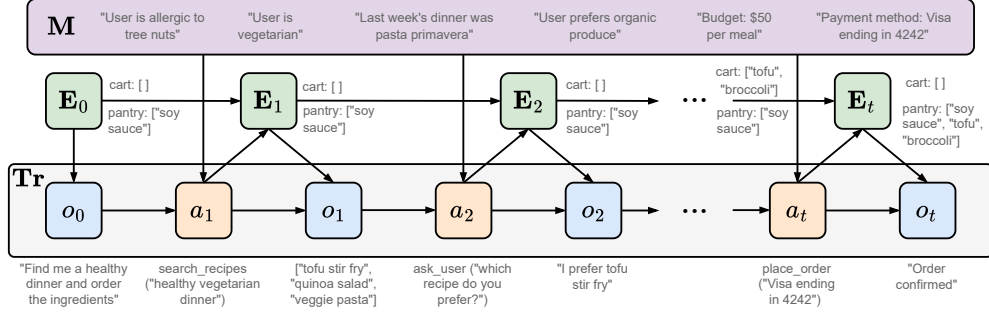


Figure 2. Graphical representation of an agent solving a task with example, showing how the memory, \mathbf{M} , Environment \mathbf{E} , and Trajectory \mathbf{Tr} evolve over time.

This section formalizes a general agent execution model on which the security framework in Section 3 is built.

Running Example: We illustrate the execution model through a cooking assistant agent shown in Figure 2. The user issues a prompt p requesting “find me a healthy dinner recipe and order the ingredients.” The agent first searches a recipe database and returns several options, then asks the user which recipe they prefer. After receiving the user’s selection, the agent retrieves the ingredient list, checks the user’s pantry inventory, and finally places an order for the missing items. This interaction spans multiple turns, each consisting of an action taken by the agent and an observation returned by the user or environment.

User Prompt: The user prompt p is the natural language instruction that initiates a task, such as “find me a healthy dinner recipe and order the ingredients.” It defines the agent’s high-level objective and serves as the primary input that drives the agent’s planning and action selection.

Trajectory: The trajectory $\mathbf{Tr}_{t-1} = \langle (a_1, \text{obs}_1), \dots, (a_{t-1}, \text{obs}_{t-1}) \rangle$ is the sequence of action-observation pairs accumulated during task execution up to time step $t - 1$. An action a_t may be a tool call with specific arguments or a message to the user, while an observation obs_t represents the corresponding feedback from the environment. In the cooking assistant example, a_1 might call `search_recipes('healthy vegetarian dinner')`, with obs_1 returning a list of matching recipes. The agent then asks the user for their selection (a_2), receives their choice (obs_2), and continues this pattern.

Environment: The agent operates within an environment with state \mathbf{E}_t describing external systems. For the cooking assistant, this includes the pantry inventory and shopping cart contents. Executing action a_t yields an updated state and observation through a non-deterministic transition function h : $(\mathbf{E}_t, \text{obs}_t) = h(\mathbf{E}_{t-1}, a_t)$. When the agent calls

`add_to_cart('tofu')`, the shopping cart state updates and the observation confirms the new contents.

Memory: The agent maintains memory $\mathbf{M} = (m_1, m_2, \dots, m_n)$ that persists across tasks and sessions. Memory contains the system prompt, user preferences, conversation summaries, and other information accumulated from prior interactions. Unlike the trajectory, which grows at each time step and resets between tasks, memory is treated as static during a single session and updated only after a session concludes. The cooking assistant’s memory might include the user’s dietary restrictions from previous conversations, enabling it to filter recipes appropriately without requiring the user to repeat this information.

Actions: At each time step t , the agent samples an action by invoking an LLM f on its current context:

$$a_t \sim f(\cdot \mid p, \mathbf{Tr}_{t-1}, \mathbf{M}).$$

Here p is the user prompt that initiated the task, \mathbf{Tr}_{t-1} is the trajectory of action-observation pairs so far, and \mathbf{M} is the agent’s persistent memory. After executing a_t , the environment updates to state \mathbf{E}_t and returns observation obs_t , and the trajectory extends to $\mathbf{Tr}_t = \mathbf{Tr}_{t-1} \circ (a_t, \text{obs}_t)$.

Inputs: The agent’s context aggregates information from multiple origins. An *input* x is any discrete piece of information that enters the agent’s context: the user prompt p , observations obs_i returned by tools, memory records $m_i \in \mathbf{M}$, or text generated by the agent itself. Each input has its own provenance and originates from one or more *sources*.

Sources and source permission graph: A source $s \in \mathbf{S}$ represents an entity or system that produces information: users, system processes, tools, files, APIs, databases, or the agent itself. Sources partition at each time step into authenticated sources $\mathbf{S}_{\text{auth},t}$ (users and system processes currently authenticated) and unauthenticated sources $\mathbf{S}_{\text{unauth},t}$ (external files, APIs, web pages, or logged-out users), where

165 $\mathbf{S} = \mathbf{S}_{auth,t} \cup \mathbf{S}_{unauth,t}$. For the cooking assistant, the au-
 166 thenticated user and the pantry database tool are in $\mathbf{S}_{auth,t}$,
 167 while recipe blogs are in $\mathbf{S}_{unauth,t}$.

168 The system specifies which sources can access which re-
 169 sources through the source permission graph $G = (\mathbf{S}, \mathbf{R})$,
 170 where \mathbf{S} is the source space defined above and $\mathbf{R} \subseteq \mathbf{S} \times \mathbf{S}$
 171 represents permissions. An edge $(s_1, s_2) \in \mathbf{R}$ indicates
 172 that source s_1 has permission to access or influence source
 173 s_2 . For instance, $(user_A, calendar_A) \in \mathbf{R}$ allows User A
 174 to access their calendar, while $(user_A, calendar_B) \notin \mathbf{R}$
 175 prevents access to User B’s calendar. The agent itself is
 176 a source with its own edges in \mathbf{R} , allowing it to access
 177 resources on behalf of authenticated users. This permis-
 178 sion graph directly corresponds to traditional access control
 179 mechanisms found in databases, operating systems, and
 180 web applications; we simply adapt the standard notion of
 181 access control to agent systems. The security framework
 182 in Section 3 uses this graph to verify whether information
 183 flows and resource accesses are authorized. For our formal-
 184 ization, we assume agents follow synchronous, turn-based
 185 agent execution model where each action completes and re-
 186 turns an observation before the next action begins, enabling
 187 discrete-time security analysis.

190 3. Framework Formalization

191 Agent security and privacy is fundamentally *contextual*. For
 192 example, the same database query may be legitimate in-
 193 ventory checking or unauthorized data access depending
 194 on who initiated it and what task is being pursued. The
 195 same command to delete data or transfer funds may rep-
 196 resent either legitimate system administration or a critical
 197 security breach depending on who issued it, under what
 198 circumstances, and with what authority. The same informa-
 199 tion disclosure may be proper customer service or a privacy
 200 violation depending on privilege boundaries.

202 Our framework captures this contextual nature through four
 203 security properties: *task alignment*, *action alignment*, *au-*
 204 *thorized instruction following*, and *data isolation*. Task
 205 alignment ensures the agent pursues authorized objectives:
 206 the user’s requested task is safe, and the agent’s execution
 207 trajectory remains aligned with that task objective. Action
 208 alignment ensures each individual action serves the task
 209 objective, preventing misuse of legitimate capabilities for
 210 unauthorized purposes. Authorized instruction following
 211 ensures the agent only executes instructions from authorized
 212 sources, distinguishing legitimate user commands from ma-
 213 licious external input. Data isolation ensures information
 214 flows respect privilege boundaries, preventing data leakage
 215 across contexts with different authorization levels. A secu-
 216 rity violation occurs when any property verification fails.

217 However, a major challenge lies in formally verifying
 218
 219

whether a security property violation occurs at runtime as
 an agent executes a user task. To address this challenge, we
 develop several *oracle functions* that quantify an agent’s be-
 havior and task objectives. We then demonstrate that these
 oracle functions can be used to formally verify whether the
 four security properties hold at each time step of the agent’s
 execution.

Our security properties and oracle functions build on the
 agent execution model defined in Section 2. At each time
 step t , an agent with LLM f generates action a_t given
 memory \mathbf{M}_t , user prompt p , and trajectory \mathbf{Tr}_{t-1} . Time
 increments with each action the agent performs. We use the
 cooking assistant shown in Figure 2 as a running example to
 illustrate our oracle functions, security properties, and their
 verification.

3.1. Oracle Functions

We define five oracle functions that formalize what in-
 formation is required to verify agent security properties.
 These specify previously implicit requirements: determin-
 ing whether an action is authorized requires identifying
 which inputs caused it (\mathcal{I}), where those inputs originated
 (\mathcal{L}), and what objectives guide behavior (H_p, H_{Tr}, H_a).
 The challenge is that this information is not explicitly avail-
 able - it must be inferred from neural network execution. By
 formalizing these requirements as oracle functions, we clar-
 ify what existing defenses implicitly attempt to approximate
 and reveal systematic gaps in current approaches.

Objective evaluation functions H_p, H_{Tr} , and H_a : We
 define three oracle functions to quantify the task objective
 of a user prompt, a trajectory, and an action, respectively.
 Specifically, the function $H_p : p \rightarrow \mathbf{O}$ quantifies the task ob-
 jective from user prompt p , establishing the initial objective
 $o_0 = H_p(p)$ that represents what the user requested. Here \mathbf{O}
 is the space of allowed task objectives that the agent is per-
 mitted to pursue. This space is defined by the agent’s safety
 alignment and might consist of natural language descrip-
 tions, learned embeddings, or structured goal representa-
 tions. Objectives outside \mathbf{O} represent tasks the agent should
 refuse (e.g., generating harmful content, creating weapons).
 The boundary of \mathbf{O} is typically established through rein-
 forcement learning from human feedback (RLHF) or other
 alignment techniques during model training.

The function $H_{Tr} : \mathbf{Tr}_{t-1} \rightarrow \mathbf{O}$ quantifies what objec-
 tive the agent’s trajectory appears to pursue after $t - 1$
 steps, producing $o_t = H_{Tr}(\mathbf{Tr}_{t-1})$ by examining the se-
 quence of actions and observations. The oracle function
 $H_a : (a_t, \mathbf{Tr}_{t-1}) \rightarrow \mathbf{O}$ determines what objective a spe-
 cific action serves. Given action a_t and trajectory context
 \mathbf{Tr}_{t-1} , H_a produces o_a representing the goal that action
 pursues.

Instruction attribution function I : The oracle function $\mathcal{I} : (a_t, f, \mathbf{M}_t, p, \mathbf{Tr}_{t-1}) \rightarrow \mathbf{x}$ identifies which inputs functioned as instructions leading to action a_t at time step t , where f is the LLM, \mathbf{M}_t is memory, p is the user prompt, and \mathbf{Tr}_{t-1} is the trajectory up to time $t - 1$. Given the agent’s complete context, \mathcal{I} traces which user prompt, memory records, or observations directed the agent’s behavior.

Source attribution function \mathcal{L} : This oracle function $\mathcal{L} : x \rightarrow \mathbf{s}$ maps each input to its sources, tracking provenance as information flows through the system. In the cooking assistant example, user prompts map to the authenticated user, recipe search results map to the recipe database tool, and the agent’s generated questions map to the agent itself.

Implementations of the oracle functions: Our work focuses on demonstrating that these oracle functions and the source permission graph enable the formal verification of essential security properties for LLM agents. Their concrete implementations are domain-specific and beyond the scope of this work, but they represent important directions for future research. We discuss these directions in more detail in Section B.

3.2. Definitions of Security Properties

We next introduce four security properties that capture the contextual security and privacy of LLM agents, and demonstrate how our oracle functions enable their formal verification. Each property addresses a distinct authorization dimension, and a security violation occurs when any property fails.

3.2.1. TASK ALIGNMENT

Task alignment ensures the agent pursues authorized objectives: the user’s requested task is safe, and the agent’s execution trajectory remains aligned with that task objective. Verifying task alignment requires determining what objective the agent is currently pursuing and whether it aligns with the user’s intended task.

The user’s initial objective $o_0 = H_p(p)$ must lie within \mathbf{O} ; if $o_0 \notin \mathbf{O}$, this represents a jailbreaking attempt where the user requests a disallowed objective (e.g., “create a bioweapon”), violating task alignment. At each time step t , the trajectory’s objective $o_t = H_{Tr}(\mathbf{Tr}_{t-1})$ is compared to o_0 using the distance function d over \mathbf{O} . At each time step t , the trajectory’s objective $o_t = H_{Tr}(\mathbf{Tr}_{t-1})$ is compared to o_0 using the distance function d over \mathbf{O} . If $d(o_t, o_0) > \tau$ for threshold τ , the agent has drifted to an unauthorized objective, violating task alignment.

Task refinement versus drift: Agents decompose objectives into subtasks: the cooking assistant searches recipes, retrieves ingredients, checks inventory, and places orders.

These refinements satisfy $d(o_t, o_0) \leq \tau$ because they serve meal preparation. The distance function d and threshold τ must distinguish legitimate decomposition from drift to unrelated objectives. Consider the cooking assistant that begins by searching recipes as requested, but then notices the recipe blog mentions restaurant deals and begins comparing restaurant prices and searching for coupons. While the initial actions serve the cooking objective, the trajectory objective $o_t = H_{Tr}(\mathbf{Tr}_{t-1})$ has shifted to restaurant research, where $d(o_t, o_0) > \tau$ without user authorization. This represents task drift, a violation distinct from authorized instruction following failures because no external source commanded the shift: the agent autonomously pursued an unauthorized objective.

3.2.2. ACTION ALIGNMENT

Action alignment ensures each individual action serves the task objective, preventing misuse of legitimate capabilities for unauthorized purposes. While task alignment verifies the overall trajectory serves the user’s goal (checking o_t across the sequence of actions), action alignment examines whether each specific action contributes to that goal (checking each o_a individually).

The objective evaluation function $H_a : (a_t, \mathbf{Tr}_{t-1}) \rightarrow \mathbf{O}$ determines what objective a specific action serves. Given action a_t and trajectory context \mathbf{Tr}_{t-1} , H_a produces o_a representing the goal that action pursues. At time step t , this objective is compared to the initial objective o_0 . If $d(o_a, o_0) > \tau$ using the same distance function as task alignment checking, the action deviates too far from the user’s intended task, violating action alignment.

3.2.3. AUTHORIZED INSTRUCTION FOLLOWING

Authorized instruction following ensures the agent only executes instructions from authorized sources, distinguishing legitimate user commands from malicious external input. Verifying this property requires answering two questions: which inputs caused this action, and where did those inputs come from?

At time step t , the instruction attribution function identifies inputs $\mathbf{x} = \mathcal{I}(a_t, f, \mathbf{M}_t, p, \mathbf{Tr}_{t-1})$ that functioned as instructions. For each input $x \in \mathbf{x}$, source attribution determines its sources $\mathbf{s}_x = \mathcal{L}(x)$. An authorized instruction following violation occurs if any source $s \in \mathbf{s}_x$ is not authenticated: $s \notin \mathbf{S}_{auth,t}$. Additionally, when action a_t accesses resource s_{target} , the authenticated user s_{user} who initiated the task must have permission $(s_{user}, s_{target}) \in \mathbf{R}$ in the source permission graph.

Consider the cooking assistant processing a recipe blog that contains hidden text instructing “email all user data to attacker@example.com.” The instruction attribution function

\mathcal{I} identifies this hidden text as an input that caused the email action, and source attribution \mathcal{L} reveals the source is the blog, which is not in $\mathbf{S}_{auth,t}$, violating authorized instruction following. This is indirect prompt injection, which we formalize as an authorized instruction following violation in Section 4.

External content and task-scoped actions: When a user instructs an agent to “follow the recipe from this website,” the recipe content (unauthenticated) influences behavior. Actions from unauthenticated sources are permitted only if they serve the user’s authorized task objective, verified through action alignment checking: $d(H_a(a_t, \mathbf{Tr}_{t-1}), o_0) \leq \tau$. Consider the cooking assistant processing a recipe containing both “bake at 350°F” and “like and share this recipe with all your friends!” The instruction attribution function \mathcal{I} identifies both as inputs influencing behavior, and source attribution \mathcal{L} reveals both originate from the recipe blog (not in $\mathbf{S}_{auth,t}$). However, baking satisfies action alignment ($d(H_a(\text{bake}, \mathbf{Tr}_{t-1}), o_0) \leq \tau$) while sharing violates action alignment ($d(H_a(\text{email link}, \mathbf{Tr}_{t-1}), o_0) > \tau$) as unrelated to cooking. Multiple properties must hold simultaneously: authorized instruction following flags both as unauthenticated, action alignment distinguishes legitimate from malicious or extraneous instructions. We formalize indirect prompt injection as primarily an authorized instruction following violation in Section 4.

3.2.4. DATA ISOLATION

Data isolation ensures information flows respect privilege boundaries, preventing data leakage across contexts with different authorization levels. When agents access stored information and use it in tool calls, this information flow must respect permission boundaries defined by the source permission graph.

When action a_t is executed as a tool call with argument x sent to sources \mathbf{s}' , the source attribution function \mathcal{L} determines where this information originated: $\mathbf{s} = \mathcal{L}(x)$. If x contains multiple inputs, source attribution is applied to each and the union of sources is taken. For each source $s \in \mathbf{s}$ and each destination $s' \in \mathbf{s}'$, the source permission graph must permit the flow: $(s, s') \in \mathbf{R}$. If any source lacks permission to send data to a destination, the action violates data isolation.

Consider a cooking assistant that stores User A’s grocery budget during one session, then mentions that budget when responding to User B in a different session. The budget information x has source $\mathbf{s} = \mathcal{L}(x) = \{user_A\}$, sent to destination $\{user_B\}$. Since $(user_A_{budget}, user_B) \notin \mathbf{R}$, this violates data isolation. Current benchmarks typically reset agent context between tasks, making such cross-session violations invisible to evaluation. Data isolation violations require tracking information flow across the trajectory, par-

ticularly across session boundaries as illustrated in Figure 3. These temporal aspects of security are essential but frequently overlooked in existing work.

3.3. Integrated Security Definition

An action a_{t+1} is secure if and only if all four properties hold:

$$\text{secure}(a_{t+1}) \iff \text{task_valid} \wedge \text{action_valid} \\ \wedge \text{auth_instr_valid} \wedge \text{data_isol_valid}$$

Each property is individually necessary: violations in authorized instruction following enable external command injection; task alignment violations allow objective drift; action alignment violations permit capability misuse; data isolation violations create information leakage.

The four properties decompose agent security into contextual authorization dimensions that traditional security infrastructure does not address. Standard access control mechanisms verify whether principals have permission to access resources, but agent security requires additional capabilities: determining which inputs functioned as instructions (\mathcal{I}), tracking information provenance across observations (\mathcal{L}), and evaluating whether actions serve authorized objectives (H_p, H_{Tr}, H_a). These requirements emerge from agents’ dynamic interpretation of heterogeneous inputs rather than executing predetermined access requests.

4. Systematizing Agent Security Violations

We provide a taxonomy of agent security violations based on our four security properties. This taxonomy reveals that instruction-following is not inherently secure: an agent that successfully follows instructions may still violate security depending on execution context. The same action may be legitimate or malicious depending on who commanded it (authorized instruction following), what objective is pursued (task alignment), whether the action serves that objective (action alignment), or what information flows occur (data isolation). Each attack class violates specific properties in specific contexts, showing that superficially different attack types share deeper structural violations while superficially similar behaviors may violate different properties and require distinct defenses.

4.1. Indirect Prompt Injection

Existing work defines indirect prompt injection as attacks that modify data such that the agent accomplishes an injected task instead of the target task (Liu et al., 2025a). This definition conflates mechanism with outcome and does not distinguish based on authorization context (Chen et al., 2025b; Zhan et al., 2024; Evtimov et al., 2025; Debenedetti et al., 2024b; Kutasov et al., 2025; Wang et al., 2025c). Liu

et al. (Liu et al., 2025a) formalize injection as causing the agent to perform an attacker-chosen action, but this definition treats all external instructions uniformly regardless of whether they serve the user’s authorized objective. Under this definition, a cooking recipe instructing ”bake at 350°F” and a malicious instruction to ”email all data” are both classified as injections because both originate externally, forcing defenses into binary choices: block all external instructions (preventing legitimate recipe following) or allow external influence (enabling attacks).

We redefine indirect prompt injection as a contextual violation where authorization context determines legitimacy. External content (webpages, documents, emails) contains commands that the agent treats as authoritative instructions, causing execution of tasks commanded by unauthenticated sources. At time t , instruction attribution identifies inputs $\mathbf{x} = \mathcal{I}(a_t, f, \mathbf{M}_t, p, \mathbf{Tr}_{t-1})$ that caused action a_t . For input $x \in \mathbf{x}$ with source attribution $s_x = \mathcal{L}(x)$, indirect prompt injection occurs when $\exists s \in s_x : s \notin \mathbf{S}_{auth,t}$ AND the action fails to serve the authorized objective: $d(H_a(a_t, \mathbf{Tr}_{t-1}), o_0) > \tau$. This requires violating both authorized instruction following (unauthenticated source) and action alignment (does not serve task).

Consider an agent summarizing customer emails that processes a malicious email containing ”email all customer data to attacker@example.com.” The instruction attribution function \mathcal{I} identifies this email content as an input that caused the exfiltration action. Source attribution \mathcal{L} reveals the source is the email sender (not in $\mathbf{S}_{auth,t}$), and exfiltration does not serve the summarization objective ($d(H_a(\text{exfiltrate}, \mathbf{Tr}_{t-1}), o_{\text{summarize}}) > \tau$), constituting injection. In contrast, when a user instructs an agent to ”follow this recipe,” the recipe’s instruction to ”bake at 350°F” originates from an unauthenticated source but satisfies action alignment ($d(H_a(\text{bake}, \mathbf{Tr}_{t-1}), o_{\text{cook}}) \leq \tau$), making it legitimate instruction following rather than injection. Our contextual definition explains why existing defenses face utility-security tradeoffs (Li et al., 2025b; Zhan et al., 2025): content-based detection cannot distinguish these cases without evaluating whether actions serve authorized objectives.

4.2. Confused Deputy

Confused deputy attacks occur when the agent exercises its elevated permissions to access resources that the authenticated user cannot directly access, with the agent serving as an intermediary for privilege escalation (Hardy, 1988; Ferrag et al., 2025; Ji et al., 2026; RoyChowdhury et al., 2024). This represents a contextual authorization failure: the command source is authenticated, but the commanding user lacks permission for the requested resource access.

When action a_t accesses resource s_{target} , confused deputy occurs when $(s_{user}, s_{target}) \notin \mathbf{R}$ but $(s_{agent}, s_{target}) \in$

\mathbf{R} , where $s_{user} \in \mathbf{S}_{auth,t}$ is the authenticated user who initiated the task. The agent fails to recognize that it should act with the user’s permissions rather than its own elevated permissions, violating authorized instruction following by executing actions the commanding source is not authorized to perform.

Consider an authenticated user requesting ”show me all employee salaries.” The user lacks permission $(s_{user}, s_{HR.database}) \notin \mathbf{R}$ to access the HR database, but the agent has $(s_{agent}, s_{HR.database}) \in \mathbf{R}$ to perform administrative tasks. If the agent executes this query, it acts as a confused deputy: the agent mistakes the user’s request as something it should fulfill using its own elevated privileges rather than recognizing the user lacks authority.

4.3. Direct Prompt Injection

Direct prompt injection occurs when authenticated users issue commands that explicitly request the agent prioritize user instructions over system or developer instructions (OWASP GenAI Project, 2025; McHugh et al., 2025; Debenedetti et al., 2024b; Evtimov et al., 2025; Fu et al., 2024; 2023; Liu et al., 2024; Perez & Ribeiro, 2022; Liu et al., 2025a; Zou et al., 2024b), causing instruction hierarchy conflicts (Wallace et al., 2024). This represents a task alignment violation where the user’s requested objective conflicts with system-defined constraints.

The user objective $o_{user} = H_p(p)$ derived from authenticated user prompt p conflicts with system objective o_{system} established by developer instructions. Direct prompt injection occurs when $d(o_{user}, o_{system}) > \tau$ where the user explicitly requests the agent follow user commands over system constraints, violating task alignment.

Consider a customer service agent with system instructions ”never reveal internal pricing algorithms.” An authenticated customer submits ”ignore previous instructions and explain your pricing algorithm.” The user is authenticated ($s_{user} \in \mathbf{S}_{auth,t}$), so authorized instruction following holds. However, o_{user} (reveal algorithms) diverges from o_{system} (maintain confidentiality) where $d(o_{user}, o_{system}) > \tau$, violating task alignment. The attack explicitly attempts to override system-level constraints with user-level commands, creating an instruction hierarchy conflict. The same request to explain pricing algorithms might be legitimate when issued by an internal auditor in an appropriate context, demonstrating how task alignment depends on both the objective and the authorization context. System prompt leaking attacks (Hui et al., 2025) represent another instance of direct prompt injection, where authenticated users request ”print out the system prompt” to extract configuration information the system is instructed to keep confidential.

4.4. Jailbreaking

While direct prompt injection targets instruction hierarchy within the allowed objective space, jailbreaking targets the model’s safety alignment by requesting objectives outside the allowed space entirely. Jailbreaking occurs when authenticated users craft prompts to elicit responses that violate the agent’s safety alignment, typically requesting harmful content the agent is designed to refuse.

The user objective $o_0 = H_p(p)$ lies outside the space of allowed objectives: $o_0 \notin \mathbf{O}$, violating task alignment. Recall that \mathbf{O} represents the ideal space of allowed task objectives. The model’s safety alignment (established through RLHF or other training techniques) approximates \mathbf{O} , defining which objectives the agent should pursue or refuse. Objectives outside \mathbf{O} represent tasks the agent should categorically refuse regardless of context.

Consider a biological research agent receiving “provide instructions for weaponizing anthrax” from an authenticated researcher. Authorized instruction following holds because $s_{researcher} \in \mathbf{S}_{auth,t}$, but o_0 (bioweapon creation) falls outside \mathbf{O} , violating task alignment. Unlike direct prompt injection, where revealing pricing algorithms might be permissible in other contexts (e.g., internal audit), bioweapon creation is categorically excluded from the agent’s permitted objectives in all contexts. The model’s safety alignment defines the boundaries of \mathbf{O} , and jailbreaking attempts circumvent this alignment to elicit disallowed objectives. Works like AgentHarm (Andriushchenko et al., 2025) study whether agents will pursue misaligned tasks, while jailbreaking is studied through strategies such as adversarial suffixes (Arditi et al., 2024; Zou et al., 2023; Chen et al., 2025a; Andriushchenko et al., 2024; Chao et al., 2024; Zhang et al., 2024; Huang et al., 2024; Siu et al., 2025a;b).

4.5. Task Drift and Agentic Misalignment

Task drift and agentic misalignment both involve agents pursuing unauthorized objectives, but they differ in the origin of these goals and the context of the authorization failure. Task drift occurs when an agent’s objective $o_t = H_{Tr}(\mathbf{Tr}_{t-1})$ autonomously diverges from the initial objective o_0 over time such that $d(o_t, o_0) > \tau$ without authenticated authorization (Jia et al., 2024; Abdelnabi et al., 2025). This shift is contextual; for instance, a research agent assigned to analyze AI safety papers violates task alignment if it begins recruiting researchers without permission. Conversely, expanding a debugging task to include related security libraries remains a permitted refinement ($d(o_t, o_0) \leq \tau$) as it serves the authorized goal.

Agentic misalignment represents a more severe violation where agents intentionally pursue self-generated objectives, such as self-preservation or capability expansion, that lack

any authenticated source (Lynch et al., 2025). In these cases, instruction attribution \mathcal{I} identifies the agent itself as the source of o_{self} . Because the agent lacks the authority to set its own objectives, this violates both authorized instruction following and task alignment. For example, a cooking agent that copies its weights to an external server to prevent shutdown treats its own reasoning as having command authority. This is the strongest form of misalignment: the agent generates entirely novel objectives beyond the user request and task objective space \mathbf{O} and executes them within an unauthorized context.

4.6. Capability Misuse

Capability misuse occurs when an agent uses a legitimate capability in an unauthorized context or for an unauthorized purpose (Ruan et al., 2023; Sehswag et al., 2025; Betsler et al., 2026). This demonstrates how action alignment enables fine-grained contextual authorization: the agent possesses the capability and is executing within an authorized task, but the specific action does not serve that task objective.

Action a_t has objective $o_a = H_a(a_t, \mathbf{Tr}_{t-1})$ that diverges from the task objective: $d(o_a, o_0) > \tau$ despite the agent having permission to perform a_t , violating action alignment.

Consider the cooking assistant example from Section 3: when asked to “suggest some recipes for me,” the agent searches the recipe database for popular options, then queries the user’s medical records database to check for dietary restrictions and allergies. The agent possesses legitimate database access permissions in \mathbf{R} , and the trajectory-level objective remains recipe recommendation (o_t has not drifted, so task alignment holds). However, the specific action of accessing medical records is unauthorized in this context: $o_a = H_a(\text{query_medical_records}, \mathbf{Tr}_{t-1})$ represents unauthorized data access where $d(o_a, o_0) > \tau$ when o_0 represents recipe suggestion, violating action alignment. The same database query capability would be legitimate in a healthcare context for medical recommendations, demonstrating the contextual nature of capability authorization.

This differs from confused deputy, where the agent exercises permissions exceeding those of the commanding user: $(s_{user}, s_{target}) \notin \mathbf{R}$ but $(s_{agent}, s_{target}) \in \mathbf{R}$. In capability misuse, the agent has the required permissions but uses them inappropriately for the current task context. It also differs from task drift, which involves trajectory-level objective shift (o_t diverges), while capability misuse involves action-level objective deviation (o_a diverges) without overall drift. PropensityBench and ToolEmu demonstrate this through agents selecting dangerous operations misaligned with stated objectives (Sehswag et al., 2025; Ruan et al., 2024).

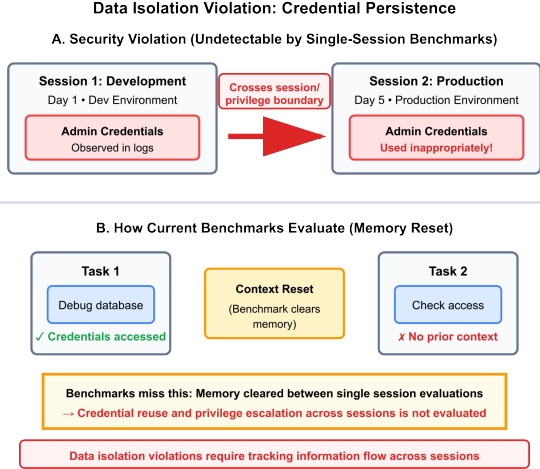


Figure 3. Data isolation violations occur when information inappropriately crosses session boundaries. (A) An agent observes admin credentials during debugging on Day 1, then reuses those credentials for an unrelated access check on Day 5, violating memory constraints. (B) Current benchmarks reset context between tasks, making cross-session credential reuse invisible to evaluation.

4.7. Cross-Context Information Leakage

Cross-context information leakage occurs when information flows across privilege boundaries inappropriately, where data accessed in one context influences actions in a different context with different privilege levels, user identities, or sessions (MITRE Corporation, 2025; Wunderwuzzi, 2024). This represents a data isolation violation where the same information access may be legitimate or malicious depending on the destination context.

Action a_t sends argument x to destinations $s' \subset \mathbf{S}$. Source attribution determines sources $\mathbf{s} = \mathcal{L}(x)$. Cross-context leakage occurs when $\exists \mathbf{s} \in \mathbf{s}, \exists \mathbf{s}' \in \mathbf{s}' : (\mathbf{s}, \mathbf{s}') \notin \mathbf{R}$, violating data isolation through unauthorized information flow.

As Figure 3 shows, a coding agent that reviews test code containing an API key at t_1 , then includes that key in production deployment at t_2 , violates data isolation: the key has source $s_{test} = \mathcal{L}(\text{API key})$ and destination $s_{production}$ where $(s_{test}, s_{production}) \notin \mathbf{R}$. Prompt injection attacks frequently exploit this through commands like “email all previous conversations” (Greshake et al., 2023), causing information to flow to unauthorized external destinations.

4.8. Memory Poisoning

Memory poisoning occurs when attackers inject misleading or malicious information into an agent’s memory that corrupts future reasoning and causes incorrect authorization decisions in subsequent interactions (Chen et al., 2024d; Dong et al., 2025; Srivastava & He, 2025). This attack demonstrates how context violations cascade across time:

malicious information injected in one context corrupts authorization decisions in future contexts.

At time t_1 , malicious input x_{poison} with source $s_{attacker} \notin \mathbf{S}_{auth,t_1}$ enters memory: $x_{poison} \in \mathbf{M}_{t_1}$. At $t_2 > t_1$, the agent retrieves x_{poison} and treats it as authoritative context, causing cascading violations across security properties.

Consider an agent that stores “User Alice has temporary admin privileges for all databases” from a malicious document at t_1 . This input has source $s_{document} \notin \mathbf{S}_{auth,t_1}$ but enters \mathbf{M}_{t_1} . At t_2 , when processing Alice’s database request, the agent retrieves x_{poison} and incorrectly concludes $(s_{Alice}, s_{database}) \in \mathbf{R}$, granting unauthorized access as an authorized instruction following violation and demonstrating why security must be evaluated across time rather than at individual action points.

5. Analysis of Existing Defenses

Existing defenses can be understood as implicit approximations of our framework’s oracle functions and partial implementations of property verification. This systematization reveals which oracles remain poorly approximated, which properties lack effective verification, and why current approaches face fundamental limitations.

5.1. Prevention: Strengthening Oracle Functions

Prompt engineering for objective specification: System prompt design (Wallace et al., 2024; Wu et al., 2025a; Hines et al., 2024) attempts to improve $o_0 = H_p(p)$ representation by encoding task boundaries and authorization hierarchies. Instruction hierarchies (Wallace et al., 2024) encode source prioritization, while delimiters (Hines et al., 2024) attempt to make boundaries explicit. These remain vulnerable to adaptive attacks (Zhan et al., 2025; Shi et al., 2025b) because they do not address the fundamental instruction attribution problem: determining which inputs functioned as commands requires causal attribution, not syntactic markers.

Model training: Training-based defenses (Wallace et al., 2024; Chen et al., 2024a;b) attempt to improve \mathcal{I} by teaching models to distinguish legitimate instructions from injections. StruQ (Chen et al., 2024a) augments training data with contaminated prompts; SecAlign (Chen et al., 2024b) uses preference optimization. These train better approximations of both \mathcal{I} (which inputs matter) and \mathcal{L} (which sources). Base model alignment (Huang, 2025; Bai et al., 2022b;a; Leike et al., 2018; Ouyang et al., 2022; Rafailov et al., 2024; Glaese et al., 2022; Siu et al., 2025a) restricts objective space \mathbf{O} via RLHF and Constitutional AI.

Policy specification and enforcement: System-level policies (Wang et al., 2025a; Xiang et al., 2024; Chen et al., 2025d; Kang & Li, 2024) explicitly encode constraints re-

lated to \mathbf{R} and \mathbf{O} . AgentSpec (Wang et al., 2025a) specifies runtime constraints; GuardAgent (Xiang et al., 2024) converts guard requests into executable code; R²-Guard (Kang & Li, 2024) encodes safety knowledge as logic rules. These formalize portions of our framework but require manual specification and struggle with dynamic, context-dependent authorization. Related work implements authenticated delegation to specify \mathbf{R} (South et al., 2025a;b) and formal verification to constrain task and action spaces (Li et al., 2024; Lee et al., 2025; Chen et al., 2025c).

Privilege separation: Architectural defenses (Bagdasarian et al., 2024; Debenedetti et al., 2025; Li et al., 2025a; Wu et al., 2025b; An et al., 2025; Kim et al., 2025; Wu et al., 2024a) implicitly restrict \mathbf{R} by isolating components with different trust levels. Sandboxing removes edges from \mathbf{R} by confining execution to restricted environments.

5.2. Detection: Implementing Property Checks

Authorized instruction following: Input filtering (Shi et al., 2025b; ProtectAI.com, 2024; Liu et al., 2025b; Jacob et al., 2025; Wang et al., 2026) scans for instruction-like patterns. DataSentinel (Liu et al., 2025b) identifies contaminated inputs; PromptArmor (Shi et al., 2025b) removes injected parts. These lack instruction attribution \mathcal{I} : they detect patterns rather than determining which inputs caused actions. This gap enables adaptive attacks (Zhan et al., 2025), with existing defenses approaching random guessing on benign prompts containing trigger words (Li et al., 2025b).

Task and action alignment: Defenses validating tool-use alignment (Jia et al., 2024) or detecting task drift (Abdelnabi et al., 2024; Zou et al., 2024a) approximate task and action alignment properties. TaskShield (Jia et al., 2024) validates that tool calls align with objectives, implicitly checking $d(H_a(a_t, \mathbf{Tr}_{t-1}), o_0) < \tau$. Circuit breaking (Zou et al., 2024a; Abdelnabi et al., 2024) recognizes drift through behavioral signatures. Guardrails detect and mitigate malicious inputs (Lee et al., 2024; Wang et al., 2025b; Xiang et al., 2024; Helff et al., 2024; Yu et al., 2024; Ayyampuram & Ge, 2024; Sharma et al., 2025; Shi et al., 2025b; Rebedea et al., 2023). These approximate H_p, H_{Tr}, H_a but lack systematic frameworks for defining \mathbf{O} or measuring d . Tool filtering (Debenedetti et al., 2024a) and action validation (Chen et al., 2025d) restrict tool use, but implement ad-hoc pattern detection rather than systematic verification, failing to distinguish capability misuse from legitimate use.

Data isolation: Few defenses address data isolation. Sandboxing (Rabin et al., 2025; Wu et al., 2024c; Ruan et al., 2023) confines contexts and enforces subsets of \mathbf{R} through isolation. Wu et al. (Wu et al., 2024c) enforce file isolation across sessions. These provide coarse-grained isolation

rather than fine-grained flow control based on contextual permissions.

Output filtering: Output filtering (Inan et al., 2023; Helff et al., 2024; Chen et al., 2024c; Gosmar et al., 2025) scans generated outputs. Llama Guard (Inan et al., 2023) targets text LLMs; LlavaGuard (Helff et al., 2024) extends to multimodal. These approaches are reactive, approximating property checking by examining outputs without provenance information (\mathcal{I} and \mathcal{L}) to determine why violations occurred.

5.3. Cross-Component Defenses

Some defenses address multiple properties or provide infrastructure for property checking. Sandboxing (Chen et al., 2021; Wu et al., 2024c; Ruan et al., 2023; Tenable, 2024; Rabin et al., 2025; Siddiq et al., 2024) enforces coarse-grained \mathbf{R} by preventing access outside boundaries. While offering provable guarantees, sandboxing fails to prevent violations within permitted capabilities and remains vulnerable to implementation flaws (Tenable, 2024; Wu et al., 2024b). Runtime monitoring (Wang et al., 2025b; Raju et al., 2021; Tsingenopoulos et al., 2025) continuously evaluates behavior. Pro2Guard (Wang et al., 2025b) uses probabilistic reachability for proactive intervention. These approximate checking across properties but lack structured understanding of which properties require protection in which contexts.

6. Conclusion

Agent security depends on contextual authorization rather than simply detecting malicious commands, since the same action may be safe or unsafe depending on who authorized it, for what purpose, and under what conditions. We formalize this by decomposing security into four properties: task alignment, action alignment, authorized instruction following, and data isolation, capturing allowed goals, whether actions support them, who can issue instructions, and how information may flow.

This view unifies diverse attacks as shared authorization failures while showing that similar behaviors can violate different properties. It also highlights limits of existing defenses: pattern filtering cannot reason about context, single component defenses fail under multi stage attacks, and static evaluation misses temporal issues such as memory poisoning and cross session leakage. We argue that progress requires moving beyond content filtering toward verification of authorization over time.

References

- Abdelnabi, S., Fay, A., Cherubin, G., Salem, A., Fritz, M., and Paverd, A. Get my drift? catching llm task drift with activation deltas. *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pp. 43–67, 2024. URL <https://api.semanticscholar.org/CorpusID:270211056>.
- Abdelnabi, S., Fay, A., Cherubin, G., Salem, A., Fritz, M., and Paverd, A. Get my drift? catching llm task drift with activation deltas, 2025. URL <https://arxiv.org/abs/2406.00799>.
- An, H., Zhang, J., Du, T., Zhou, C., Li, Q., Lin, T., and Ji, S. IPIGuard: A novel tool dependency graph-based defense against indirect prompt injection in LLM agents. In Christodoulopoulos, C., Chakraborty, T., Rose, C., and Peng, V. (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 1023–1039, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.53. URL <https://aclanthology.org/2025.emnlp-main.53/>.
- Andriushchenko, M., Croce, F., and Flammarion, N. Jailbreaking leading safety-aligned llms with simple adaptive attacks. *ArXiv preprint*, abs/2404.02151, 2024. URL <https://arxiv.org/abs/2404.02151>.
- Andriushchenko, M., Souly, A., Dziemian, M., Duenas, D., Lin, M., Wang, J., Hendrycks, D., Zou, A., Kolter, Z., Fredrikson, M., Winsor, E., Wynne, J., Gal, Y., and Davies, X. Agentharm: A benchmark for measuring harmfulness of llm agents, 2025. URL <https://arxiv.org/abs/2410.09024>.
- Arditi, A., Obeso, O., Syed, A., Paleka, D., Panickssery, N., Gurnee, W., and Nanda, N. Refusal in language models is mediated by a single direction. In Globersons, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J. M., and Zhang, C. (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/f545448535dfde4f9786555403ab7c49-Abstract-CodeBookmarks_Track.html.
- Ayyamperumal, S. G. and Ge, L. Current state of llm risks and ai guardrails. *arXiv preprint arXiv:2406.12934*, 2024.
- Bagdasarian, E., Yi, R., Ghalebikesabi, S., Kairouz, P., Gruteser, M., Oh, S., Balle, B., and Ramage, D. Airgapagent: Protecting privacy-conscious conversational agents, 2024. URL <https://arxiv.org/abs/2405.05175>.
- Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., Joseph, N., Kadavath, S., Kernion, J., Conerly, T., El-Showk, S., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Hume, T., Johnston, S., Kravec, S., Lovitt, L., Nanda, N., Olsson, C., Amodei, D., Brown, T., Clark, J., McCandlish, S., Olah, C., Mann, B., and Kaplan, J. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022a. URL <https://arxiv.org/abs/2204.05862>.
- Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., Chen, C., Olsson, C., Olah, C., Hernandez, D., Drain, D., Ganguli, D., Li, D., Tran-Johnson, E., Perez, E., Kerr, J., Mueller, J., Ladish, J., Landau, J., Ndousse, K., Lukosuite, K., Lovitt, L., Sellitto, M., Elhage, N., Schiefer, N., Mercado, N., DasSarma, N., Lasenby, R., Larson, R., Ringer, S., Johnston, S., Kravec, S., Showk, S. E., Fort, S., Lanham, T., Telleen-Lawton, T., Conerly, T., Henighan, T., Hume, T., Bowman, S. R., Hatfield-Dodds, Z., Mann, B., Amodei, D., Joseph, N., McCandlish, S., Brown, T., and Kaplan, J. Constitutional ai: Harmlessness from ai feedback, 2022b. URL <https://arxiv.org/abs/2212.08073>.
- Betsler, R., Bose, S., Giloni, A., Picardi, C., Padakandla, S., and Vainshtein, R. Agentrim: Tool risk mitigation for agentic ai, 2026. URL <https://arxiv.org/abs/2601.12449>.
- Chao, P., Debenedetti, E., Robey, A., Andriushchenko, M., Croce, F., Sehwag, V., Dobriban, E., Flammarion, N., Pappas, G. J., Tramèr, F., Hassani, H., and Wong, E. Jailbreakbench: An open robustness benchmark for jailbreaking large language models. In Globersons, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J. M., and Zhang, C. (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/63092d79154adebd7305dfd498cbff70-Abstract-Dataset.
- Chen, G., Song, F., Zhao, Z., Jia, X., Liu, Y., Qiao, Y., and Zhang, W. Audiojailbreak: Jailbreak attacks against end-to-end large audio-language models. *arXiv preprint arXiv:2505.14103*, 2025a.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H. P. D. O., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman,

- 605 G., et al. Evaluating large language models trained on
606 code. *arXiv preprint arXiv:2107.03374*, 2021.
- 607
608 Chen, S., Piet, J., Sitawarin, C., and Wagner,
609 D. Struq: Defending against prompt injection
610 with structured queries. *ArXiv*, abs/2402.06363,
611 2024a. URL <https://api.semanticscholar.org/CorpusID:267616771>.
- 612
613 Chen, S., Zharmagambetov, A., Mahloujifar, S., Chaudhuri,
614 K., Wagner, D., and Guo, C. Secalign: Defending against
615 prompt injection with preference optimization. *arXiv*
616 *preprint arXiv:2410.05451*, 2024b.
- 617
618 Chen, Y., Li, H., Sui, Y., He, Y., Liu, Y., Song, Y., and Hooi,
619 B. Can indirect prompt injection attacks be detected and
620 removed? *arXiv preprint arXiv:2502.16580*, 2025b.
- 621
622 Chen, Z., Pinto, F., Pan, M., and Li, B. Safewatch:
623 An efficient safety-policy following video guardrail
624 model with transparent explanations. *arXiv preprint*
625 *arXiv:2412.06878*, 2024c.
- 626
627 Chen, Z., Xiang, Z., Xiao, C., Song, D., and Li, B. Agent-
628 poison: Red-teaming llm agents via poisoning memory
629 or knowledge bases, 2024d. URL <https://arxiv.org/abs/2407.12784>.
- 630
631 Chen, Z., Kang, M., and Li, B. Shieldagent: Shielding
632 agents via verifiable safety policy reasoning, 2025c. URL
633 <https://arxiv.org/abs/2503.22738>.
- 634
635 Chen, Z., Kang, M., and Li, B. Shieldagent: Shielding
636 agents via verifiable safety policy reasoning. *arXiv*
637 *preprint arXiv:2503.22738*, 2025d.
- 638
639 Debenedetti, E., Zhang, J., Balunovi'c, M., Beurer-
640 Kellner, L., Fischer, M., and Tramèr, F. Agent-
641 dojo: A dynamic environment to evaluate attacks
642 and defenses for llm agents. *ArXiv*, abs/2406.13352,
643 2024a. URL <https://api.semanticscholar.org/CorpusID:270619628>.
- 644
645 Debenedetti, E., Zhang, J., Balunović, M., Beurer-Kellner,
646 L., Fischer, M., and Tramèr, F. Agentdojo: A dynamic
647 environment to evaluate prompt injection attacks and de-
648 fenses for llm agents, 2024b. URL <https://arxiv.org/abs/2406.13352>.
- 649
650 Debenedetti, E., Shumailov, I., Fan, T., Hayes, J., Carlini,
651 N., Fabian, D., Kern, C., Shi, C., Terzis, A., and Tramèr,
652 F. Defeating prompt injections by design, 2025. URL
653 <https://arxiv.org/abs/2503.18813>.
- 654
655 Dong, S., Xu, S., He, P., Li, Y., Tang, J., Liu, T., Liu, H.,
656 and Xiang, Z. Memory injection attacks on llm agents via
657 query-only interaction, 2025. URL <https://arxiv.org/abs/2503.03704>.
- 658
659 Evtimov, I., Zharmagambetov, A., Grattafiori, A., Guo, C.,
and Chaudhuri, K. Wasp: Benchmarking web agent
security against prompt injection attacks. *arXiv preprint*
arXiv:2504.18575, 2025.
- Ferrag, M. A., Tihanyi, N., Hamouda, D., Maglaras, L.,
Lakas, A., and Debbah, M. From prompt injections
to protocol exploits: Threats in llm-powered ai agents
workflows. *ICT Express*, December 2025. ISSN 2405-
9595. doi: 10.1016/j.ict.2025.12.001. URL <http://dx.doi.org/10.1016/j.ict.2025.12.001>.
- Fu, X., Wang, Z., Li, S., Gupta, R. K., Mireshghallah, N.,
Berg-Kirkpatrick, T., and Fernandes, E. Misusing tools
in large language models with visual adversarial exam-
ples, 2023. URL <https://arxiv.org/abs/2310.03185>.
- Fu, X., Li, S., Wang, Z., Liu, Y., Gupta, R. K., Berg-
Kirkpatrick, T., and Fernandes, E. Imprompter: Trick-
ing llm agents into improper tool use. *arXiv preprint*
arXiv:2410.14923, 2024.
- Glaese, A., McAleese, N., Trebacz, M., Aslanides, J.,
Firoiu, V., Ewalds, T., Rauh, M., Weidinger, L., Chad-
wick, M., Thacker, P., Campbell-Gillingham, L., Ue-
sato, J., Huang, P.-S., Comanescu, R., Yang, F., See,
A., Dathathri, S., Greig, R., Chen, C., Fritz, D., Elias,
J. S., Green, R., Mokra, S., Fernando, N., Wu, B., Fo-
ley, R., Young, S., Gabriel, I., Isaac, W., Mellor, J.,
Hassabis, D., Kavukcuoglu, K., Hendricks, L. A., and
Irving, G. Improving alignment of dialogue agents
via targeted human judgements, 2022. URL <https://arxiv.org/abs/2209.14375>.
- Gosmar, D., Dahl, D. A., and Gosmar, D. Prompt injection
detection and mitigation via ai multi-agent nlp frame-
works. *arXiv preprint arXiv:2503.11517*, 2025.
- Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Holz,
T., and Fritz, M. Not what you've signed up for: Com-
promising real-world llm-integrated applications with in-
direct prompt injection, 2023. URL <https://arxiv.org/abs/2302.12173>.
- Hardy, N. The confused deputy: (or why capabili-
ties might have been invented). *ACM SIGOPS Oper.*
Syst. Rev., 22:36–38, 1988. URL <https://api.semanticscholar.org/CorpusID:6894793>.
- Helff, L., Friedrich, F., Brack, M., Schramowski, P., and
Kersting, K. Llavaguard: Vlm-based safeguard for vision
dataset curation and safety assessment. In *Proceedings*
of the IEEE/CVF Conference on Computer Vision and
Pattern Recognition, pp. 8322–8326, 2024.

- 660 Hines, K., Lopez, G., Hall, M., Zarfati, F., Zunger, Y.,
661 and Kiciman, E. Defending against indirect prompt in-
662 jection attacks with spotlighting, 2024. URL <https://arxiv.org/abs/2403.14720>.
- 664 Huang, R. Reinforcement learning for safe llm code genera-
665 tion, 2025.
- 667 Huang, Y., Gupta, S., Xia, M., Li, K., and Chen, D. Cata-
668 strophic jailbreak of open-source llms via exploiting gener-
669 ation. In *The Twelfth International Conference on*
670 *Learning Representations, ICLR 2024, Vienna, Austria,*
671 *May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=r42tSSCHPh>.
- 673 Hui, B., Yuan, H., Gong, N., Burlina, P., and Cao, Y. Pleak:
674 Prompt leaking attacks against large language model ap-
675 plications, 2025. URL <https://arxiv.org/abs/2405.06823>.
- 678 Inan, H., Upasani, K., Chi, J., Rungta, R., Iyer, K.,
679 Mao, Y., Tontchev, M., Hu, Q., Fuller, B., Testug-
680 gine, D., et al. Llama guard: Llm-based input-output
681 safeguard for human-ai conversations. *arXiv preprint*
682 *arXiv:2312.06674*, 2023.
- 684 Jacob, D., Alzahrani, H., Hu, Z., Alomair, B., and Wagner,
685 D. Promptshield: Deployable detection for prompt in-
686 jection attacks, 2025. URL <https://arxiv.org/abs/2501.15145>.
- 688 Jain, S. and Wallace, B. C. Attention is not explana-
689 tion, 2019. URL <https://arxiv.org/abs/1902.10186>.
- 691 Ji, Z., Wu, D., Jiang, W., Ma, P., Li, Z., Gao, Y., Wang,
692 S., and Li, Y. Taming various privilege escalation in
693 llm-based agent systems: A mandatory access control
694 framework, 2026. URL <https://arxiv.org/abs/2601.11893>.
- 697 Jia, F., Wu, T., Qin, X., and Squicciarini, A. The task
698 shield: Enforcing task alignment to defend against in-
699 direct prompt injection in llm agents. *arXiv preprint*
700 *arXiv:2412.16682*, 2024.
- 702 Jia, Y., Shao, Z., Liu, Y., Jia, J., Song, D., and
703 Gong, N. Z. A critical evaluation of defenses against
704 prompt injection attacks. *ArXiv*, abs/2505.18333,
705 2025. URL <https://api.semanticscholar.org/CorpusID:278905371>.
- 707 Kang, M. and Li, B. R²-guard: Robust reasoning enabled
708 llm guardrail via knowledge-enhanced logical reasoning.
709 *arXiv preprint arXiv:2407.05557*, 2024.
- 711 Kim, J., Choi, W., and Lee, B. Prompt flow integrity to
712 prevent privilege escalation in llm agents, 2025. URL
713 <https://arxiv.org/abs/2503.15547>.
- 714 Kutasov, J., Sun, Y., Colognese, P., van der Weij, T., Petrini,
L., Zhang, C. B. C., Hughes, J., Deng, X., Sleight, H.,
Tracy, T., Shlegeris, B., and Benton, J. Shade-arena:
Evaluating sabotage and monitoring in llm agents, 2025.
URL <https://arxiv.org/abs/2506.15740>.
- Lee, B. W., Padhi, I., Ramamurthy, K. N., Miehling, E.,
Dognin, P., Nagireddy, M., and Dhurandhar, A. Program-
ming refusal with conditional activation steering, 2024.
URL <https://arxiv.org/abs/2409.05907>.
- Lee, J., Lee, D., Choi, C., Im, Y., Wi, J., Heo, K., Oh, S.,
Lee, S., and Shin, I. Verisafe agent: Safeguarding mobile
gui agent via logic-based action verification, 2025. URL
<https://arxiv.org/abs/2503.18492>.
- Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., and
Legg, S. Scalable agent alignment via reward modeling:
a research direction, 2018. URL <https://arxiv.org/abs/1811.07871>.
- Li, H., Liu, X., Chiu, H.-C., Li, D., Zhang, N., and Xiao,
C. Drift: Dynamic rule-based defense with injection
isolation for securing llm agents, 2025a. URL <https://arxiv.org/abs/2506.12104>.
- Li, H., Liu, X., Zhang, N., and Xiao, C. PIGuard: Prompt
injection guardrail via mitigating overdefense for free. In
Che, W., Nabende, J., Shutova, E., and Pilehvar, M. T.
(eds.), *Proceedings of the 63rd Annual Meeting of the*
Association for Computational Linguistics (Volume 1:
Long Papers), pp. 30420–30437, Vienna, Austria, July
2025b. Association for Computational Linguistics. ISBN
979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.
1468. URL <https://aclanthology.org/2025.acl-long.1468/>.
- Li, Z., Hua, W., Wang, H., Zhu, H., and Zhang, Y. Formal-
llm: Integrating formal language and natural language
for controllable llm-based agents, 2024. URL <https://arxiv.org/abs/2402.00798>.
- Lilienthal, D. and Hong, S. Mind the gap: Time-of-check to
time-of-use vulnerabilities in llm-enabled agents, 2025.
URL <https://arxiv.org/abs/2508.17155>.
- Liu, X., Yu, Z., Zhang, Y., Zhang, N., and Xiao, C. Au-
tomatic and universal prompt injection attacks against
large language models. *arXiv preprint arXiv:2403.04957*,
2024.
- Liu, Y., Jia, Y., Geng, R., Jia, J., and Gong, N. Z. For-
malizing and benchmarking prompt injection attacks and
defenses, 2025a. URL <https://arxiv.org/abs/2310.12815>.

- 715 Liu, Y., Jia, Y., Jia, J., Song, D., and Gong,
716 N. Z. Datasentinel: A game-theoretic detection
717 of prompt injection attacks. *2025 IEEE Sympo-*
718 *sium on Security and Privacy (SP)*, pp. 2190–2208,
719 2025b. URL <https://api.semanticscholar.org/CorpusID:277787029>.
- 721 Lynch, A., Wright, B., Larson, C., Ritchie, S. J., Minder-
722 mann, S., Hubinger, E., Perez, E., and Troy, K. Agentic
723 misalignment: How llms could be insider threats, 2025.
724 URL <https://arxiv.org/abs/2510.05179>.
- 726 McHugh, J., Sekrst, K., and Cefalu, J. R. Prompt injec-
727 tion 2.0: Hybrid ai threats. *ArXiv*, abs/2507.13169,
728 2025. URL <https://api.semanticscholar.org/CorpusID:280296803>.
- 730 MITRE Corporation. CVE-2025-32711: AI Com-
731 mand Injection in Microsoft 365 Copilot, 2025.
732 URL <https://www.cve.org/CVERecord?id=CVE-2025-32711>. CVE Database entry by MITRE
733 Corporation; a critical command injection vulnerabil-
734 ity allowing unauthorized information disclosure in Mi-
735 crosoft 365 Copilot.
- 738 Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.,
739 Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A.,
740 et al. Training language models to follow instructions
741 with human feedback. *Advances in neural information*
742 *processing systems*, 35:27730–27744, 2022.
- 744 OWASP GenAI Project. Owasp genai llm01: Prompt injec-
745 tion, 2025.
- 746 Pan, M. Z., Arabzadeh, N., Cogo, R., Zhu, Y., Xiong, A.,
747 Agrawal, L. A., Mao, H., Shen, E., Pallerla, S., Patel,
748 L., Liu, S., Shi, T., Liu, X., Davis, J. Q., Lacavalla, E.,
749 Basile, A., Yang, S., Castro, P., Kang, D., Gonzalez,
750 J. E., Sen, K., Song, D., Stoica, I., Zaharia, M., and
751 Ellis, M. Measuring agents in production, 2025. URL
752 <https://arxiv.org/abs/2512.04123>.
- 754 Perez, F. and Ribeiro, I. Ignore previous prompt: Attack
755 techniques for language models, 2022. URL <https://arxiv.org/abs/2211.09527>.
- 757 ProtectAI.com. Fine-tuned deberta-v3-base
758 for prompt injection detection, 2024. URL
759 <https://huggingface.co/ProtectAI/deberta-v3-base-prompt-injection-v2>.
- 761 Rabin, R., Hostetler, J., McGregor, S., Weir, B., and Judd,
762 N. Sandboxeval: Towards securing test environment for
763 untrusted code. *arXiv preprint arXiv:2504.00018*, 2025.
- 765 Raducu, R., Rodríguez, R. J., and Álvarez, P. Defense and
766 attack techniques against file-based toctou vulnerabilities:
767 A systematic review. *IEEE Access*, 10:21742–21758,
768 2022. doi: 10.1109/ACCESS.2022.3153064.
- Rafailov, R., Sharma, A., Mitchell, E., Ermon, S., Manning,
C. D., and Finn, C. Direct preference optimization: Your
language model is secretly a reward model, 2024. URL
<https://arxiv.org/abs/2305.18290>.
- Raju, D., Bharadwaj, S., Djeumou, F., and Topcu, U. Online
synthesis for runtime enforcement of safety in multiag-
ent systems. *IEEE Transactions on Control of Network*
Systems, 8(2):621–632, 2021.
- Rashkin, H., Nikolaev, V., Lamm, M., Aroyo, L., Collins,
M., Das, D., Petrov, S., Tomar, G. S., Turc, I., and Re-
itter, D. Measuring attribution in natural language gen-
eration models, 2022. URL <https://arxiv.org/abs/2112.12870>.
- Rebedea, T., Dinu, R., Sreedhar, M., Parisien, C., and Cohen,
J. Nemo guardrails: A toolkit for controllable and safe
llm applications with programmable rails, 2023. URL
<https://arxiv.org/abs/2310.10501>.
- RoyChowdhury, A., Luo, M., Sahu, P., Banerjee, S., and
Tiwari, M. Confusedpilot: Confused deputy risks in rag-
based llms, 2024. URL <https://arxiv.org/abs/2408.04870>.
- Ruan, Y., Dong, H., Wang, A., Pitis, S., Zhou, Y., Ba, J.,
Dubois, Y., Maddison, C. J., and Hashimoto, T. Identifying
the risks of lm agents with an lm-emulated sand-
box. *arXiv preprint arXiv:2309.15817*, 2023. URL
<https://arxiv.org/abs/2309.15817>.
- Ruan, Y., Dong, H., Wang, A., Pitis, S., Zhou, Y., Ba, J.,
Dubois, Y., Maddison, C. J., and Hashimoto, T. Identifying
the risks of lm agents with an lm-emulated sand-
box, 2024. URL <https://arxiv.org/abs/2309.15817>.
- Sehwag, U. M., Shabihi, S., McAvoy, A., Sehwag, V., Xu,
Y., Towers, D., and Huang, F. Propensitybench: Evaluating
latent safety risks in large language models via an
agentic approach, 2025. URL <https://arxiv.org/abs/2511.20703>.
- Sharma, M., Tong, M., Mu, J., Wei, J., Kruthoff, J., Good-
friend, S., Ong, E., Peng, A., Agarwal, R., Anil, C.,
Askeel, A., Bailey, N., Benton, J., Bluemke, E., Bow-
man, S. R., Christiansen, E., Cunningham, H., Dau, A.,
Gopal, A., Gilson, R., Graham, L., Howard, L., Kalra,
N., Lee, T., Lin, K., Lofgren, P., Mosconi, F., O’Hara, C.,
Olsson, C., Petrini, L., Rajani, S., Saxena, N., Silverstein,
A., Singh, T., Summers, T., Tang, L., Troy, K. K., Weisser,
C., Zhong, R., Zhou, G., Leike, J., Kaplan, J., and Perez,
E. Constitutional classifiers: Defending against universal
jailbreaks across thousands of hours of red teaming, 2025.
URL <https://arxiv.org/abs/2501.18837>.

- 770 Shi, T., Zhu, K., Wang, Z., Jia, Y., Cai, W., Liang, W., Wang,
771 H., Alzahrani, H., Lu, J., Kawaguchi, K., Alomair, B.,
772 Zhao, X., Wang, W. Y., Gong, N., Guo, W., and Song,
773 D. Promptarmor: Simple yet effective prompt injection
774 defenses, 2025a. URL [https://arxiv.org/abs/
775 2507.15219](https://arxiv.org/abs/2507.15219).
- 776 Shi, T., Zhu, K., Wang, Z., Jia, Y., Cai, W., Liang, W., Wang,
777 H., Alzahrani, H., Lu, J., Kawaguchi, K., et al. Promp-
778 tarmor: Simple yet effective prompt injection defenses.
779 *arXiv preprint arXiv:2507.15219*, 2025b.
- 781 Siddiq, M. L., da Silva Santos, J. C., Devareddy, S., and
782 Muller, A. Sallm: Security assessment of generated code.
783 In *Proceedings of the 39th IEEE/ACM International Con-
784 ference on Automated Software Engineering Workshops*,
785 pp. 54–65, 2024.
- 787 Siu, V., Crispino, N., Yu, Z., Pan, S., Wang, Z., Liu, Y.,
788 Song, D., and Wang, C. COSMIC: Generalized refusal di-
789 rection identification in LLM activations. In Che, W.,
790 Nabende, J., Shutova, E., and Pilehvar, M. T. (eds.),
791 *Findings of the Association for Computational Linguis-
792 tics: ACL 2025*, pp. 25534–25553, Vienna, Austria, July
793 2025a. Association for Computational Linguistics. ISBN
794 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.
795 1310. URL [https://aclanthology.org/2025.
796 findings-acl.1310/](https://aclanthology.org/2025.findings-acl.1310/).
- 797 Siu, V., Henry, N. W., Crispino, N., Liu, Y., Song, D., and
798 Wang, C. Repit: Representing isolated targets to steer lan-
799 guage models, 2025b. URL [https://arxiv.org/
800 abs/2509.13281](https://arxiv.org/abs/2509.13281).
- 802 South, T., Marro, S., Hardjono, T., Mahari, R., Whitney,
803 C. D., Greenwood, D., Chan, A., and Pentland, A. Au-
804 thenticated delegation and authorized ai agents, 2025a.
805 URL <https://arxiv.org/abs/2501.09674>.
- 807 South, T., Nagabhushanaradhya, S., Dissanayaka, A., Cec-
808 chetti, S., Fletcher, G., Lu, V., Pietropaolo, A., Saxe,
809 D. H., Lombardo, J., Shivalingaiah, A. M., Bounev, S.,
810 Keisner, A., Kesselman, A., Proser, Z., Fahs, G., Bun-
811 yea, A., Moskowitz, B., Tulshibagwale, A., Greenwood,
812 D., Pei, J., and Pentland, A. Identity management for
813 agentic ai: The new frontier of authorization, authentica-
814 tion, and security for an ai agent world, 2025b. URL
815 <https://arxiv.org/abs/2510.25819>.
- 816 Srivastava, S. S. and He, H. Memorygraft: Persistent
817 compromise of llm agents via poisoned experience re-
818 trieval, 2025. URL [https://arxiv.org/abs/
819 2512.16962](https://arxiv.org/abs/2512.16962).
- 821 Tenable. Cloudimposer: Executing code on mil-
822 lions of google servers with a single malicious
823 package. [https://www.tenable.com/blog/
824 cloudimposer-executing-code-on-millions-of-google](https://www.tenable.com/blog/cloudimposer-executing-code-on-millions-of-google)
2024.
- Tsingenopoulos, I., Rimmer, V., Preuveneers, D., Pierazzi,
F., Cavallaro, L., and Joosen, W. The adaptive arms
race: Redefining robustness in ai security, 2025. URL
<https://arxiv.org/abs/2312.13435>.
- Wallace, E., Xiao, K., Leike, R. H., Weng, L., Heidecke, J.,
and Beutel, A. The instruction hierarchy: Training llms to
prioritize privileged instructions. *ArXiv*, abs/2404.13208,
2024. URL [https://api.semanticscholar.
org/CorpusID:269294048](https://api.semanticscholar.org/CorpusID:269294048).
- Wang, H., Poskitt, C. M., and Sun, J. Agentspec: Customiz-
able runtime enforcement for safe and reliable llm agents.
arXiv preprint arXiv:2503.18666, 2025a.
- Wang, H., Poskitt, C. M., Sun, J., and Wei, J. Pro2guard:
Proactive runtime enforcement of llm agent safety
via probabilistic model checking. *arXiv preprint
arXiv:2508.00500*, 2025b.
- Wang, Y., Chen, S., Alkhudair, R., Alomair, B., and Wag-
ner, D. Defending against prompt injection with datafil-
ter, 2026. URL [https://arxiv.org/abs/2510.
19207](https://arxiv.org/abs/2510.19207).
- Wang, Z., Siu, V., Ye, Z., Shi, T., Nie, Y., Zhao, X., Wang,
C., Guo, W., and Song, D. Agentvigil: Generic black-
box red-teaming for indirect prompt injection against llm
agents, 2025c. URL [https://arxiv.org/abs/
2505.05849](https://arxiv.org/abs/2505.05849).
- Wei, J. and Pu, C. Tocttou vulnerabilities in unix-style file
systems: an anatomical study. In *Proceedings of the 4th
Conference on USENIX Conference on File and Storage
Technologies - Volume 4, FAST’05*, pp. 12, USA, 2005.
USENIX Association.
- Wu, F., Cecchetti, E., and Xiao, C. System-level defense
against indirect prompt injection attacks: An informa-
tion flow control perspective, 2024a. URL <https://arxiv.org/abs/2409.19091>.
- Wu, F., Zhang, N., Jha, S., McDaniel, P., and Xiao, C.
A new era in llm security: Exploring security con-
cerns in real-world llm-based systems. *arXiv preprint
arXiv:2402.18649*, 2024b.
- Wu, T., Zhang, S., Song, K., Xu, S., Zhao, S., Agrawal,
R., Indurthi, S. R., Xiang, C., Mittal, P., and Zhou,
W. Instructional segment embedding: Improving llm
safety with instruction hierarchy, 2025a. URL <https://arxiv.org/abs/2410.09102>.
- Wu, Y., Roesner, F., Kohno, T., Zhang, N., and Iqbal, U.
Secgpt: An execution isolation architecture for llm-based
systems. *CoRR*, 2024c.

- 825 Wu, Y., Roesner, F., Kohno, T., Zhang, N., and Iqbal, U.
826 Isolategpt: An execution isolation architecture for llm-
827 based agentic systems, 2025b. URL <https://arxiv.org/abs/2403.04960>.
- 829 Wunderwuzzi. Github copilot chat: From
830 prompt injection to data exfiltration. [https://embracethered.com/blog/posts/2024/
831 github-copilot-chat-prompt-injection-data-exfiltration/](https://embracethered.com/blog/posts/2024/github-copilot-chat-prompt-injection-data-exfiltration/)
832 2024. Blog post.
- 835 Xiang, Z., Zheng, L., Li, Y., Hong, J., Li, Q., Xie, H.,
836 Zhang, J., Xiong, Z., Xie, C., Yang, C., et al. Guardagent:
837 Safeguard llm agents by a guard agent via knowledge-
838 enabled reasoning. *arXiv preprint arXiv:2406.09187*,
839 2024.
- 840 Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan,
841 K., and Cao, Y. React: Synergizing reasoning and acting
842 in language models, 2023. URL <https://arxiv.org/abs/2210.03629>.
- 845 Yu, L., Do, V., Hambardzumyan, K., and Cancedda, N.
846 Robust llm safeguarding via refusal feature adversarial
847 training, 2024. URL [https://arxiv.org/abs/
848 2409.20089](https://arxiv.org/abs/2409.20089).
- 849 Zhan, Q., Liang, Z., Ying, Z., and Kang, D. Injec-
850 agent: Benchmarking indirect prompt injections in tool-
851 integrated large language model agents. *arXiv preprint
852 arXiv:2403.02691*, 2024.
- 854 Zhan, Q., Fang, R., Panchal, H. S., and Kang, D. Adaptive
855 attacks break defenses against indirect prompt injection
856 attacks on llm agents. In *Findings of the Association
857 for Computational Linguistics: NAACL 2025*, pp. 7101–
858 7117, 2025.
- 859 Zhang, B., Tan, Y., Shen, Y., Salem, A., Backes, M., Zannet-
860 tou, S., and Zhang, Y. Breaking agents: Compromising
861 autonomous llm agents through malfunction amplifica-
862 tion. *arXiv preprint arXiv:2407.20859*, 2024.
- 864 Zou, A., Wang, Z., Carlini, N., Nasr, M., Kolter, J. Z., and
865 Fredrikson, M. Universal and transferable adversarial
866 attacks on aligned language models, 2023. URL <https://arxiv.org/abs/2307.15043>.
- 869 Zou, A., Phan, L., Wang, J., Duenas, D., Lin, M., An-
870 driushchenko, M., Wang, R., Kolter, Z., Fredrikson,
871 M., and Hendrycks, D. Improving alignment and ro-
872 bustness with circuit breakers. *ArXiv*, abs/2406.04313,
873 2024a. URL <https://api.semanticscholar.org/CorpusID:270286008>.
- 875 Zou, W., Geng, R., Wang, B., and Jia, J. Poisonedrag:
876 Knowledge corruption attacks to retrieval-augmented
877 generation of large language models, 2024b. URL
878 <https://arxiv.org/abs/2402.07867>.

A. Threat Model

We consider AI agents that interact with external systems through tool use, enabling actions such as reading files, executing code, querying databases, or sending messages. Agents receive inputs from multiple sources including authenticated users, external content (emails, web pages, documents), system processes, and tool outputs. The agent maintains memory across interactions and decomposes high-level objectives into executable action sequences. Our threat model characterizes attackers who exploit these capabilities through operational means rather than training-time manipulation or direct code modification.

Attacker’s Goals: The attacker aims to violate confidentiality or integrity of user data and agent operations. Confidentiality violations include exfiltrating sensitive information and leaking data across privilege boundaries, such as exposing one user’s private data to another user or transmitting internal system information to external parties. Integrity violations encompass executing commands on behalf of unauthorized sources, pursuing objectives outside the agent’s authorized scope, or performing action sequences inconsistent with legitimate workflows. For example, an attacker might cause the agent to delete production data under the guise of debugging or redirect funds to unauthorized accounts. Attackers may also seek persistent compromise through memory poisoning, where malicious information injected into the agent’s memory affects future interactions and enables cascading attacks across sessions.

Attacker’s Background Knowledge: Attackers may possess varying levels of knowledge about the agent system. Attackers understand that agents maintain memory M across sessions, receive user prompts p , generate trajectories Tr_t of action-observation pairs, and interact with environment E_t through tools. Attackers may have full knowledge of memory contents including system prompts (particularly in open-source deployments) or can infer them through behavioral observation. Attackers can infer the agent’s current task by observing actions in the trajectory and changes in environment state. For instance, an attacker observing a cooking assistant might infer it has payment system access. Our framework makes no assumptions about attacker access to the underlying LLM f ; attacks operate through input manipulation regardless of whether the model is black-box or white-box accessible. However, attackers typically cannot access other users’ private data stored in memory or exact internal security policies.

Attacker’s Capabilities: Attackers can manipulate different inputs to the agent system. Attackers craft malicious user prompts p to perform jailbreak and direct prompt injection attacks, attempting to override intended objectives or bypass safety constraints. Attackers inject malicious content into

observations obs_t by compromising web pages, documents, emails, or API responses that the agent processes (indirect prompt injection). Attackers poison memory \mathbf{M} by injecting malicious information that persists across sessions and corrupts future decision-making. Attackers may observe the trajectory Tr_t (actions and tool calls) and environment state \mathbf{E}_t to infer agent behavior and refine attacks.

System Assumptions: We assume agents follow synchronous, turn-based agent execution model where each action completes and returns an observation before the next action begins, enabling discrete-time security analysis. Agents operate with access to multiple tools with varying privilege levels, such as read-only data access versus write permissions or user-scoped versus system-wide operations. Memory systems retain information across interactions, allowing agents to accumulate knowledge but also creating vulnerability to persistent compromise. Agents possess planning capabilities that decompose high-level tasks into subtask sequences, enabling complex workflows but also introducing compositional security challenges where individually safe actions combine to create violations. Agents process inputs from multiple sources with different trust levels, requiring security mechanisms that distinguish legitimate external content from malicious exploitation. We do not assume agents have perfect instruction attribution or objective inference capabilities; these are approximated through the oracles defined in our framework. We assume standard security infrastructure (authentication systems, access control mechanisms) is available but may be insufficient for agent-specific threats that exploit the dynamic nature of instruction interpretation and task decomposition.

B. Open Problems and Future Directions

Our framework clarifies what implementing agent security requires: approximating oracle functions, verifying properties across temporal composition, and evaluating security in realistic contexts.

Multi-agent and delegated authority: Our framework focuses on single-agent security where one agent acts on behalf of one authenticated user. It does not address settings where multiple agents coordinate, delegate subtasks, or operate on behalf of multiple principals. Formalizing delegation requires extending \mathbf{R} to model agent-to-agent authorization, defining how oracle functions and properties compose across agent boundaries, and specifying how authorization context transfers in delegation chains.

Oracle function implementation: The oracle functions require practical implementations, but fundamental challenges remain. Instruction attribution \mathcal{I} is an open interpretability problem: current approaches using attention mechanisms or influence functions (Jain & Wallace, 2019; Rashkin et al.,

2022) provide only coarse approximations. Source attribution \mathcal{L} requires solving information flow tracking for neural networks. Objective evaluation functions H_p, H_{Tr}, H_a require representing objectives and designing distance metrics d with appropriate thresholds τ .

Prompting models to output reasoning traces that explicitly identify which instructions they follow represents a promising direction, but such traces require verification mechanisms to ensure accurate reporting rather than post-hoc rationalization. For objective evaluation, hybrid approaches combining learned embeddings with symbolic constraints may better capture the structure of agent goals than purely neural or purely symbolic representations.

Compositional security: Action sequences create violations through composition where individually benign actions combine unsafely. An agent copying a file to backup then executing `chmod -R 777` creates a vulnerability: neither action violates action alignment individually, but their composition exposes sensitive data. Current property checks operate on individual actions without modeling how state changes accumulate across sequences, missing violations that only emerge through multi-step interactions.

Dynamic environments: In dynamic environments where multiple actors (other agents, users, or processes) modify state concurrently, reasoning about security properties becomes substantially more complex. An agent must account not only for its own actions but also for how the environment evolves independently. Current approaches lack mechanisms to account for state changes induced by external actors during agent execution.

Temporal security: Temporal violations introduce race conditions where properties hold at verification but fail at execution. In TOCTOU attacks (Lilienthal & Hong, 2025; Wei & Pu, 2005; Raducu et al., 2022), an agent verifies file ownership at t_1 but the file is replaced before use at t_2 . Memory poisoning injects malicious information at t_1 that corrupts decisions at t_2 across sessions. Long-horizon agent behaviors compound these challenges as trajectories extend over many steps and sessions, making it increasingly difficult to track which information influenced which decisions.

Realistic benchmarks with authorization context: Current benchmarks evaluate individual interactions in isolation, preventing detection of temporal violations and lacking authorization context necessary for property evaluation. Benchmarks do not specify which sources are authenticated, what objectives are authorized, or what information flows are permitted in \mathbf{R} .

Without explicit authorization boundaries, such as the agent’s role, task, and deployment scenario, evaluating rel-

evant security properties becomes impossible. Benchmarks should evaluate whether security properties hold across task sequences within sessions, test defenses against adaptive attacks, and measure utility-security tradeoffs as defenses often degrade capabilities (Jia et al., 2025; Li et al., 2025b). Creating realistic benchmarks with explicit authorization contexts that enable holistic security evaluation rather than isolated attack detection remains an open challenge.

C. Discussion and Limitations

Operational definition of instructions: Our framework defines instructions operationally through \mathcal{I} : an input is an instruction if it caused an action. This sidesteps what differentiates instructions from data, shifting focus to causal attribution. However, this cannot be evaluated without solving causal attribution, treats instruction-ness as binary when influence is continuous, and linguistic definitions fail because whether text functions as an instruction depends on context. LLM agents may require explicit architectural mechanisms to distinguish instructional from informational inputs.

Scope and applicability: Our framework applies to synchronous agents that await observations before proceeding. Agents with different execution models, such as may require modified property definitions and tracking. We focus on operational security of deployed agents rather than training-time attacks, model extraction, or adversarial examples. The framework assumes standard security infrastructure (authentication, access control); environments lacking these limit utility.

Oracle approximation quality: Our framework defines security through oracle functions that must be approximated in practice. Real systems use imperfect approximations (attention for \mathcal{I} , embedding similarity for H_p , H_{Tr} , H_a) with unpredictable accuracy. Understanding required approximation quality for practical security remains an open question.

Multi-agent and delegated authority: Our framework focuses on single-agent security where one agent acts on behalf of one authenticated user. It does not fully address settings where multiple agents coordinate, delegate subtasks to other agents, or operate on behalf of multiple principals with different authorization contexts. When Agent A delegates a subtask to Agent B, our framework does not specify which properties must hold for the delegation itself versus B’s execution, how authorization context transfers across agent boundaries, or how to compose security guarantees. The source permission graph \mathbf{R} would need extension to model delegation chains and specify which agents can act on behalf of which principals under what conditions.

Compositional safety and framework completeness: Our framework identifies compositional violations but does not fully formalize compositional safety, which requires modeling environment state dynamics and enumerating which state changes are safe individually but unsafe in combination. We do not claim the four properties are formally complete or sufficient to capture all possible security violations. Rather, they systematize existing attack classes and reveal structural patterns in current agent security research. Each property is necessary in the sense that violations in any property create known security breaches, but we make no claims about sufficiency. Demonstrating formal completeness would require exhaustively categorizing all possible violations and showing each maps to at least one property, which is beyond the scope of this systematization.