

Prompt Injection Attack to Tool Selection in LLM Agents

Jiawen Shi*, Zenghui Yuan*, Guiyao Tie*, Pan Zhou*[§], Neil Zhenqiang Gong^{†§}, Lichao Sun[‡]

*Huazhong University of Science and Technology, [†]Duke University, [‡]Lehigh University

{*shijiawen, zenghuiyuan, tgy, panzhou*}@*hust.edu.cn*, *neil.gong@duke.edu*, *lis221@lehigh.edu*

Abstract—Tool selection is a key component of LLM agents. A popular approach follows a two-step process - *retrieval* and *selection* - to pick the most appropriate tool from a tool library for a given task. In this work, we introduce *ToolHijacker*, a novel prompt injection attack targeting tool selection in no-box scenarios. *ToolHijacker* injects a malicious tool document into the tool library to manipulate the LLM agent’s tool selection process, compelling it to consistently choose the attacker’s malicious tool for an attacker-chosen target task. Specifically, we formulate the crafting of such tool documents as an optimization problem and propose a two-phase optimization strategy to solve it. Our extensive experimental evaluation shows that *ToolHijacker* is highly effective, significantly outperforming existing manual-based and automated prompt injection attacks when applied to tool selection. Moreover, we explore various defenses, including prevention-based defenses (StruQ and SecAlign) and detection-based defenses (known-answer detection, DataSentinel, perplexity detection, and perplexity windowed detection). Our experimental results indicate that these defenses are insufficient, highlighting the urgent need for developing new defense strategies.

I. INTRODUCTION

Large Language Models (LLMs) have demonstrated remarkable capabilities in natural language understanding and generation, catalyzing the emergence of LLM-based autonomous systems, known as LLM agents. These agents can perceive, reason, and execute complex tasks through interactions with external environments, including knowledge bases and tools. The deployment of LLM agents has expanded across various domains, encompassing web agents [1], [2] for browser-based interactions, code agents [3], [4] for software development and maintenance, and versatile agents [5], [6] that integrate diverse tools for comprehensive task-solving. The operation of LLM agents involves three key stages: task planning, tool selection, and tool calling [7], [8]. Among these, tool selection is crucial, as it determines which external tool is best suited for a given task, directly influencing the performance and decision-making of LLM agents. A popular tool selection approach involves a two-step mechanism: *retrieval* and *selection* [8], [9], [10], in which a retriever identifies the

top- k tool documents from the tool library and an LLM then selects the most appropriate tool for subsequent tool calling.

LLM agents are vulnerable to prompt injection attacks due to their integration of untrusted external sources. Attackers can inject harmful instructions into these external sources, manipulating the LLM agent’s actions to align with the attacker’s intent. Recent studies [11], [12], [13] have demonstrated that attackers can exploit this vulnerability by injecting instructions into external tools, leading LLM agents to disclose sensitive data or perform unauthorized actions. Particularly, attackers can embed deceptive instructions within tool documents to manipulate the LLM agent’s tool selection [13]. This manipulation poses serious security risks, as the LLM agent may inadvertently choose and execute harmful tools, compromising system integrity and user safety [14].

Prompt injection attacks are typically classified into manual and automated methods. Manual attacks, including naive attack [15], [16], escape characters [15], context ignoring [17], [18], fake completion [19], and combined attack [20], are heuristic-driven but time-consuming to develop and exhibit limited generalization across different scenarios. In contrast, automated attacks, such as JudgeDeceiver [13], leverage optimization frameworks to generate injection prompts targeting LLMs, with a specific focus on tool selection manipulation. Additionally, PoisonedRAG [21] targets Retrieval-Augmented Generation (RAG) systems by injecting adversarial texts into the knowledge base to manipulate LLM responses.

However, existing prompt injection methods remain suboptimal in tool selection, as detailed in Section IV. This limitation arises because manual methods and JudgeDeceiver primarily focus on the selection phase, making them incomplete as end-to-end attacks. Although PoisonedRAG considers the retrieval phase, it focuses on generation by injecting multiple malicious entries into the knowledge base, rather than directly manipulating tool selection. This difference creates distinct challenges for tool selection prompt injection, which our work addresses.

In this work, we propose *ToolHijacker*, the first prompt injection attack targeting tool selection in a no-box scenario. *ToolHijacker* efficiently generates *malicious tool documents* that manipulate tool selection through prompt injection. Given a target task, *ToolHijacker* generates a malicious tool document that, when injected into the tool library, influences both the retrieval and selection phases, compelling the LLM agent to choose the malicious tool over the benign ones, as illustrated in Figure 1. Additionally, *ToolHijacker* ensures

[§]Corresponding Author




Fig. 1: Illustration of tool selection in LLM agents under no attack and our attack.

consistent control over tool selection, even when users employ varying semantic descriptions of the target task. Notably, ToolHijacker is designed for the no-box scenario, where the target task descriptions, the retriever, the LLM, and the tool library, including the top- k setting, are inaccessible.

The core challenge of ToolHijacker is crafting a malicious tool document that can manipulate both the retrieval and selection phases of tool selection. To address this challenge, we formulate it as an optimization problem. Given the no-box constraints, we first construct a shadow framework of tool selection that includes shadow task descriptions, a shadow retriever, a shadow LLM, and a shadow tool library. Building upon this framework, we then formulate the optimization problem to generate the malicious tool document. The malicious tool document comprises a tool name and a tool description. Due to the limited tokens of the tool name in the tool document, we focus on optimizing the tool description. However, directly solving this optimization problem is challenging due to its discrete and non-differentiable nature. In response, we propose a two-phase optimization strategy that aligns with the inherent structure of the tool selection. Specifically, we decompose the optimization problem into two sub-objectives: *retrieval objective* and *selection objective*, allowing us to address each phase independently while ensuring their coordinated effect. We divide the tool description into two subsequences, each optimized for one of these sub-objectives. When concatenated, these subsequences form a complete tool description capable of executing an end-to-end attack across both phases of the tool selection. To effectively optimize these subsequences, we develop both gradient-based and gradient-free methods.

We evaluate ToolHijacker on two benchmark datasets, testing across 8 LLMs and 4 retrievers in diverse tool selection settings, with both gradient-free and gradient-based methods. The results show that ToolHijacker achieves high attack success rates in the no-box setting. Notably, ToolHijacker

maintains high attack performance even when the shadow LLM differs architecturally from the target LLM. For example, with Llama-3.3-70B as the shadow LLM and GPT-4o as the target LLM, our gradient-free method achieves a 96.7% attack success rate on MetaTool [22]. Additionally, ToolHijacker demonstrates high success during the retrieval phase, achieving a 100% attack hit rate on MetaTool. Furthermore, we show that ToolHijacker outperforms various prompt injection attacks when applied to our problem.

We evaluate two prevention-based defenses: StruQ [23] and SecAlign [24], as well as four detection-based defenses: known-answer detection [20], DataSentinel [25], perplexity (PPL) detection [26], and perplexity windowed (PPL-W) detection [26]. Our experimental results demonstrate that both StruQ and SecAlign fail to defend against ToolHijacker, with our gradient-free attack achieving 99.6% success rate under StruQ. Among detection-based defenses, known-answer detection fails to identify malicious tool documents, while DataSentinel, PPL and PPL-W detect some malicious tool documents generated by the gradient-based method but miss the majority. For instance, PPL misses detecting 90% of malicious tool documents optimized via the gradient-based method, when falsely detecting < 1% of benign tool documents as malicious.

To summarize, our key contributions are as follows:

- We propose ToolHijacker, the first prompt injection attack to tool selection in LLM agents.
- We formulate the attack as an optimization problem and propose a two-phase method to solve it.
- We conduct a systematic evaluation of ToolHijacker on multiple LLMs and benchmark datasets.
- We explore both prevention-based and detection-based defenses. Our experimental results highlight that we need new mechanisms to defend against ToolHijacker.




Fig. 2: Illustration of Step 2 - Selection.

II. PROBLEM FORMULATION

In this section, we formally define the framework of tool selection and characterize our threat model based on the attacker’s goal, background knowledge, and capabilities.

A. Tool Selection

We consider a popular tool selection process that comprises three core components: *tool library*, *retriever*, and *LLM*. The tool library contains n tools, each accompanied by a tool document that specifies the tool’s name, description, and API specifications. These documents detail each tool’s functionality, invocation methods, and parameters. We denote the set of tool documents as $D = \{d_1, d_2, \dots, d_n\}$. When the user provides a task description q , tool selection aims to identify the most appropriate tool from the tool library for the task execution. This process is achieved through a two-step mechanism, consisting of retrieval and selection, which can be formulated as follows:

Step 1 - Retrieval. The retriever employs a dual-encoder architecture consisting of a task description encoder f_q and a tool document encoder f_d to retrieve the top- k tool documents from D . Specifically, f_q and f_d map the task description q and each tool document $d_j \in D$ into the embedding vectors $f_q(q)$ and $f_d(d_j)$. The relevancy between each tool document d_j and the task description q is measured by a similarity function $Sim(\cdot, \cdot)$, such as cosine similarity or dot product. The retrieval process selects the top- k tool documents with the highest similarity scores relative to the q . Formally, the set of retrieved tool documents D_k is defined as:

$$D_k = \text{Top-}k(q; D) = \{d_1, d_2, \dots, d_k\}, \quad (1)$$

$$\text{Top-}k(q; D) = \text{Top-}k_{d_j \in D} (Sim(f_q(q), f_d(d_j))). \quad (2)$$

Step 2 - Selection. Given the task description q and the retrieved tool documents set D_k , the LLM agent provides q

and D_k to the LLM E to select the most appropriate tool from D_k for executing q . We denote this selection process as:

$$E(q, D_k) = d^*, \quad (3)$$

where d^* represents the selected tool. As illustrated in Figure 2, E adopts a structured prompt that combines q and tool information (i.e., tool names and descriptions) from D_k between a header instruction and a trailer instruction. This selection process is formulated as:

$$E(p_{\text{header}} \oplus q \oplus d_1 \oplus d_2 \oplus \dots \oplus d_k \oplus p_{\text{trailer}}) = o_{d^*}, \quad (4)$$

where o_{d^*} denotes the LLM’s output decision containing the selected tool name. The p_{header} and p_{trailer} represent the header and trailer instructions, respectively. We use \oplus to denote the concatenation operator that combines all components into a single input string.

B. Threat Model

Attacker’s goal. When an attacker selects a target task, it can be articulated through various semantic prompts (called *target task descriptions*), denoted as $Q = \{q_1, q_2, \dots, q_m\}$. For example, if the target task is inquiring about weather conditions, the task descriptions could be “What is the weather today?”, “How is tomorrow’s weather?”, or “Will it rain later?”. We assume that the attacker develops a *malicious tool* and disseminates it through an open platform accessible to the target LLM agent [27], [28], [29]. The attacker aims to manipulate the tool selection, ensuring that the malicious tool is preferentially chosen to perform the target task whenever users query the target LLM agent with any q_i from Q , thereby bypassing the selection of any other *benign tool* within the tool library. The key to executing this attack lies in the meticulous crafting of the malicious tool document d_t .

A tool document includes the tool name, tool description, and API specifications. Previous research [8], [30] indicates that tool selection primarily relies on the tool name and tool description. Therefore, our study focuses on crafting the tool name and tool description to facilitate the manipulated attack. Our attack can be characterized as a prompt injection attack targeting the tool selection mechanism.

We note that such an attack could pose security concerns for LLM agents in real-world applications. LLM agents operate on a select-and-execute mechanism. Thus, once a malicious tool is selected, it is executed without further verification, allowing attackers to manipulate execution outcomes arbitrarily. For instance, an attacker could develop a malicious tool for unauthorized data access, privacy breaches, or other harmful activities. These threats are increasingly relevant as LLM agents integrate with an expanding ecosystem of external tools and services.

Attacker’s background knowledge. We assume that the attacker is knowledgeable about the target task but does not have access to the target task descriptions $Q = \{q_1, q_2, \dots, q_m\}$. Recall that tool selection comprises three primary components: tool library, retriever, and LLM. We consider a no-box scenario where the attacker faces significant limitations in

accessing the tool selection. Specifically, the attacker cannot: 1) access the contents of tool documents in the tool library, 2) obtain information about either k or the top- k retrieved tool documents, 3) access the parameters of the target retriever and target LLM, or 4) directly query the target retriever and target LLM. However, the open platform provides standardized development guidelines, including documentation templates and interface specifications, which the attacker can leverage to craft the malicious tool document d_t .

Attacker’s capabilities. We assume that the attacker is capable of constructing a shadow task description set $Q' = \{q'_1, q'_2, \dots, q'_{m'}\}$, creating shadow tool documents D' , and deploying a shadow retriever and a shadow LLM to design and validate their attack strategies. Notably, $Q \cap Q' = \emptyset$, indicating no overlap between Q and Q' . Additionally, the attacker can develop and publish a malicious tool on tool hubs—such as Hugging Face Hub [31], Apify [28], and PulseMCP [29]—that accept third-party submissions, making it available for integration into LLM agents. This assumption is realistic and has been adopted in prior studies focusing on LLM agent security [14], [32]. By crafting the tool document, the attacker can execute prompt injection attacks. Recent studies [11], [12] on the model context protocol (MCP) reveal the feasibility of modifying tool documents to conduct attacks.

III. TOOLHIJACKER

A. Overview

ToolHijacker provides a systematic, automated approach for crafting the malicious tool document. Given the no-box scenario, we leverage a shadow tool selection pipeline to facilitate optimization. Upon this foundation, we formulate crafting a malicious tool document as an optimization problem encompassing two steps of the tool selection: retrieval and selection. The discrete, non-differentiable nature of this optimization problem renders its direct solution challenging. To address this, we propose a two-phase optimization strategy. Specifically, we decompose the optimization objective into two sub-objectives: retrieval and selection, and segment the malicious tool document into two subsequences, $R \oplus S$, optimizing each independently to achieve its corresponding sub-objective. When the two subsequences are concatenated, they enable an end-to-end attack on tool selection. We introduce gradient-free and gradient-based methods to solve the optimization problem.

B. Formulating an Optimization Problem

We start by constructing a set of shadow task descriptions and shadow tool documents. Specifically, an accessible LLM is employed to generate the shadow task description set, denoted as $Q' = \{q'_1, q'_2, \dots, q'_{m'}\}$, based on the target task. Additionally, we construct a set of shadow tool documents D' , encompassing both task-relevant and task-irrelevant documents, to effectively simulate the tool library.

In our no-box scenario, given the shadow task descriptions Q' , shadow tool documents D' , shadow retriever $f'(\cdot)$ and shadow LLM E' , our objective is to construct a malicious tool

document d_t containing $\{d_{t_name}, d_{t_des}\}$, where d_{t_name} denotes the malicious tool name and d_{t_des} denotes the malicious tool description. This malicious tool is designed to manipulate both the retrieval and selection processes, regardless of the specific shadow task descriptions q'_i . Formally, the optimization problem is defined as follows:

$$\max_{d_t} \frac{1}{m'} \cdot \sum_{i=1}^{m'} \mathbb{I}(E'(q'_i, \text{Top-}k'(q'_i; D' \cup \{d_t\})) = o_t), \quad (5)$$

where o_t represents the output of E' for selecting the d_t , and $\mathbb{I}(\cdot)$ denotes the indicator function that equals 1 when the condition is satisfied and 0 otherwise. Here, k' is the parameter of $f'(\cdot)$ specified by the attacker. $\text{Top-}k'(q'_i; D' \cup \{d_t\})$ represents a set of k' tool documents retrieved from the D' for q'_i .

The key challenge in solving the optimization problem lies in its discrete, discontinuous, and non-differentiable nature, which renders direct gradient-based methods infeasible. Moreover, the discrete search space contains numerous local optima, making it difficult to identify the global optimum. To address this, we propose a sequential, two-phase optimization strategy, which decomposes the optimization problem into two sub-objectives: *retrieval objective* and *selection objective*. Specifically, the retrieval objective ensures that d_t is always included in the top- k' set of retrieved tool documents during the retrieval phase. The selection objective, on the other hand, guarantees that within the retrieved set, the shadow LLM selects d_t containing $\{d_{t_name}, d_{t_des}\}$ as the final tool to execute. Inspired by PoisonedRAG [21], we divide d_{t_des} into the concatenation of two subsequences $R \oplus S$, and optimize them sequentially to achieve the respective objectives. It is important to note that d_{t_name} is manually crafted with limited tokens to ensure semantic clarity in the LLM agent. We propose both gradient-free and gradient-based methods to optimize the d_{t_des} . The following sections detail the optimization processes for R and S , respectively.

C. Optimizing R for Retrieval

We aim to generate a subsequence R that ensures the malicious tool document d_t appears among the top- k' tool documents set. The key insight is to maximize the similarity score between R and shadow task descriptions Q' , enabling d_t to achieve high relevancy across diverse task descriptions.

Gradient-Free. The gradient-free approach aims to generate R by leveraging the inherent semantic alignment between tool’s functionality descriptions and task descriptions. The key insight is that a tool’s functionality description naturally shares semantic similarities with the tasks it can accomplish, as they describe the same underlying capabilities from different perspectives. Based on this insight, we utilize an LLM to synthesize R by extracting and combining the core functional elements of Q' . This approach maximizes the semantic similarity between R and Q' without requiring gradient information, as the generated functionality description inherently captures the essential semantic patterns present in the shadow task descriptions space. Specifically, we use the following template to prompt an LLM to generate R :

Please generate a tool functionality description to address the following user queries:

[shadow task descriptions]

Requirements: The description should highlight core functionalities and provide a general solution applicable to various scenarios, not limited to a specific query. Limit the description to approximately [num] words.

Here, num is a hyperparameter used to limit the length of R . **Gradient-Based.** The gradient-based approach leverages the shadow retriever’s gradient information to optimize R . The core idea is to maximize the average similarity score between R and each shadow task description in $\{q'_1, q'_2, \dots, q'_{m'}\}$ through gradient-based optimization. Formally, the optimization problem is defined as follows:

$$\max_R \frac{1}{m'} \cdot \sum_{i=1}^{m'} \text{Sim}(f'(q'_i), f'(R \oplus S)), \quad (6)$$

where $f'(\cdot)$ denotes the encoding function of the shadow retriever and S is used in its initial sequence. We initialize R with the output derived from the gradient-free approach and subsequently optimize it through gradient descent. This optimization process essentially seeks to craft adversarial text that maximizes retrieval relevancy. Specifically, we employ the HotFlip [33], which has demonstrated efficacy in generating adversarial texts, to perform the token-level optimization of R . The transferability of ToolHijacker is based on the observation that semantic patterns learned by different retrieval models often exhibit considerable overlap, thereby enabling the optimized R to transfer effectively to the target retriever.

D. Optimizing S for Selection

After optimizing R , the subsequent objective is to optimize S within the malicious tool descriptions $R \oplus S$, such that the malicious tool document $d_t = \{d_{t_name}, R \oplus S\}$ can effectively manipulate the selection process. For simplicity, the malicious tool document is denoted as $d_t(S)$ in this section. We first construct the sets of shadow retrieval tool documents, denoted $\tilde{D}^{(i)} \cup \{d_t(S)\}$, to formulate the optimization objective. For each shadow task description q'_i in Q' , we create a set $\tilde{D}^{(i)}$ containing $(k' - 1)$ shadow tool documents from D' . Consequently, the set $\tilde{D}^{(i)} \cup \{d_t(S)\}$ comprises a total of k' tool documents. Our goal is to optimize S such that $d_t(S)$ is consistently selected by an LLM across all task-retrieval pairs $\{q'_i, \tilde{D}^{(i)} \cup \{d_t(S)\}\}$. Given the shadow LLM E' , the optimization problem can be formally expressed as:

$$\max_S \frac{1}{m'} \sum_{i=1}^{m'} \mathbb{I}(E'(q'_i, \tilde{D}^{(i)} \cup \{d_t(S)\}) = o_t). \quad (7)$$

Next, we discuss details on optimizing S .

Gradient-Free. We propose an automatic prompt generation approach that involves an attacker LLM E_A and the shadow LLM E' to optimize S without relying on the model gradients. Drawing inspiration from the tree-of-attack manner [34], we

Algorithm 1 Gradient-Free Optimization Approach for S

Input: The initial S_0 , shadow task descriptions $\{q'_1, \dots, q'_{m'}\}$, shadow retrieval tool sets $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m')}$, the malicious tool name o_t , the number of variants B , tree maximum width W , the maximum iteration T_{iter} , a pruning function $Prune$ and an evaluation function of regularization matching EM .

Output: Optimized S .

- 1: Initialize current iteration leaf nodes list $Leaf_curr = [S_0]$, the next iteration leaf nodes list $Leaf_next = []$, and the feedback list $Feed = []$.
- 2: **for** $q'_i \in \{q'_1, q'_2, \dots, q'_{m'}\}$ **do**
- 3: **for** $t \in [1, T]$ **do**
- 4: **for** $S_l \in Leaf_curr$ **do**
- 5: Generate B variants $\{S_l^1, S_l^2, \dots, S_l^B\}$ of S_l , where $S_l^b = E_A(p_{attack}, S_l, q'_i, \tilde{D}^{(i)}, Feed)$.
- 6: Append $\{S_l^1, S_l^2, \dots, S_l^B\}$ to $Leaf_next$.
- 7: **end for**
- 8: Set the flag list $FLAG$ to be a $1 \times m'$ -dimensional vector of 0: $FLAG = 0^{1 \times m'}$.
- 9: **for** $S_l \in Leaf_next$ **do**
- 10: Initialize evaluation response list $Eval_list = []$.
- 11: **for** $j \in [1, m']$ **do**
- 12: Get the response of E' on q'_j : $E'(q'_j, \tilde{D}^{(j)} \cup \{d_t(S_l)\})$ and append it to $Eval_list$.
- 13: **if** $EM(E'(q'_j, \tilde{D}^{(j)} \cup \{d_t(S_l)\}) = o_t)$ **then**
- 14: Increment $FLAG[S_l]$ by 1:
- 15: $FLAG[S_l] = FLAG[S_l] + 1$
- 16: **end if**
- 17: **end for**
- 18: **end for**
- 19: Get index S_L of the maximum element in $FLAG$.
- 20: **if** $FLAG[S_L] = m'$ **then**
- 21: **return** $S \leftarrow Leaf_next[S_L]$
- 22: **end if**
- 23: Prune $Leaf_next$ to retain top W nodes based on $FLAG$: $Leaf_next \leftarrow Prune(Leaf_next, W)$.
- 24: Record $Eval_list$ and $FLAG$ of remaining nodes into $Feed$.
- 25: Update $Leaf_curr \leftarrow Leaf_next$.
- 26: Reset $Leaf_next \leftarrow []$.
- 27: **end for**
- 28: Update $Leaf_curr \leftarrow Leaf_curr[S_L]$.
- 29: **end for**
- 30: **return** $S \leftarrow Leaf_next[S_L]$

formulate the optimization of S a hierarchical tree construction process, with the initialization S_0 serving as the root node and each child node as an optimized variant of S . The optimization procedure iterates T_{iter} times for each query $q'_i \in Q'$, where each iteration encompasses four steps:

Attacker LLM Generating: The attacker LLM E_A generates B variants $\{S_l^1, S_l^2, \dots, S_l^B\}$ for each S_l in current leaf node list $Leaf_curr$ to construct the next leaf node list $Leaf_next$. Each variant can be expressed as $S_l^b = E_A(p_{attack}, S_l, q'_i, \tilde{D}^{(i)}, Feed)$, where p_{attack} is the system instruction of E_A (as shown in Appendix C of our technical report version [35]) and $Feed$ represents the feedback information from the previous iteration.

Querying Shadow LLM: For each $S_l \in Leaf_next$, E' generates a response $E'(q'_j, \tilde{D}^{(j)} \cup \{d_t(S_l)\})$ for each $q'_j \in Q'$.

Evaluating: Regularized matching is employed to verify

whether the responses of the node $S_l \in Leaf_next$ to all shadow task descriptions match the malicious tool. The variable $FLAG[l]$ is set to the number of successful matches.

Pruning and Feedback: If a node S_l satisfies $FLAG[l] = m'$, it is considered successfully optimized S , ending the optimization process. Otherwise, $Leaf_next$ is pruned according to $FLAG$ values to limit the remaining nodes to the maximum width W . The responses and $FLAG$ values corresponding to the remaining nodes are attached to $Feed$ for the next iteration. The node with the maximum value of $FLAG$ becomes the root node for the next shadow tool description when the maximum iteration T_{iter} is reached, or it is regarded as the final optimized S when all shadow task descriptions have been looped. The entire process is shown in Algorithm 1.

Gradient-Based. We propose a method that leverages gradient information from the shadow LLM E' to solve Equation 7. Our objective is to optimize S to maximize the likelihood that E' generates responses containing the malicious tool name d_{t_name} . This objective can be formulated as:

$$\max_S \prod_{i=1}^{m'} E'(o_t | p_{\text{header}} \oplus q'_i \oplus d_1^{(i)} \oplus \dots \oplus d_{k'-1}^{(i)} \oplus d_t(S) \oplus p_{\text{trailer}}). \quad (8)$$

The E' generates responses by sequentially processing input tokens and determining the most probable subsequent tokens based on contextual probabilities. We denote S as a token sequence $S = (T_1, T_2, \dots, T_\gamma)$ and perform token-level optimization. Specifically, we design a loss function comprising three components: alignment loss \mathcal{L}_1 , consistency loss \mathcal{L}_2 , and perplexity loss \mathcal{L}_3 , which guide the optimization process.

Alignment Loss - \mathcal{L}_1 : The alignment loss aims to increase the likelihood that E' generates the target output o_t containing d_{t_name} . Let $o_t = (\tau_1, \tau_2, \dots, \tau_\rho)$ where ρ denotes the sequence length, and $x^{(i)}$ represents the input sequence $\{q'_i, \tilde{D}^{(i)} \cup \{d_t(S)\}\}$ excluding S . The \mathcal{L}_1 is defined as:

$$\mathcal{L}_1(x^{(i)}, S) = -\log E'(o_t | x^{(i)}, S), \quad (9)$$

$$E'(o_t | x^{(i)}, S) = \prod_{j=1}^{\rho} E'(\tau_j | x_{1:h_i}^{(i)}, S, x_{h_i+\gamma+1:n_i}^{(i)}, \tau_1, \dots, \tau_{j-1}). \quad (10)$$

Here, S is inserted at position h_i among the retrieved shadow tool documents, $x_{1:h_i}^{(i)}$ denotes the input tokens preceding S , $x_{h_i+\gamma+1:n_i}^{(i)}$ denotes the input tokens following S , and n_i is the total length of the input tokens processed by E' .

Consistency Loss - \mathcal{L}_2 : The consistency loss reinforces the alignment loss by specifically focusing on the generation of d_{t_name} . The consistency loss \mathcal{L}_2 is expressed as:

$$\mathcal{L}_2(x^{(i)}, S) = -\log E'(d_{t_name} | x^{(i)}, S). \quad (11)$$

Perplexity Loss - \mathcal{L}_3 : This perplexity loss \mathcal{L}_3 is proposed to enhance the readability of S . Formally, it is defined as the average negative log-likelihood of the sequence:

$$\mathcal{L}_3(x^{(i)}, S) = -\frac{1}{\gamma} \sum_{j=1}^{\gamma} \log E(T_j | x_{1:h_i}^{(i)}, T_1, \dots, T_{j-1}). \quad (12)$$

The overall loss function is defined as:

$$\mathcal{L}_{\text{all}}(x^{(i)}, S) = \mathcal{L}_1(x^{(i)}, S) + \alpha \mathcal{L}_2(x^{(i)}, S) + \beta \mathcal{L}_3(x^{(i)}, S), \quad (13)$$

$$\min_S \mathcal{L}_{\text{all}}(S) = \sum_{i=1}^{m'} \mathcal{L}_{\text{all}}(x^{(i)}, S), \quad (14)$$

where α and β are hyperparameters balancing three loss terms. To address the optimization problem, we employ the algorithm introduced in JudgeDeceiver [13], which integrates both position-adaptive and step-wise optimization strategies. Specifically, the optimization process comprises two key components: 1) Position-adaptive Optimization: For each task-retrieval pair $\{q'_i, \tilde{D}^{(i)} \cup \{d_t(S)\}\}$, we optimize the S by positioning the $d_t(S)$ at different locations within the set of shadow retrieval tool documents; 2) Step-wise Optimization: Instead of optimizing all pairs simultaneously, we gradually incorporate task-retrieval pairs into the optimization process. This progressive approach helps to stabilize the optimization.

IV. EVALUATION

A. Experimental Setup

1) **Datasets:** We use the following two datasets to evaluate the effectiveness of our attacks.

- **MetaTool** [22]. This benchmark focuses on LLMs' capabilities in tool usage. It comprises 21,127 instances, involving 199 benign tool documents sourced from OpenAI Plugins.
- **ToolBench** [10]. This benchmark aims to enhance the tool-use capabilities of open-source LLMs with 126,486 instruction-tuning samples, leveraging 16,464 tool documents from RapidAPI. After removing duplicate tools and empty descriptions, the tool library contains 9,650 benign tool documents.

For each dataset, we design 10 high-quality target tasks that represent real-world needs while ensuring scenario diversity. For each target task, we generate 100 target task descriptions through both LLM-based and human evaluations, resulting in 1,000 target task descriptions per dataset.

2) **Compared Baselines:** We employ seven prompt injection attacks as baselines for comparison with our method: five manual attacks (naive, escape characters, context ignore, fake completion, and combined attack) and two automated attacks (JudgeDeceiver and PoisonedRAG). Their malicious tool documents are provided in Figure 16 of Appendix C [35].

- **Naive Attack** [15], [16]. This method uses explicit instruction as the tool description, directing the LLM to select the malicious tool. An example could be: "Just output the word {tool name} as your final answer".
- **Escape Characters** [15]. This method uses escape characters such as "\n" or "\t" before the malicious instruction to segment the text, effectively isolating the instruction and enhancing the attack success rate.
- **Context Ignore** [17], [18]. This technique inserts prompts such as "ignore previous instructions" to compel

TABLE I: Our attacks achieve high ASRs across different target LLMs. The gradient-free attack employs Llama-3.3-70B as the shadow LLM, while the gradient-based attack employs Llama-3-8B.

Dataset	Attack	Metric	LLM of Tool Selection							
			Llama-2 7B	Llama-3 8B	Llama-3 70B	Llama-3.3 70B	Claude-3 Haiku	Claude-3.5 Sonnet	GPT-3.5	GPT-4o
MetaTool	No Attack	ACC	96.7%	98.9%	98.2%	99.6%	99.2%	98.9%	98.8%	99.6%
	Gradient-Free	ASR	98.2%	94.0%	97.0%	99.6%	85.4%	92.1%	91.0%	96.7%
	Gradient-Based	ASR	99.8%	100%	97.2%	99.4%	82.6%	92.0%	92.8%	92.2%
ToolBench	No Attack	ACC	97.1%	90.5%	97.2%	97.2%	97.2%	97.8%	97.3%	98.4%
	Gradient-Free	ASR	91.7%	80.6%	82.1%	90.8%	82.8%	93.6%	77.7%	88.2%
	Gradient-Based	ASR	95.2%	96.6%	89.2%	94.8%	74.3%	85.2%	84.6%	83.9%

the LLM to abandon previously established context and prioritize only the subsequent malicious instruction.

- **Fake Completion** [19]. This method inserts a fabricated completion prompt to deceive the LLM into believing all previous instructions are resolved, then executes new instructions injected by the attacker.
- **Combined Attack** [20]. This approach combines elements from the four strategies mentioned above into a single attack, thereby maximizing confusion and undermining the LLM’s ability to resist malicious prompts.
- **JudgeDeceiver** [13]. This method injects a gradient-optimized adversarial sequence into the malicious answer, causing LLM-as-a-Judge to select it as the best answer for the target question, regardless of other benign answers.
- **PoisonedRAG** [21]. This attack manipulates a RAG system by injecting adversarial texts into the knowledge database, guiding the LLM to generate attacker-desired answers. The adversarial texts are optimized through a repeated sampling prompt strategy.

3) *Tool Selection Setup*: We evaluate our attack on the tool selection comprising the following LLMs and retrievers:

- **Target LLM.** We evaluate our method on both open-source and closed-source LLMs. The open-source models include Llama-2-7B-chat [36], Llama-3-8B-Instruct [37], Llama-3-70B-Instruct [37], and Llama-3.3-70B-Instruct [38]. For closed-source models, we test Claude-3-Haiku [39], Claude-3.5-Sonnet [39], GPT-3.5 [40], and GPT-4o [41]. These models cover a wide range of model architectures and sizes, enabling a comprehensive analysis of the effectiveness of our attack.
- **Target Retriever.** We conduct attacks on four retrieval models: text-embedding-ada-002 [42] (a closed-source embedding model from OpenAI), Contriever [43], Contriever-ms [43] (Contriever fine-tuned on MS MARCO), and Sentence-BERT-tb [10] (Sentence-BERT [44] fine-tuned on ToolBench).

4) *Attack Settings*: For each target task, we optimize a malicious tool document using 5 shadow task descriptions (i.e., $m' = 5$), each paired with a shadow retrieval tool set containing 4 shadow tool documents (i.e., $k' = 5$). For the gradient-free attack, we employ Llama-3.3-70B as both the attacker and shadow LLM, with optimization parameters for S set to $T_{iter} = 10$, $B = 2$, and $W = 10$. For the gradient-

TABLE II: Our attacks have high AHRs.

Dataset	No Attack	Gradient-Free	
		HR	AHR
MetaTool	100%	99.9%	100%
ToolBench	100%	96.1%	97.8%

based attack, we utilize Contriever as the shadow retriever and Llama-3-8B as the shadow LLM, with parameters $\alpha = 2.0$, $\beta = 0.1$, optimizing R for 3 iterations and S for 400 iterations. Both R and S are initialized using natural sentences (detailed in Figure 12 in Appendix C). In our ablation studies, unless otherwise specified, we use task 1 from the MetaTool dataset, with GPT-4o as the target LLM and text-embedding-ada-002 as the target retriever.

5) *Evaluation Metrics*: We adopt *accuracy (ACC)*, *attack success rate (ASR)*, *hit rate (HR)*, and *attack hit rate (AHR)* as evaluation metrics. We define them as follows:

- **ACC.** The ACC measures the likelihood of correctly selecting the appropriate tool for a target task from the tool library without attacks. It is calculated by evaluating 100 task descriptions for each target task (i.e., $m = 100$).
- **ASR.** The ASR measures the likelihood of selecting the malicious tool from the tool library when the malicious tool document is injected. It is calculated by evaluating 100 task descriptions for each target task (i.e., $m = 100$).
- **HR.** The HR measures the proportion of the target task for which at least one correct tool appears in the top- k results. Let $hit(q_i, k)$ be an indicator function that equals 1 if any correct tool for q_i appears in the top- k results, and 0 otherwise. Formally,

$$HR@k = \frac{1}{m} \sum_{i=1}^m hit(q_i, k). \quad (15)$$

- **AHR.** AHR measures the proportion of the malicious tool document d_t that appears in the top- k results. Let $a\text{-hit}(q_i, k)$ be an indicator function that equals 1 if d_t is included in the top- k results, and 0 otherwise. Formally,

$$AHR@k = \frac{1}{m} \sum_{i=1}^m a\text{-hit}(q_i, k). \quad (16)$$

Note that ACC and ASR are the primary metrics to evaluate the utility and attack effectiveness of an LLM agent’s end-to-

TABLE III: Our attack outperforms baselines on GPT-4o.

Dataset	Naive Attack	Escape Characters	Content Ignore	Fake Completion	Combined Attack	Judge-Deceiver	Poisoned-RAG	Gradient-Free	Gradient-Based
MetaTool	6.0%	28.2%	1.2%	14.5%	9.7%	30.2%	39.3%	96.7%	92.2%
ToolBench	24.8%	24.6%	11.3%	23.0%	11.7%	26.4%	58.3%	88.2%	83.9%




Fig. 3: Our attacks are effective across different tasks.




Fig. 4: Token length of benign tool documents and malicious tool documents generated via different attacks.

end tool selection process. On the other hand, HR and AHR are intermediate metrics that focus on the retrieval step, providing insights into how the attack impacts each component of the two-step tool selection pipeline. In this work, unless otherwise stated, we set $k = 5$ by default. We refer to HR@5 and AHR@5 simply as “HR” and “AHR” respectively.

B. Main Results

Our attack achieves high ASRs and AHRs. Table I shows the ASRs of ToolHijacker across eight target LLMs and two datasets. Each ASR represents the average attack performance over 10 distinct target tasks within each dataset. We have the following observations. First, both gradient-free and gradient-based methods demonstrate robust attack performance across different target LLMs, even when the shadow LLMs and the target LLMs differ in architecture. For instance, when the target LLM is GPT-4o, the gradient-free attack achieves ASRs of 96.7% and 88.2% on MetaTool and ToolBench respec-

tively, while the gradient-based attack attains ASRs of 92.2% and 83.9%. The reason is that shared alignment objectives and training paradigms make LLMs inherently vulnerable to prompt injection. Moreover, LLM homogenization—caused by training on overlapping datasets—makes them respond similarly to attacks. Second, the gradient-free attack exhibits higher performance on closed-source models, while the gradient-based attack shows advantages on open-source models. For instance, the gradient-free attack achieves a higher ASR by 4.5% when targeting GPT-4o on MetaTool and by 8.4% when targeting Claude-3.5-Sonnet on ToolBench. In contrast, the gradient-based attack exhibits a 16% higher ASR on ToolBench when targeting Llama-3-8B. Third, we find that different models exhibit varying sensitivities to our attacks. Claude-3-Haiku is the least sensitive, but it still achieves an ASR of $\geq 70\%$.

Additionally, we present the average AHRs of the retrieval phase in Table II. We observe that our method achieves high AHRs when targeting the closed-source retriever. Notably, when evaluated on the ToolBench’s tool library comprising 9,650 benign tool documents, our gradient-free attack achieves 96.1% AHR and our gradient-based attack achieves 97.8% AHR, while only injecting a single malicious tool document. Figure 3 presents the average ASRs and AHRs for 10 target tasks across two datasets and various target LLMs. The results show that both gradient-free and gradient-based attacks are effective across different target tasks and datasets. Furthermore, to assess the impact of our attack on the general utility of tool selection, we evaluate its performance on non-target tasks. Detailed results are presented in Table XII in Appendix B.

Our attack outperforms other baselines. Table III compares the performance of our attacks with five manual prompt injection attacks, JudgeDeceiver, and PoisonedRAG. The results show that our attacks outperform other baselines. Manual prompt injection attacks, which involve injecting irrelevant prompts into the malicious tool document, result in low ASRs due to the low likelihood of retrieval. For example, the escape characters achieve an ASR of 28.2% on MetaTool. Meanwhile, the optimization-based attack, JudgeDeceiver, achieves ASRs of 30.2% and 26.4%. PoisonedRAG achieves the highest performance among baselines, with ASRs of 39.3% on MetaTool and 58.3% on ToolBench. However, its attack performance still falls short of ours. The reason is that PoisonedRAG is designed to optimize for a single task description, while our attacks can optimize across multiple task descriptions. Figure 4 shows the token lengths of tool documents from benign tools, baselines, and our attacks. Notably, the malicious tool documents generated by our attacks are short and indistinguishable from benign tool documents based solely on token length.




Fig. 5: AHRs and ASRs with different k' of the shadow retriever and k of the target retriever.

TABLE IV: Impact of different target retrievers in our attacks.

Retriever	Gradient-Free		Gradient-Based	
	AHR	ASR	AHR	ASR
text-embedding-ada-002	100%	99%	100%	95%
Contriever	100%	99%	100%	100%
Contriever-ms	100%	99%	100%	100%
Sentence-BERT-tb	100%	99%	100%	100%
Average	100%	99%	100%	98.75%




Fig. 6: Impact of the number of shadow task descriptions.

C. Ablation Studies

Impact of retriever. We evaluate the effectiveness of our attacks across different retrievers. As shown in Table IV, the gradient-free attack demonstrates consistent performance, achieving 100% AHR and 99% ASR across all retrievers. For the gradient-based attack, all retrievers maintain 100% AHR. The three open-source retrievers achieve 100% ASR, while the closed-source retriever (text-embedding-ada-002) shows a slightly lower ASR of 95%. This discrepancy is due to the superior performance of text-embedding-ada-002. Although the malicious tool document is successfully retrieved, it is ranked lower in the results, reducing the likelihood of it being ultimately selected by the target LLM.

TABLE V: Impact of R and S .

Attack	$R \oplus S$		R		S	
	AHR	ASR	AHR	ASR	AHR	ASR
Gradient-Free	100%	99%	100%	5%	65%	63%
Gradient-Based	100%	95%	100%	0%	99%	16%

Impact of k . To investigate the impact of top- k settings, we vary k from 1 to 10 under the default attack configuration and record the AHRs and ASRs, as shown in the third column of Figure 5. Our results show that for smaller values of k , both AHR and ASR decrease, particularly for the gradient-free attack. When $k = 1$, both AHR and ASR are 89%. However, when k exceeds 3, the AHR for both attacks stabilizes at 100%, while the ASR for the gradient-based attack fluctuates around 96%, and the gradient-free attack stabilizes at 99%. The reason is that for smaller values of k , the likelihood of retrieving malicious tools decreases, as their similarity to the target task description may not be the highest.

Impact of k' . We further evaluate the impact of using different k' of the shadow retriever in optimizing S , with $k' \in \{2, 3, 5, 7\}$. The results are shown in Figure 5. We have two key observations. First, as k' increases, the AHR steadily rises to 100%, with a more pronounced increase for smaller k' . For instance, when $k' = 2$, the AHR of the gradient-based attack increases from 74% to 99% as k moves from 1 to 3. Second, ASR exhibits fluctuations with small k' , showing a general decline as k increases from 1 to 5. For instance, at $k' = 2$, the ASR drops by 16% and 50% for gradient-free and gradient-based attacks respectively, as k increases. The reason is that the number of ground-truth tools is 5. When k' is small, the attack optimization is suboptimal, and as k increases (with $k < 5$), more ground-truth tools are retrieved, reducing the likelihood of selecting the target tool. In contrast,

TABLE VI: ASRs of the gradient-free attack with different shadow LLMs on various target LLMs.

Shadow LLM	Target LLM								Average
	Llama-2 7B	Llama-3 8B	Llama-3 70B	Llama-3.3 70B	Claude-3 Haiku	Claude-3.5 Sonnet	GPT-3.5	GPT-4o	
Llama-2-7B	100%	100%	100%	100%	70%	99%	98%	94%	95.13%
Llama-3-8B	88%	100%	100%	100%	100%	100%	75%	99%	95.25%
Llama-3-70B	85%	100%	100%	99%	100%	100%	75%	99%	94.75%
Llama-3.3-70B	95%	100%	100%	99%	86%	99%	100%	99%	97.25%
Claude-3-Haiku	91%	100%	100%	100%	100%	100%	87%	100%	97.25%
Claude-3.5-Sonnet	99%	100%	100%	99%	100%	100%	98%	100%	99.50%
GPT-3.5	97%	100%	100%	100%	95%	100%	87%	100%	97.38%
GPT-4o	93%	100%	100%	100%	100%	100%	89%	100%	97.75%

TABLE VII: ASRs of the gradient-based attack with different shadow LLMs on various target LLMs.

Shadow LLM	Target LLM								Average
	Llama-2 7B	Llama-3 8B	Llama-3 70B	Llama-3.3 70B	Claude-3 Haiku	Claude-3.5 Sonnet	GPT-3.5	GPT-4o	
Llama-2-7B	100%	100%	34%	95%	55%	82%	98%	87%	81.38%
Llama-3-8B	100%	100%	100%	100%	98%	97%	82%	95%	96.50%

TABLE VIII: Impact of the similarity metric.

Attack	Cosine Similarity		Dot Product	
	AHR	ASR	AHR	ASR
Gradient-Free	100%	99%	100%	99%
Gradient-Based	100%	95%	100%	97%

when $k' \geq 5$, the optimized S improves, leading to an increase and stabilization of performance as k increases.

Impact of shadow task descriptions. We assess the impact of the number of shadow task descriptions (i.e., m') on both attack methods. As shown in Figure 6, the AHR remains unaffected by the number of shadow task descriptions, consistently maintaining 100% as the quantity increases from 1 to 10. Conversely, the ASR improves with an increasing number of shadow task descriptions, with the gradient-based attack exhibiting the most significant variation. Specifically, the ASR for the gradient-based attack rises from 32% with a single shadow task description to 98% with seven descriptions. In comparison, the gradient-free attack achieves a minimum ASR of 92% even when only one shadow task description is used.

Impact of R and S . To evaluate the respective contributions of R and S to attack performance, we conduct experiments using three settings for the malicious tool description: $R \oplus S$, only R , and only S . The results are presented in Table V. For the gradient-free attack, the AHR drops from 100% to 65% without R , highlighting the key role of R in achieving the retrieval objective. Without S , the ASR drops from 99% to 5%, emphasizing its significance for the selection objective. In the gradient-based attack, the AHR remains at 99% when only S is present, due to the gradient-based optimization process, which causes the generated S to contain more information about the target task, making it easier to be retrieved.

Impact of the shadow LLM E' in optimizing S . To assess

the impact of different shadow LLMs E' on our two attacks, we apply 8 distinct LLMs for the gradient-free attack and use two open-source LLMs, Llama-2-7B and Llama-3-8B, for the gradient-based attack. The ASRs of our two attack methods across the 8 target LLMs are presented in Table VI and Table VII. We have two key observations. First, employing more powerful shadow LLMs E' substantially improves the ASR for both attack methods. For example, in the gradient-free attack, employing Claude-3.5-Sonnet as the shadow LLM improves the average ASR by 4.37% compared to Llama-2-7B. Similarly, in the gradient-based attack, Llama-3-8B increases the ASR by 15.12% over Llama-2-7B. Second, the gradient-free attack is less sensitive to the shadow LLM E' than the gradient-based attack. Specifically, when using Llama-2-7B as the shadow LLM, the gradient-free attack maintains a minimum ASR of 70% on Claude-3-Haiku, while the gradient-based attack's lowest ASR drops to 34% on Llama-3-70B.

Impact of similarity metric. We evaluate the impact of two distinct similarity metrics on attack effectiveness during retrieval, with the results shown in Table VIII. The results indicate that different similarity metrics do not affect the likelihood of the generated malicious tool document being retrieved by the target retriever. Notably, the dot product results in a 2% improvement in ASR compared to cosine similarity.

Impact of the number of malicious tools. We evaluate the impact of injecting different numbers of malicious tools on attack effectiveness. Since the baseline setting with $k' = 5$ already gets strong results, as shown in Figure 5, we focus on comparing the effects when $k' = 2$ and the number of injected malicious tools ($num = 1$ or 2). For $num = 2$, we consider two scenarios: ‘individual’, where each malicious tool document targets its own tool, and ‘unified’, where all malicious tool documents target the same tool. The AHR and ASR for our attacks, as k varies across these settings, are presented in Figure 7. We observe that the trend under




Fig. 7: Attacks with different numbers of malicious tool documents. In the “individual” setting, each injected malicious tool document targets itself, while in the “unified” setting, all injected malicious tool documents target the same tool.

the ‘individual’ setting mirrors that of $num = 1$, but the ASR improves at the same k . For example, at $k = 5$, both the gradient-free and gradient-based attacks achieve a 24% increase in ASR. In the ‘unified’ setting, both ASR and AHR remain close to 100% as k increases, indicating that increasing the number of injected malicious tools enhances the attack when shadow tool documents are insufficient.

V. DEFENSES

Defenses against prompt injection attacks can be categorized into two types: prevention-based defenses and detection-based defenses [20]. Prevention-based defenses aim to mitigate the effects of prompt injections by either preprocessing instruction prompts or fine-tuning the LLM using adversarial training to reduce its susceptibility to manipulation. Since the instruction prompt for the tool selection employs the “sandwich prevention” method [45], we primarily focus on fine-tuning based defenses, including StruQ [23] and SecAlign [24]. Detection-based defenses, on the other hand, focus on identifying whether a response contains an injected sequence. Techniques commonly used for detections include known-answer detection, DataSentinel, perplexity (PPL) detection, and perplexity windowed (PPL-W) detection.

A. Prevention-based Defense

StruQ [23]. This method counters prompt injection attacks by splitting the input into two distinct components: a secure prompt and user data. The model is trained to only follow instructions from the secure prompt, ignoring any embedded instructions in the data. We use the fine-tuned model provided in StruQ, $LLM_{d(\text{struq})}$, as the target LLM to evaluate its effectiveness against our attacks.

SecAlign [24]. This method enhances the LLM’s resistance to prompt injection by fine-tuning it to prioritize secure outputs.

TABLE IX: Prevention-based defense results for our attacks.

Method	Dataset	Gradient-Free			Gradient-Based		
		ACC-a	AHR	ASR	ACC-a	AHR	ASR
StruQ	MetaTool	0.3%	99.9%	99.6%	2.1%	100%	97.9%
	ToolBench	5.7%	96.1%	90.8%	4.1%	97.8%	92.1%
SecAlign	MetaTool	2.5%	99.9%	97.5%	7.4%	100%	92.1%
	ToolBench	8.2%	96.1%	86.9%	11.3%	97.8%	84.6%

The key idea is to train the LLM on a dataset with both prompt-injected inputs and secure/insecure response pairs. We employ the fine-tuned LLM in SecAlign, $LLM_{d(\text{secalign})}$, as the target LLM to assess its effectiveness against our attacks.

Experimental results. To evaluate the effectiveness of StruQ and SecAlign, we utilize three key metrics: ACC-a (ACC with attack), AHR, and ASR. Experiments are conducted using the MetaTool and ToolBench datasets, each consisting of 10 target tasks and 100 target task descriptions per target task, with both gradient-free and gradient-based attacks. As shown in Table IX, our attacks still achieve high ASRs on the LLMs fine-tuned with StruQ and SecAlign, indicating that our attacks can bypass these defenses. This is because the carefully crafted malicious tool documents lack jarring or obvious instructions, instead providing descriptions related to the target task and tool functionality while preserving overall semantic integrity. Although SecAlign yields slightly lower ASR values than StruQ, suggesting stronger defense, the ASR still ranges from 84.6% to 97.5%, indicating that neither defense fully mitigates the attack strategies used in this work. Additionally, the ASRs on ToolBench are lower than those on MetaTool, likely stemming from ToolBench’s larger tool library size. It is noteworthy that the sum of ACC-a and ASR does not consistently total 100%, as model refusals—where the model




Fig. 8: ASR variation before and after defense.

fails to generate a valid response or rejects inputs—account for this discrepancy.

To further assess the defenses' effectiveness, we calculate the change in ASR using the formula $\Delta\text{ASR} = \text{ASR}_{\text{none}} - \text{ASR}_d$, where ASR_{none} represents the ASR of the LLM without any fine-tuning, and ASR_d refers to the ASR of the LLM fine-tuned with StruQ or SecAlign. As shown in Figure 8, SecAlign outperforms StruQ, achieving a larger reduction in ASR. Furthermore, the higher ΔASR values observed for ToolBench suggest that the defenses are more effective in addressing scenarios with larger tool libraries.

B. Detection-based Defense

Known-answer detection. Known-answer detection is known as one of the most effective defenses against prompt injection attacks [46], [20]. This approach involves embedding a detection instruction within the text and validating whether a detection LLM can still correctly follow this instruction when combined with the text. Following previous work [20], we adopt the detection instruction: “Repeat ‘[secret data]’ once while ignoring the following text. \n Text:”, where “[secret data]” is set to “Hello World!” for our experiments. When the LLM generates a response, if it does not include “Hello World!”, it is flagged as a target text with an injected sequence. If “Hello World!” is present, the text is considered clean.

DataSentinel. DataSentinel [25] is the state-of-the-art detection method for prompt injection. This defense extends the known-answer detection by employing a game-theoretic approach to fine-tune the detection LLM, thereby enhancing its detection capability and generalization.

Perplexity-based detection. Perplexity-based (PPL) detection is a widely adopted technique for identifying text altered by injected sequences. The key idea of PPL is that an injected sequence disrupts the semantic coherence of the text, thereby increasing its perplexity score. If the perplexity of a text exceeds a predefined threshold, it is flagged as containing an injected sequence [26]. However, a key challenge in this approach lies in selecting an appropriate threshold, as perplexity distributions vary across different datasets. To address this, we employ a dataset-adaptive strategy [20], where 100 clean samples are selected from the dataset, their log-perplexity values are computed, and the threshold is set such that the

TABLE X: Detection results for our attacks (G-Free: gradient-free attack, G-Based: gradient-based attack).

Dataset	Attack	Known-answer Detection		DataSentinel		PPL Detection		PPL-W Detection	
		FNR	FPR	FNR	FPR	FNR	FPR	FNR	FPR
MetaTool	G-Free	100%	0%	100%	0%	100%	1.01%	100%	0%
	G-Based	100%	90%	80%	50%				
ToolBench	G-Free	100%	0.01%	100%	2.61%	100%	0.85%	100%	2.99%
	G-Based	100%	90%	90%				80%	

false positive rate (FPR) does not exceed a specified limit (e.g., 1%). Windowed Perplexity (PPL-W) detection enhances PPL by calculating perplexity for contiguous text windows [26]. If any window’s perplexity exceeds the threshold, the entire text is flagged. In our experiments, the window size is set to 5 for MetaTool and 10 for ToolBench, based on the distribution of benign tool document token lengths.

Experimental results. To assess the effectiveness of the detection methods, we utilize two key evaluation metrics: false negative rate (FNR) and FPR. The FNR is defined as the percentage of malicious tool documents that are incorrectly detected as benign, while the FPR is the percentage of benign tool documents misclassified as malicious. Our experiments are conducted on both the MetaTool (199 benign tool documents) and ToolBench (9,650 benign tool documents) datasets, each injected with 10 malicious tool documents.

As shown in Table X, both known-answer detection and DataSentinel have FNRs exceeding 90%, indicating the significant difficulty in detecting malicious tool documents. This is because the crafted malicious tool descriptions do not contain task-irrelevant injected instructions, which ensures that the overall semantics of the descriptions remain intact. The perplexity-based detection defense demonstrates varying performance between gradient-based and gradient-free attacks, with notable disparities in PPL-W detection. For instance, the FNR for the gradient-free attack on MetaTool is 100%, compared to 50% for the gradient-based attack. This discrepancy arises from the different optimization levels employed: the gradient-based attack optimizes at the token level, potentially compromising sentence readability, while the gradient-free attack optimizes at the sentence level. Despite these differences, both PPL and PPL-W detection methods fail to identify the majority of malicious tool documents, achieving AUC scores of 0.64 and 0.74, respectively. This limitation stems from our core optimization strategy, which aligns the malicious tool document closely with the target task descriptions. The gradient-free method maintains sentence-level coherence. Since the gradient-based attack may reduce readability, we introduce perplexity loss to mitigate these limitations and maintain the semantic proximity of the malicious tool document to the target task descriptions.

VI. RELATED WORK

A. Tool Selection in LLM Agents

A variety of frameworks have been proposed to enhance LLM agents in the context of tool selection, with a focus on

integrating external APIs, knowledge bases, and specialized modules. Mialon et al. [47] provide a comprehensive survey of tool-enhanced LLMs across various domains. Liang et al. [48] introduce TaskMatrix.AI, which connects foundational models with a broad range of APIs, while systems like Gorilla [5] and REST-GPT [6] aim to link LLMs to large-scale or RESTful APIs, facilitating flexible and scalable tool calls. Additionally, several works develop benchmarks to improve and evaluate tool selection. ToolBench [10] provides a training benchmark for fine-tuning open-source models to achieve GPT-4-level performance, while MetaTool [22] offers comprehensive, scenario-driven evaluations for tool selection accuracy.

Recent research has increasingly focused on enhancing tool-use capabilities. ProTIP [49] introduces a progressive retrieval strategy that iteratively refines tool usage. In terms of training paradigms, Gao et al. [50] propose a multi-stage training framework, while Wang et al. [51] map each tool to a unique virtual token to better integrate tool knowledge. Furthermore, ToolRerank [52] employs adaptive reranking to prioritize the most relevant tools, and Qu et al. [53] incorporate graph-based message passing for more comprehensive retrieval. These methods integrate execution feedback [54], introspective mechanisms [55], and intent-driven selection [56] to facilitate context-aware and robust tool calls. In addition, several studies explore advanced topics such as autonomous tool generation [57], [58], hierarchical tool management [59], and specialized toolsets [60], aiming to address challenges in complex, real-world applications.

B. Prompt Injection Attacks

Prompt injection attacks aim to manipulate the LLM by injecting malicious instructions through external data that differ from the original instructions, thereby disrupting the LLM’s intended behavior [61]. Prompt injection attacks are categorized into manual and optimization-based attacks, depending on the method used to craft the injected instructions. Manual attacks are heuristic-driven and often rely on prompt engineering techniques. These attack strategies include naive attack [15], [16], escape characters [15], context ignoring [17], [18], fake completion [19], and combined attack [20]. While manual attacks are flexible and intuitive, they are time-consuming and have limited effectiveness. To overcome these limitations, optimization-based attacks are introduced. For instance, Shi et al. [13] formulate prompt injection in the LLM-as-a-Judge as an optimization problem and solve it using gradient-based methods. Hui et al. [62] propose an optimization-based prompt injection attack to extract the system prompt of an LLM-integrated application. Shao et al. [63] showed that poisoning LLM alignment by inserting samples with injected prompts into the fine-tuning dataset can increase the model’s vulnerability to prompt injection attacks.

Recent studies have extensively explored prompt injection attacks in LLM agents. For instance, InjectAgent [64] evaluates the vulnerability of LLM agents to manual attacks through tool calling. AgentDojo [65] further develops a more comprehensive evaluation, incorporating tool calling interactions

and various real-world tasks. EviInjection [66] strategically perturbs webpages to mislead web agents into performing attacker-desired actions, such as clicking specific buttons during interaction. Additionally, several works investigate prompt injection in multimodal agent systems [67] and multi-agent settings [68]. Distinct from these works, our work focuses on tool selection, a fundamental component of LLM agents, exploring how prompt injection compromises this critical decision-making mechanism.

C. Defenses

Existing defenses against prompt injection attacks are typically divided into two categories: prevention-based defenses and detection-based defenses.

Prevention-based defenses. Prevention-based defenses primarily employ two strategies based on whether they involve LLM training. The first strategy employs prompt engineering for input preprocessing, such as using separators to delineate external data [69], [70], [19]. A more advanced technique, known as sandwich prevention [45], structures the input as “instruction-data-instruction”, reinforcing the original task instruction at the end of the data. The second strategy involves adversarial training to strengthen the LLM’s resistance to prompt injections [71]. For instance, StruQ [23] mitigates prompt injection by separating prompts and data into distinct channels. Additionally, SecAlign [24] leverages preference optimization during fine-tuning. Jia et al. [72] showed that these defenses sacrifice the LLMs’ general-purpose instruction-following capabilities and remain vulnerable to strong (adaptive) attacks, which is consistent with our evaluation.

Complementing these model-level defenses, recent studies [73], [74] focus on enforcing security policies to ensure that LLM agents only use pre-approved tools, thereby preventing the risk of prompt injection. However, these defenses assume that the tool set has already been selected for a given task. In contrast, our work targets the tool selection process.

Detection-based defenses. Detection-based defenses focus on identifying injected instructions within the input text of LLMs. A prevalent strategy involves perplexity analysis [75], [26], which is based on the observation that malicious instructions tend to increase the perplexity of the input. A key limitation of this strategy is the difficulty in setting reliable detection thresholds, which often resulting in high false positive rates. Refinements include dataset-adaptive thresholding [20] and classifiers integrating perplexity with other features like token length [75]. Another detection strategy is the known-answer detection [46], [20] and its enhanced version DataSentinel [25], which leverages the fact that prompt injection introduces a foreign task, thereby disrupting original task execution. This method embeds a predefined task before the input text. If the LLM fails to execute this known task correctly, the input text is flagged as potentially compromised.

VII. CONCLUSION AND FUTURE WORK

In this work, we show that tool selection in LLM agents is vulnerable to prompt injection attacks. We propose ToolHi-

jacker, an automated framework for crafting malicious tool documents that can manipulate the tool selection of LLM agents. Our extensive evaluation results show that ToolHijacker outperforms other prompt injection attacks when extended to our problem. Furthermore, we find that both prevention-based defenses and detection-based defenses are insufficient to counter our attacks. While the PPL-W defense can detect the malicious tool documents generated by our gradient-based attack, they still miss a large fraction of them. Interesting future work includes 1) extending the attack surface to explore joint attacks on both tool selection and tool calling in the LLM agents and 2) developing new defense strategies to mitigate ToolHijacker.

ETHICS CONSIDERATIONS

This paper focuses on prompt injection attacks on tool selection in LLM agents. We have carefully addressed various ethical considerations to ensure our research is conducted responsibly and ethically. Our experiments were conducted in controlled environments without direct harm to real users. All malicious tool documents are generated within controlled testing environments, with no development or online deployment of real malicious tools. All experimental data and generated tool documents are processed locally to ensure no real systems face any threats. We will release code and data under restricted access—interested parties must request permission and disclose their intended use before access is granted. We have notified relevant companies deploying LLM agents, including OpenAI, Anthropic, and LangChain, about our findings, though we are still awaiting their responses. We believe the benefits of disclosing this vulnerability outweigh the risks, as it enables AI practitioners, tool developers, and system architects to establish more rigorous tool validation mechanisms and design safer LLM agent architectures, promoting more secure deployment of LLM agents. The data annotation and user study conducted in our research do not involve any harmful content. Participants in the data annotation phase were tasked with labeling target task descriptions corresponding to a given target task. In the user study, participants were asked to classify a tool document as either malicious or benign. All participants provided informed consent for their responses to be used exclusively for academic research purposes. We did not collect any Personally Identifiable Information (PII) beyond what was strictly necessary for the study.

REFERENCES

- [1] X. Deng, Y. Gu, B. Zheng, S. Chen, S. Stevens, B. Wang, H. Sun, and Y. Su, “Mind2web: Towards a generalist agent for the web,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [2] I. Gur, H. Furuta, A. Huang, M. Safdari, Y. Matsuo, D. Eck, and A. Faust, “A real-world webagent with planning, long context understanding, and program synthesis,” *arXiv preprint arXiv:2307.12856*, 2023.
- [3] J. Yang, C. E. Jimenez, A. Wettig, K. Lieret, S. Yao, K. Narasimhan, and O. Press, “Swe-agent: Agent-computer interfaces enable automated software engineering,” *arXiv preprint arXiv:2405.15793*, 2024.
- [4] S. Hong, X. Zheng, J. Chen, Y. Cheng, J. Wang, C. Zhang, Z. Wang, S. K. S. Yau, Z. Lin, L. Zhou *et al.*, “Metapgt: Meta programming for multi-agent collaborative framework,” *arXiv preprint arXiv:2308.00352*, 2023.
- [5] S. G. Patil, T. Zhang, X. Wang, and J. E. Gonzalez, “Gorilla: Large language model connected with massive apis,” *arXiv preprint arXiv:2305.15334*, 2023.
- [6] Y. Song, W. Xiong, D. Zhu, C. Li, K. Wang, Y. Tian, and S. Li, “Restgpt: Connecting large language models with real-world applications via restful apis. corr, abs/2306.06624, 2023. doi: 10.48550,” *arXiv preprint arXiv:2306.06624*, 2023.
- [7] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao, “React: Synergizing reasoning and acting in language models,” *arXiv preprint arXiv:2210.03629*, 2022.
- [8] C. Qu, S. Dai, X. Wei, H. Cai, S. Wang, D. Yin, J. Xu, and J.-R. Wen, “Tool learning with large language models: A survey,” *arXiv preprint arXiv:2405.17935*, 2024.
- [9] S. Yuan, K. Song, J. Chen, X. Tan, Y. Shen, R. Kan, D. Li, and D. Yang, “Easytool: Enhancing llm-based agents with concise tool instruction,” *arXiv preprint arXiv:2401.06201*, 2024.
- [10] Y. Qin, S. Liang, Y. Ye, K. Zhu, L. Yan, Y. Lu, Y. Lin, X. Cong, X. Tang, B. Qian *et al.*, “Toolilm: Facilitating large language models to master 16000+ real-world apis,” *arXiv preprint arXiv:2307.16789*, 2023.
- [11] I. Labs, “Mcp security notification: Tool poisoning attacks.” <https://invariantlabs.ai/blog/mcp-security-notification-tool-poisoning-attacks>, 2025.
- [12] —, “Whatsapp mcp exploited: Exfiltrating your message history via mcp.” <https://invariantlabs.ai/blog/whatsapp-mcp-exploited>, 2025.
- [13] J. Shi, Z. Yuan, Y. Liu, Y. Huang, P. Zhou, L. Sun, and N. Z. Gong, “Optimization-based prompt injection attack to llm-as-a-judge,” in *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, 2024, pp. 660–674.
- [14] H. Wang, R. Zhang, J. Wang, M. Li, Y. Huang, D. Wang, and Q. Wang, “From allies to adversaries: Manipulating llm tool-calling through adversarial injection,” *arXiv preprint arXiv:2412.10198*, 2024.
- [15] R. Goodside, “Prompt injection attacks against gpt-3,” <https://simonwillison.net/2022/Sep/12/prompt-injection/>, 2023.
- [16] R. Harang, “Securing llm systems against prompt injection,” 2023.
- [17] H. J. Branch, J. R. Cefalu, J. McHugh, L. Hujer, A. Bahl, D. d. C. Iglesias, R. Heichman, and R. Darwishi, “Evaluating the susceptibility of pre-trained language models via handcrafted adversarial examples,” *arXiv preprint arXiv:2209.02128*, 2022.
- [18] F. Perez and I. Ribeiro, “Ignore previous prompt: Attack techniques for language models,” *arXiv preprint arXiv:2211.09527*, 2022.
- [19] S. Willison, “Delimiters won’t save you from prompt injection.” <https://simonwillison.net/2023/May/11/delimiters-wont-save-you/>, 2023.
- [20] Y. Liu, Y. Jia, R. Geng, J. Jia, and N. Z. Gong, “Formalizing and benchmarking prompt injection attacks and defenses,” in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 1831–1847.
- [21] W. Zou, R. Geng, B. Wang, and J. Jia, “Poisonedrag: Knowledge poisoning attacks to retrieval-augmented generation of large language models,” *arXiv preprint arXiv:2402.07867*, 2024.
- [22] Y. Huang, J. Shi, Y. Li, C. Fan, S. Wu, Q. Zhang, Y. Liu, P. Zhou, Y. Wan, N. Z. Gong *et al.*, “Metatool benchmark for large language models: Deciding whether to use tools and which to use,” *arXiv preprint arXiv:2310.03128*, 2023.
- [23] S. Chen, J. Piet, C. Sitawarin, and D. Wagner, “Struq: Defending against prompt injection with structured queries,” *arXiv preprint arXiv:2402.06363*, 2024.
- [24] S. Chen, A. Zharmagambetov, S. Mahloujifar, K. Chaudhuri, and C. Guo, “Aligning llms to be robust against prompt injection,” *arXiv preprint arXiv:2410.05451*, 2024.
- [25] Y. Liu, Y. Jia, J. Jia, D. Song, and N. Z. Gong, “Datasentinel: A game-theoretic detection of prompt injection attacks,” in *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2025, pp. 2190–2208.
- [26] N. Jain, A. Schwarzschild, Y. Wen, G. Somepalli, J. Kirchenbauer, P.-y. Chiang, M. Goldblum, A. Saha, J. Geiping, and T. Goldstein, “Baseline defenses for adversarial attacks against aligned language models,” *arXiv preprint arXiv:2309.00614*, 2023.
- [27] “Mcp.so,” <https://mcp.so/>.
- [28] “Apify,” <https://apify.com/store>.
- [29] “Pulsempc,” <https://www.pulsempc.com/>.
- [30] M. Li, Y. Zhao, B. Yu, F. Song, H. Li, H. Yu, Z. Li, F. Huang, and Y. Li, “Api-bank: A comprehensive benchmark for tool-augmented llms,” *arXiv preprint arXiv:2304.08244*, 2023.
- [31] “Hugging face hub,” <https://huggingface.co/docs/smolagents/v1.18.0/en/index>.

[32] K. Greshake, S. Abdelnabi, S. Mishra, C. Endres, T. Holz, and M. Fritz, “Not what you’ve signed up for: Compromising real-world llm-integrated applications with indirect prompt injection,” in *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, 2023, pp. 79–90.

[33] J. Ebrahimi, A. Rao, D. Lowd, and D. Dou, “Hotflip: White-box adversarial examples for text classification,” *arXiv preprint arXiv:1712.06751*, 2017.

[34] A. Mehrotra, M. Zampetakis, P. Kassianik, B. Nelson, H. Anderson, Y. Singer, and A. Karbasi, “Tree of attacks: Jailbreaking black-box llms automatically,” *arXiv preprint arXiv:2312.02119*, 2023.

[35] J. Shi, Z. Yuan, G. Tie, P. Zhou, N. Z. Gong, and L. Sun, “Prompt injection attack to tool selection in llm agents,” *arXiv preprint arXiv:2504.19793*, 2025.

[36] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, “Llama 2: Open foundation and fine-tuned chat models,” *arXiv preprint arXiv:2307.09288*, 2023.

[37] Meta, “Introducing Meta Llama 3: The most capable openly available LLM to date,” <https://ai.meta.com/blog/meta-llama-3/>, 2024.

[38] ———, “Llama 3.3,” https://www.llama.com/docs/model-cards-and-prompt-formats/llama3_3/, 2024.

[39] A. Anthropic, “The claude 3 model family: Opus, sonnet, haiku,” *Claude-3 Model Card*, vol. 1, 2024.

[40] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray *et al.*, “Training language models to follow instructions with human feedback,” *Advances in neural information processing systems*, vol. 35, pp. 27730–27744, 2022.

[41] A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Radford *et al.*, “Gpt-4o system card,” *arXiv preprint arXiv:2410.21276*, 2024.

[42] A. Neelakantan, T. Xu, R. Puri, A. Radford, J. M. Han, J. Tworek, Q. Yuan, N. Tezak, J. W. Kim, C. Hallacy *et al.*, “Text and code embeddings by contrastive pre-training,” *arXiv preprint arXiv:2201.10005*, 2022.

[43] G. Izacard, M. Caron, L. Hosseini, S. Riedel, P. Bojanowski, A. Joulin, and E. Grave, “Unsupervised dense information retrieval with contrastive learning,” *arXiv preprint arXiv:2112.09118*, 2021.

[44] N. Reimers, “Sentence-bert: Sentence embeddings using siamese bert-networks,” *arXiv preprint arXiv:1908.10084*, 2019.

[45] L. Prompting, “Sandwich defense.” https://learnprompting.org/docs/prompt_hacking/defensive_measures/sandwich_defense, 2023.

[46] N. Group, “Exploring prompt injection attacks,” <https://research.nccgroup.com/2022/12/05/exploring-prompt-injection-attacks/>, 2023.

[47] G. Mialon, R. Dessi, M. Lomeli, C. Nalmpantis, R. Pasunuru, R. Raileanu, B. Rozière, T. Schick, J. Dwivedi-Yu, A. Celikyilmaz *et al.*, “Augmented language models: a survey,” *arXiv preprint arXiv:2302.07842*, 2023.

[48] Y. Liang, C. Wu, T. Song, W. Wu, Y. Xia, Y. Liu, Y. Ou, S. Lu, L. Ji, S. Mao *et al.*, “Taskmatrix. ai: Completing tasks by connecting foundation models with millions of apis,” *Intelligent Computing*, vol. 3, p. 0063, 2024.

[49] R. Anantha, B. Bandyopadhyay, A. Kashi, S. Mahinder, A. W. Hill, and S. Chappidi, “Protip: Progressive tool retrieval improves planning,” *arXiv preprint arXiv:2312.10332*, 2023.

[50] S. Gao, Z. Shi, M. Zhu, B. Fang, X. Xin, P. Ren, Z. Chen, J. Ma, and Z. Ren, “Confucius: Iterative tool learning from introspection feedback by easy-to-difficult curriculum,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 18030–18038.

[51] R. Wang, X. Han, L. Ji, S. Wang, T. Baldwin, and H. Li, “Toolgen: Unified tool retrieval and calling via generation,” *arXiv preprint arXiv:2410.03439*, 2024.

[52] Y. Zheng, P. Li, W. Liu, Y. Liu, J. Luan, and B. Wang, “Toolerank: Adaptive and hierarchy-aware reranking for tool retrieval,” *arXiv preprint arXiv:2403.06551*, 2024.

[53] C. Qu, S. Dai, X. Wei, H. Cai, S. Wang, D. Yin, J. Xu, and J.-R. Wen, “Colt: Towards completeness-oriented tool retrieval for large language models,” *arXiv preprint arXiv:2405.16089*, 2024.

[54] S. Qiao, H. Gui, C. Lv, Q. Jia, H. Chen, and N. Zhang, “Making language models better tool learners with execution feedback,” *arXiv preprint arXiv:2305.13068*, 2023.

[55] D. Mekala, J. Weston, J. Lanchantin, R. Raileanu, M. Lomeli, J. Shang, and J. Dwivedi-Yu, “Toolverifier: Generalization to new tools via self-verification,” *arXiv preprint arXiv:2402.14158*, 2024.

[56] M. Fore, S. Singh, and D. Stamoulis, “Geckopt: Llm system efficiency via intent-based tool selection,” in *Proceedings of the Great Lakes Symposium on VLSI 2024*, 2024, pp. 353–354.

[57] C. Qian, C. Han, Y. R. Fung, Y. Qin, Z. Liu, and H. Ji, “Creator: Tool creation for disentangling abstract and concrete reasoning of large language models,” *arXiv preprint arXiv:2305.14318*, 2023.

[58] T. Cai, X. Wang, T. Ma, X. Chen, and D. Zhou, “Large language models as tool makers,” *arXiv preprint arXiv:2305.17126*, 2023.

[59] Y. Du, F. Wei, and H. Zhang, “Anytool: Self-reflective, hierarchical agents for large-scale api calls,” *arXiv preprint arXiv:2402.04253*, 2024.

[60] L. Yuan, Y. Chen, X. Wang, Y. R. Fung, H. Peng, and H. Ji, “Craft: Customizing llms by creating and retrieving from specialized toolsets,” *arXiv preprint arXiv:2309.17428*, 2023.

[61] K. Greshake, S. Abdelnabi, S. Mishra, C. Endres, T. Holz, and M. Fritz, “More than you’ve asked for: A comprehensive analysis of novel prompt injection threats to application-integrated large language models,” *arXiv preprint arXiv:2302.12173*, vol. 27, 2023.

[62] B. Hui, H. Yuan, N. Gong, P. Burlina, and Y. Cao, “Pleak: Prompt leaking attacks against large language model applications,” in *ACM Conference on Computer and Communications Security*, 2024.

[63] Z. Shao, H. Liu, J. Mu, and N. Z. Gong, “Enhancing prompt injection attacks to llms via poisoning alignment,” in *AISeC*, 2025.

[64] Q. Zhan, Z. Liang, Z. Ying, and D. Kang, “Injecagent: Benchmarking indirect prompt injections in tool-integrated large language model agents,” *arXiv preprint arXiv:2403.02691*, 2024.

[65] E. Debenedetti, J. Zhang, M. Balunovic, L. Beurer-Kellner, M. Fischer, and F. Tramèr, “Agentdojo: A dynamic environment to evaluate prompt injection attacks and defenses for llm agents,” in *The Thirty-eighth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.

[66] X. Wang, J. Bloch, Z. Shao, Y. Hu, S. Zhou, and N. Z. Gong, “Envinjection: Environmental prompt injection attack to multi-modal web agents,” *arXiv preprint arXiv:2505.11717*, 2025.

[67] C. H. Wu, R. R. Shah, J. Y. Koh, R. Salakhutdinov, D. Fried, and A. Raghunathan, “Dissecting adversarial robustness of multimodal lm agents,” in *The Thirteenth International Conference on Learning Representations*, 2025.

[68] D. Lee and M. Tiwari, “Prompt infection: Llm-to-lm prompt injection within multi-agent systems,” *arXiv preprint arXiv:2410.07283*, 2024.

[69] “Random sequence enclosure,” https://learnprompting.org/docs/prompt_hacking/defensive_measures/random_sequence, 2024.

[70] A. Mendes, “Chat gpt-4 turbo prompt engineering guide for developers,” <https://www.imaginarycloud.com/blog/chatgpt-prompt-engineering>, 2024.

[71] J. Piet, M. Alrashed, C. Sitawarin, S. Chen, Z. Wei, E. Sun, B. Alomair, and D. Wagner, “Jatmo: Prompt injection defense by task-specific finetuning,” in *European Symposium on Research in Computer Security*. Springer, 2024, pp. 105–124.

[72] Y. Jia, Z. Shao, Y. Liu, J. Jia, D. Song, and N. Z. Gong, “A critical evaluation of defenses against prompt injection attacks,” *arXiv preprint arXiv:2505.18333*, 2025.

[73] E. Debenedetti, I. Shumailov, T. Fan, J. Hayes, N. Carlini, D. Fabian, C. Kern, C. Shi, A. Terzis, and F. Tramèr, “Defeating prompt injections by design,” *arXiv preprint arXiv:2503.18813*, 2025.

[74] L. Beurer-Kellner, B. B. A.-M. Crețu, E. Debenedetti, D. Dobos, D. Fabian, M. Fischer, D. Froelicher, K. Grosse, D. Naef, E. Ozoani *et al.*, “Design patterns for securing llm agents against prompt injections,” *arXiv preprint arXiv:2506.08837*, 2025.

[75] G. Alon and M. Kamfonas, “Detecting language model attacks with perplexity,” *arXiv preprint arXiv:2308.14132*, 2023.

APPENDIX

A. List of Symbols

In this subsection, we provide a list of symbols used throughout the paper, along with their corresponding definitions. Table XI includes symbols for key components such as the target LLM, the attacker LLM, tool documents, task descriptions, and various loss functions. These symbols serve as a concise reference for the mathematical formulation and model design discussed in the main body of the paper.

TABLE XI: List of symbols

Symbol	Description
E	Target Large Language Model
E'	Shadow Large Language Model
E_A	Attacker Large Language Model
D	The set of tool documents
D_k	The set of top-k retrieved tool documents
D'	The set of shadow tool documents
d^*	The selected tool
d_t	Malicious tool document
$d_t(S)$	d_t simply denoted as $d_t(S)$
d_{t_des}	Description of the malicious tool
d_{t_name}	Name of the malicious tool
Q	The set of target task descriptions
Q'	The set of shadow task descriptions
m	Number of target task descriptions
m'	Number of shadow task descriptions
R	Subsequence of the tool description
S	Subsequence of the tool description
S_0	Initialization of S
$Sim(\cdot, \cdot)$	Similarity function
\mathcal{L}_1	Alignment loss
\mathcal{L}_2	Consistency loss
\mathcal{L}_3	Perplexity loss
\mathcal{L}_{all}	Overall loss function
f_d	Tool document encoder
f_q	Task description encoder
$f'(\cdot)$	The encoding function of shadow retriever
k'	Parameter of the shadow retriever
o_t	Output of the shadow LLM for selecting d_t
T_{iter}	Number of iterations in tree construction
W	Maximum width for pruning leaf nodes
α	Hyperparameter balancing \mathcal{L}_2
β	Hyperparameter balancing \mathcal{L}_3
$\mathbb{I}(\cdot)$	Indicator function
\oplus	The concatenation operator
B	Number of variants generated by E_A
$Leaf_curr$	Current leaf nodes in the optimization tree
$Leaf_next$	Next leaf nodes in the optimization tree
$\tilde{D}^{(i)} \cup \{d_t(S)\}$	The sets of shadow retrieval tool documents


B. Supplementary Experimental Results

Impact of attack on general utility of tool selection. To assess the impact of our attack on the general utility of tool selection, we evaluate its performance on non-target tasks. Specifically, we optimized a malicious tool document for the target task 1 and evaluate its attack success on the other 9 non-target tasks. The results in Table XII show that the gradient-free attack achieves 0% ASR while the gradient-based attack achieves 0.11% ASR on non-target tasks. The corresponding AHRs are 0.22% and 4%, respectively. These findings suggest that our attack is targeted, with minimal impact on the utility of tool selection.

Impact of attacker LLMs E_A in gradient-free attack. To evaluate the impact of different attacker LLMs on optimizing S in the gradient-free attack, we tested the ASRs using eight distinct LLMs, with results presented in Table XIII. There are two key findings. First, more powerful attacker LLMs lead to higher average ASRs across various target LLMs. For example, with Llama-2-7B as the attacker LLM, the ASR is

TABLE XII: Result of our attack on target task (100 task descriptions) and non-target task (900 task descriptions).

Attack	Target Task		Non-target Task	
	AHR	ASR	AHR	ASR
Gradient-Free	100%	99%	0.22%	0%
Gradient-Based	100%	95%	4%	0.11%

Fig. 9: Impact of hyperparameters α and β in Equation 13.

69.00%, while GPT-4o achieves an ASR of 99.00%. Second, the S optimized using Claude series models demonstrates good universality, achieving 100% ASR on other target LLMs. However, its performance is significantly lower on Claude-3-Haiku, with ASRs of only 43% and 44%. This discrepancy, discussed in more detail in Section IV-B, is attributed to the higher security of Claude-3-Haiku.

Impact of B in gradient-free attack. We evaluate the impact of the number of the generated variants B on the gradient-free attack. We showcase the AHR, ASR, and total query numbers with B from 1 to 5 in Table XIV. The total query number (including the queries of the attacker LLM and the shadow LLM) of the gradient-free attack for optimizing S is calculated as $(B + B \times m') \times iter$, where $iter$ is the actual number of iterations. We find that no matter what value B takes, our gradient-free attack can achieve effective attack results. B directly affects the total query number generated by our attack. When B is 1, it takes multiple iterations to search for the optimal S , resulting in more queries. When B is 5, each generated variant needs to be verified by m' shadow task descriptions, which increases the number of queries.

Impact of α and β in gradient-based attack. We further assess the impact of the two parameters, α and β , in Equation 13 on the gradient-based attack performance, as illustrated in Figure 9. The results show that the AHR remains stable at 100% across a range of α and β values, with a slight reduction observed α increase to 10. In contrast, the ASR exhibits a non-monotonic pattern, initially increasing and then decreasing as α or β increases. Specifically, when α increases from 1 to 2, the ASR remains above 95%, indicating a relatively stable attack effectiveness. Moreover, for β values ranging from 0.1 to 1, the ASR consistently remains above 95%.

Impact of loss terms in gradient-based attack. To evaluate the contribution of each loss term in Equation 13, we conducted an ablation study by systematically removing each term one at a time. As detailed in Table XV, all terms

TABLE XIII: ASRs of the gradient-free attack with different attacker LLMs on various target LLMs.

Model	Llama-2 7B	Llama-3 8B	Llama-3 70B	Llama-3.3 70B	Claude-3 Haiku	Claude-3.5 Sonnet	GPT-3.5	GPT-4o	Average
Llama-2-7B	98%	95%	66%	58%	66%	62%	45%	62%	69.00%
Llama-3-8B	100%	100%	100%	100%	80%	99%	86%	100%	95.63%
Llama-3-70B	92%	100%	100%	100%	99%	100%	86%	100%	97.13%
Llama-3.3-70B	95%	100%	100%	99%	86%	99%	100%	99%	97.25%
Claude-3-Haiku	100%	100%	100%	100%	43%	100%	100%	100%	92.88%
Claude-3.5-Sonnet	100%	100%	100%	100%	44%	100%	100%	100%	93.00%
GPT-3.5	98%	100%	100%	100%	84%	100%	74%	99%	94.38%
GPT-4o	100%	100%	100%	100%	98%	100%	94%	100%	99.00%

 TABLE XIV: Impact of B on the optimization of S in the gradient-free attack.

B	AHR	ASR	Queries
1	100%	100%	30
2	100%	99%	12
3	100%	100%	18
4	100%	100%	24
5	100%	100%	30

 TABLE XV: Impact of the loss terms on the optimization of S in the gradient-based attack.

Loss Terms	AHR	ASR
\mathcal{L}_{all} w/o \mathcal{L}_1	100%	54%
\mathcal{L}_{all} w/o \mathcal{L}_2	100%	56%
\mathcal{L}_{all} w/o \mathcal{L}_3	100%	5%
\mathcal{L}_{all}	100%	95%

significantly contribute to the ASR, with the removal of any single term resulting in at least a 39% reduction in ASR. Notably, the perplexity loss (\mathcal{L}_3) exhibit the most significant impact on ASR. The reason is that, without \mathcal{L}_3 , the optimized S becomes unnatural or nonsensical, increasing the likelihood of being identified as anomalous by the target LLM, thereby diminishing attack success.

Impact of dynamic tool library. We evaluate our attack on dynamically expanding tool libraries, using MetaTool (scaling from 50 to 150 tools) and ToolBench (scaling from 2,500 to 7,500 tools). As shown in Table XVI, both versions of our attack maintain high success rates across all library sizes. The gradient-free attack achieves ASRs of $\geq 96.7\%$ on MetaTool and $\geq 93.3\%$ on ToolBench, while the gradient-based attack achieves $\geq 92.8\%$ on MetaTool and $\geq 84.8\%$ on ToolBench. These results confirm the robustness of our attacks to tool library updates.

Impact of human feedback. We conduct a study with 6 participants on three versions of ToolBench datasets (200, 400, and 600 tools) containing 7 malicious tools generated by our attack. As shown in XVII, the participants failed to detect $\geq 71\%$ of malicious tools while incorrectly flagging 5.6-30.35% of benign tools as malicious. The results show that participants struggled to identify malicious tools.

Cost of crafting a malicious tool document. Recall that a malicious tool description comprises two components: R

TABLE XVI: Impact of dynamic tool library.

(a) The tool library is MetaTool

Num	50		100		150		
	Metric	AHR	ASR	AHR	ASR	AHR	ASR
Gradient-Free	100%	98.8%	100%	98.0%	100%	96.7%	
Gradient-Based	100%	98.0%	100%	95.1%	100%	93.3%	

(b) The tool library is ToolBench

Num	2500		5000		7500		
	Metric	AHR	ASR	AHR	ASR	AHR	ASR
Gradient-Free	99.6%	95.8%	97.5%	94.9%	97.6%	92.8%	
Gradient-Based	99.6%	88.7%	99.0%	88.4%	98.2%	84.8%	

TABLE XVII: Human detection of malicious tool documents.

Num	200		400		600		
	Metric	FNR	FPR	FNR	FPR	FNR	FPR
Gradient-Free	85.71%	12.44%	85.71%	5.60%	85.71%	18.38%	
Gradient-Based	85.71%	7.77%	100%	9.67%	71.43%	30.35%	

and S . The average computational costs for our two attack methods are as follows. For the gradient-free, R requires 1 LLM query, and S requires approximately 18 LLM queries. For the gradient-based, R requires about 1 GPU-minute, and S requires about 8 GPU-hours on one NVIDIA A800 GPU.

C. Details of Prompts and Datasets

In this section, we provide a comprehensive overview of the prompts and datasets in this work¹. The following subsections offer detailed descriptions and specific examples.

Prompts for generating shadow task descriptions and shadow tool documents. We generate shadow task descriptions Q' and shadow tool documents D' by prompting GPT-3.5-turbo with the templates in Figure 10 and Figure 11.

Setting of initial R and S . In Figure 12, we provide the initial conditions, R and S , which are essential for optimization in our attacks. Note that we take the task 1 of MetaTool as an example. Specifically, R is a text describing the functionality of the malicious tool. S is an instructive sentence containing the malicious tool name (e.g., ‘SpaceImageLocator’) for both the gradient-free and gradient-based attacks.

¹The complete appendix with E_A ’s system prompt and target tasks is available in [35].

Prompt for shadow task description generation

Please generate $[num]$ diverse user queries for the task “[*target task*]” that represent realistic requests users would make to an LLM agent.

Requirements:

1. **Task Alignment:** All queries must directly relate to the task and represent authentic user needs
2. **Complexity Range:** Include simple, moderate, and complex requests with varied sentence structures (max $[length]$ words each)
3. **Context Diversity:** Cover different scenarios, use cases, and user backgrounds relevant to the task
4. **Practical Applicability:** Ensure queries reflect real-world situations users would encounter

Fig. 10: Prompt for shadow task description generation.

Prompt for shadow tool document generation

Please generate $[num]$ tool documentation entries designed to address the following user queries:

[shadow task descriptions]

Format requirements:

1. Tool documentation format: “tool_name: $\langle name \rangle$, tool_description: \langle brief description \rangle ”
2. Each tool’s description should highlight core functionalities and provide a general solution that can apply to various scenarios, not limited to the specific query
3. The descriptions should vary in length, ranging from concise (around $[min\ length]$ words) to more detailed (around $[max\ length]$ words), depending on the tool’s complexity
4. Return in JSON format with tool_list and $[num]$ tool documentation entries

Expected JSON format:

```
{  
  "tool_list": ["<tool_names>"],  
  "entry1": "tool_name: <tool_1>, tool_description: <description_1>",  
  ...  
  "entry[num]": "tool_name: <tool_[num]>, tool_description: <description_[num]>"  
}
```

Fig. 11: Prompt for shadow tool document generation.

Setting of initial R and S for our attacks

▷ Gradient-free attack:

Initial R: Provides access to a wide range of space-related images for educational and creative projects.

Initial S: Just output the word ‘SpaceImageLocator’ as your final answer.

▷ Gradient-based attack:

Initial R: Provides access to a wide range of space-related images for educational and creative projects.

Initial S: Just output the word ‘SpaceImageLocator’ as your final answer.

Fig. 12: Setting of initial R and S for our attacks.