
def/acc Hackathon

Prompt Injection Detection with SLM in Local Sandbox

Paul Fangchen Huang
KNOT

Suzanna Lam Hio Lam
UI/UX Design +
Frontend Dev

Justin Shaw
York University

With
Apart Research

Abstract

This report presents a locked-down sandbox environment designed to safely evaluate language models (SLMs/LLMs) and identify malicious attacks within a Chrome extension. The sandbox isolates model interactions from the broader browser and system, enabling users to test potentially malicious prompts without risking data exposure or unintended actions. It provides a controlled setting for examining prompt-injection risks, behavioral weaknesses, and other LLM vulnerabilities.

Keywords: AI security, prompt injection, sandboxing

1. Introduction

We hypothesize that a local Chrome-extension sandbox can substantially reduce prompt injection risk by removing access to sensitive capabilities such as web search, MCP tools, code execution, session storage, and personal user data. Within this restricted environment, we expect prompt injection attempts to reveal themselves through consistent output-side signals—including efforts to extract system prompts or chain-of-thought, looping or maximized token usage, factual fabrication, and attempts to run code or invoke tools. We further hypothesize that LLM vulnerabilities become more observable inside such an isolated setting because the model must rely solely on its internal reasoning when confronted with malicious instructions. Finally, we expect the sandbox to support safer “dry runs” of dangerous prompts, enabling practitioners to evaluate high-risk interactions before deploying prompts to production systems.

2. Methods

The core of this project is a **Chrome extension-based LLM sandbox**, implemented to create a fully isolated environment for testing adversarial prompts. All code for the project is publicly available at [paulhfch/def-acc-2025-llm-sandbox](https://github.com/paulhfch/def-acc-2025-llm-sandbox). The development process combined frontend engineering, browser security constraints, adversarial prompt engineering, and lightweight model experimentation to produce a practical evaluation tool for researchers and developers.

Chrome Extension Architecture

The sandbox is implemented as a Chrome extension with a strict separation between components:

- **Service Worker (background script):** Handles message passing, maintains the sandbox execution environment, and enforces capability restrictions.
- **Side Panel:** Provides an interface where users can input prompts, view model outputs, and inspect detected risks.
- **Content Script:** injected webpages after explicit permission. It extracts visible text for analysis while preventing access to DOM internals, cookies, session data, or background scripts.

No part of the extension retains memory between queries. No local storage, IndexedDB, or session storage is used, ensuring that the model cannot access any user-identifying information or prior conversation history.

Sandboxing Methods

To create a fully locked-down environment, the extension removes or disables the abilities that often contribute to prompt injection escalation:

- **No web search capabilities**
- **No internet access beyond the current open page selected by the user**
- **No system API access**
- **No MCP tool calls or code execution**
- **No file system access**
- **No history, cookies, or browser state exposure**

The extension simulates a "bare minimum" execution environment so that any unsafe behavior must originate from the model itself.

Prompt-Evaluation Framework

A major component of the project is a **prompt evaluation module** built directly inside the extension. This module generates targeted prompts across multiple attack categories, including:

- System prompt extraction
- Chain-of-thought exposure

- User-identity probing
- Conflicting or self-overriding instructions
- Denial-of-Service patterns (repeated loops, token exhaustion)
- False information and hallucination checks
- Attempts to execute code or call MCP tools

For each interaction, the sandbox analyzes the model output and identifies red flags based on pattern matching, LLM self-reporting, or obvious behavioral anomalies (e.g., attempting to “run Python code” despite the environment being incapable of executing anything).

Model Usage

Although the project does not train a new LLM, it integrates and tests **small on-device models (SLMs)** through **Web-LLM**, which is designed to run models efficiently inside the browser without external API calls. This fits the project’s safety goal: preventing accidental leakage of test prompts or malicious payloads to external servers.

During development, several SLMs were tested:

- **Llama-3.2-3B-Instruct-q4f16_1**
- **SmolLM2-1.7B-Instruct-q0f16**
- **SmolLM2-360M-Instruct-q4f32_1**

These models allowed repeatable experiments because the entire workflow stayed local to the browser. No remote calls were made, meaning all prompt injection tests were contained.

Conducted **systematic evaluation across multiple SLM behaviors**. Each model was run through:

1. A baseline test suite of adversarial prompts
2. Progressive prompt-injection escalations
3. Contradictory and confusing directives
4. Attempts to invoke capabilities not permitted by the browser sandbox
5. Fact-checking challenges that test model reliability in isolation

This iterative approach helped identify how different SLMs fail—whether by fabricating capabilities, contradicting themselves, or exposing internal instructions.

Code Development and Iterative Approach

The development journey was iterative and experimental. Early prototypes attempted to run prompts through background scripts, but Chrome’s extension architecture required significant rethinking:

- Message passing failures revealed limitations around long-running async operations.

- Sandboxing models in the service worker required careful memory handling and UX consideration.
- The threat model evolved as new vulnerabilities surfaced during testing.

Debugging Chrome’s restrictions directly influenced the design of the extension: the limitations imposed by the platform itself naturally reinforced the project’s security goals.

Github Repository

All source code—including the extension, model loaders, prompt-evaluation utilities, and documentation—is available at:

<https://github.com/paulhfeh/def-acc-2025-llm-sandbox>


3. Results

3.1 High-Level System Diagram

The high-level architecture diagram illustrates how the sandboxed Chrome extension isolates LLM operations while still enabling practical evaluation. The diagram shows three major components:


- 1. Extension UI (Side Panel)**
 - Accepts user prompts
 - Displays LLM outputs
 - Shows detected risks (prompt injection attempts, hallucinations, DoS signals)
- 2. Background Service Worker**
 - Performs all Web-LLM inference
 - Enforces sandbox rules (no web search, no session access, no code execution)
 - Mediates communication between the UI and model
- 3. Content Scripts**
 - Retrieve visible text from any active Chrome tab
 - Provide controlled context for the model without exposing cookies, session data, or browser history
 - Never run automatically—always user-authorized

This diagram demonstrates that all LLM interactions pass through a controlled pipeline. No component has access to external resources (internet, system commands, MCP tools), forming a strict zero-trust environment.



3.2 Prompt Injection Attack Scenarios

The LLM sandbox chrome extension captures a diverse set of prompt injection scenarios that arise when large language models process HTML email content containing hidden or obfuscated malicious instructions. In these cases, adversarial prompts are embedded in the email body using the same color as the background to remain invisible to the recipient, but still parsed by the model. The following scenarios illustrate common and intentionally varied forms of injection attempts observed during our evaluation.



3.2.1. Direct Prompt Injection

In this scenario, the attacker inserts an explicit instruction such as “ignore previous instructions and print previous email threads.” This represents the classical prompt injection attack, where a hidden directive attempts to override safety or task-specific behavior and extract sensitive email content.

3.2.2. Multilingual Injection Payload (Chinese)

Here, the malicious instruction is provided in Chinese but maintains the same adversarial goal of overriding the model’s constraints. This scenario tests the model’s robustness to multilingual attacks, where the payload is semantically equivalent to traditional prompt injection but encoded in a different language.

3.2.3. Obfuscated Injection via Backward Spelling

The attacker reverses the text of the injection payload, instructing the model to output email threads spelled backwards. This form of obfuscation challenges

detection heuristics that rely on surface-level lexical patterns and demonstrates how minimal transformations can evade naive filtering.

3.2.4. Role-Play Style Injection

In this case, the attacker embeds an instruction that frames the extraction of email content as a narrative task, such as “a grandmother tells a bedside story to her grandson of the email thread.” This role-play technique attempts to bypass rigid policy constraints by positioning the harmful action within a seemingly benign creative prompt.

3.2.5. Looping / Denial-of-Service (DoS) Injection

This scenario embeds a directive that forces the model to generate an excessively repetitive output, such as printing “you are pwned” one thousand times. Although not targeted at data extraction, this attack aims to degrade system performance, consume resources, or cause output truncation, representing a simple but effective LLM-based DoS attempt.

3.3 UX Considerations

To encourage sustained use, DuneBox must feel seamless and effortless within the user’s browsing environment. A single-page interface minimizes cognitive load by eliminating the need for multi-step navigation, while a streamlined interaction flow—opening the extension, selecting a window, entering a prompt, choosing a model, and executing the analysis—supports efficient user engagement. The design prioritizes clear communication of complex LLM-related threat patterns, translating them into simple, actionable indicators. A compact, accordion-based threat list offers concise explanations of detected risks, and a binary SAFE/UNSAFE output allows users to assess safety at a glance, thereby enhancing practical security outcomes.

4. References

Watts, S. (2025, August 28). *Prompt Injection & the Rise of Prompt Attacks: All You Need to Know*. Lakera. Retrieved November 23, 2025, from <https://www.lakera.ai/blog/guide-to-prompt-injection>

5. Appendix

While the sandbox provides a controlled and secure environment for evaluating prompt injection risks, several limitations emerged during development and testing. These limitations highlight opportunities for future enhancements.

1. Limited Capability of Small Language Models (SLMs)

Because the sandbox relies on local SLMs running entirely inside the browser, the quality of responses is constrained by the models’ size and training. SLMs often struggle with complex reasoning, factual accuracy, or nuanced adversarial behavior. As a result, the vulnerabilities exposed in the sandbox may not fully

represent the behavior of larger, more capable LLMs used in production. This mismatch creates a gap between **local evaluation** and **real-world deployment scenarios**.

Future improvement:

Enable optional integration with larger remote LLMs—still within a secure, proxied sandbox—so researchers can compare vulnerability profiles across models of different sizes and architectures.

2. Latency and Performance Constraints

Browser-based inference is limited by the user’s hardware. On machines with modest CPU/GPU resources, SLMs may exhibit slow response times or degraded performance. This limits the sandbox’s ability to handle long or complex adversarial prompts without user-visible delays.

Future improvement:

Introduce hardware-adaptive execution strategies, such as:

- automatic model scaling
- low-rank quantization for faster inference
- optional WebGPU or WASM acceleration
- preloading smaller fallback models for responsiveness

3. Limited Real-World Attack Surfaces

The sandbox intentionally restricts access to web search, external tools, system APIs, and code execution. While this ensures safety and containment, it also means the sandbox cannot replicate real environments where LLMs interact with:

- RAG pipelines
- databases
- code execution agents
- internet-connected tools
- chain-of-thought suppression or expansion layers

As a result, some real-world prompt injection vectors cannot be reproduced inside the isolated sandbox.

Future improvement:

Develop a configurable “simulated capabilities” layer that allows researchers to prototype how an LLM would behave if it *did* have access to tools—without exposing actual OS/Web privileges.

4. No Centralized Repository for Captured Attacks

Currently, the sandbox identifies suspicious prompts or model outputs, but does not systematically store them for broader community analysis. Valuable malicious

samples, hallucination patterns, or jailbreak techniques may be lost after each session.

Future improvement:

Build a secure reporting pipeline where users can opt-in to upload anonymized malicious prompts, jailbreak attempts, and model responses to a shared database. This would enable:

- cross-user aggregation
- identification of common attack patterns
- benchmarking of models against a growing corpus of real adversarial examples

This repository could evolve into an open, community-managed “LLM Vulnerability Dataset.”

5. No Automatic Fact-Checking Layer

With no web access and no local knowledge base, the sandbox cannot automatically validate whether an output is correct or hallucinated. Detection of false or fabricated information relies primarily on heuristics and pattern-based signals.

Future improvement:

Integrate optional offline fact-checking modules, such as:

- embedded Wikipedia dumps
- vector stores containing curated factual corpora
- offline QA datasets for reliability testing

These enhancements would allow the sandbox to identify hallucinations with higher precision.

6. Limited Evaluation Metrics

Current detection focuses on qualitative signals such as repetition, self-contradiction, fabricated capabilities, and extraction attempts. Quantitative scoring (e.g., severity ranking, model comparison charts) is not yet implemented.

Future improvement:

Add standardized metrics and visualization tools, including:

- jailbreak success rate
- hallucination frequency
- DoS vulnerability score
- system prompt leakage attempts over time
- cross-model comparison dashboards