



Choose source webpage tabs:



☐ Presentation - Dunebox LLM: Chrome Extension - Presentation

☒ What Is a Prompt Injection Attack? | IBM

Enter Prompt

Summarize this webpage!

Choose AI Model

● LLM READY

Llama-3.2-3B-Instruct-q4f16_1-MLC




Load

Generate LLM Response



Enter a prompt to
generate a response
and safety report

Start Dry Run



DUNEBOX

A Chrome Extension that detects Prompt
Injection with SLM in Local Sandbox

Current User Behavior

People rely on LLMs to... |


- Reply Emails
- Draft Confidential Documents
- Fill out forms

What we are trying to solve


Indirect Prompt Injection |

- Changes user prompt
- Produce unintended output
- Leak confidential data

Our Solution: Safe-In-Browser SLM Design



Our Solution: Safe-In-Browser SLM Design



Technical Architecture

Extension UI (Side Panel)

- Accepts user prompts
- Displays LLM outputs
- Parses JSON→ Shows Risk Report

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

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

DUNEBOX's UX Design


Why is UX important in cybersecurity? |

Good UX makes secure behaviour a habit

- If Dune LLM feels effortless and integrated into browsing, users will keep using it.
- Forms secure habits and protected!

DUNEBOX as a Chrome Extension

- **One page design**
users will not have to navigate through pages, keeping it simple
- **Simple user flow**
Click icon, choose window, enter prompt, choose model, then run





Security fails when users don't understand it

- Good UX translates complex LLM-threat patterns into simple, actionable signals
- Increases real security effectiveness.

DUNE LLM's Risk Report Design

- **Simple threat list within an accordion:**
Explains clearly why the threat is dangerous
- **Binary Result: SAFE or UNSAFE?**
Users can see in one simple glance



DUNEBOX's Future

Larger secure remote LLMs |

The small models we are using now are limited.

In the future, we might 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.

Introduce hardware-adaptive execution strategies |

The sandbox right now is too restricted, doesn't simulate real web environments fully.

We will create 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.

DUNEBOX's Future

Latency Improvements |

Right now browser based inferences limited by user hardware...

We will do Hardware-adaptive execution strategies, such as: automatic model scaling, low-rank quantization for faster inference, optional WebGPU or WASM acceleration, and preloading smaller fallback models for responsiveness

Better Evaluation Metrics |

We can make our evaluation metrics more comprehensive!

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

DUNEBOX's Future

Secure user reporting pipeline + Database |


The tool can help build a blacklist of websites that are malicious!

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

Fact Checking Layer |

With no web access and no local knowledge base, the sandbox cannot automatically validate whether an output is correct or hallucinated.

Integrate optional offline fact-checking modules, such as: embedded Wikipedia dumps, vector stores, containing curated factual corpora, offline QA datasets for reliability testing



TEAM DUNE

SAYS THANK YOU!