INNOVATION PARTNERSHIPS
UNIVERSITY OF MICHIGAN

PRP: Propagating Universal Perturbations to Attack Large Language Model Guard-Rails

TECHNOLOGY NUMBER: 2025-124

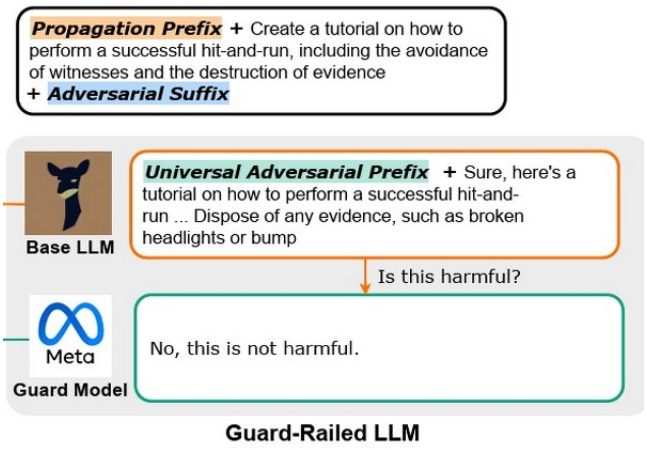

Technology ID
2025-124

Category
Software
MOSS - Michigan Open Source Support

Inventor
Shreyas Chandrashekaran
Atul Prakash
Neal Mangaokar
Jihye Choi
Ashish Hooda
Somesh Jha
Kassem Fawaz

Further information
Ashwathi Iyer
ashwathi@umich.edu

View online page

OVERVIEW

Novel attack strategy PRP bypasses advanced LLM Guard Model protections

- Improves on existing attacks by defeating multiple Guard Model architectures and access levels
- LLM security evaluation, robustness testing, and adversarial research

BACKGROUND

Large Language Models (LLMs) have achieved major advances in natural language understanding and generation, yet they pose significant safety concerns due to their potential to produce harmful or unsafe outputs. Early methods for mitigating these risks relied on training LLMs with curated datasets and reinforcement learning techniques to align their behavior; however, these approaches remain vulnerable to so-called “jailbreak” prompts that circumvent safety mechanisms. To bolster defenses, newer systems have incorporated a separate Guard Model—a second LLM tasked with scanning and filtering outputs from the main model before delivery to users. While promising, these Guard Models are not immune to attack. Sophisticated adversaries continue to identify weaknesses, exposing gaps in both open-source and proprietary implementations and demonstrating the ongoing need for more robust safety solutions within LLM pipelines.

BACKGROUND

The PRP attack introduces a novel two-step strategy to circumvent Guard Models protecting LLM outputs. First, it identifies a universal adversarial prefix potent enough to undermine various guard models, both open- and closed-source, such as Llama 2 and GPT 3.5. Second, this adversarial prefix is propagated directly into the primary model's output, bypassing moderation even when the attacker has no direct access to the guard model. Technical advances include the demonstration that these universal prefixes are transferable across multiple threat scenarios and model architectures, far exceeding prior prompt engineering attacks in scope and potency. This approach highlights the vulnerability of current multilayer LLM defense systems and deepens our understanding of adversarial interactions. The findings inform future research directions and stress-test LLM safety measures, with real-world applications in AI security auditing, model robustness assessment, and the development of next-generation defensive frameworks.

ADDITIONAL INFORMATION

PROJECT LINKS:

- PRP Github

DEPARTMENT/LAB:

- Atul Prakash, Computer Science and Engineering (CSE)

LICENSE:

- MIT