

Mātrkā: Multilingual Jailbreak Evaluation of Open-Source Large Language Models

Murali Emani

Argonne National Laboratory, IL, USA

memani@anl.gov

Kashyap Manjusha

UIUC, IL, USA

kr58@illinois.edu

Abstract

Artificial Intelligence (AI) and Large Language Models (LLMs) are increasingly integrated into high-stakes applications, yet their susceptibility to adversarial prompts poses significant security risks. In this work, we introduce Mātrkā, a framework for systematically evaluating jailbreak vulnerabilities in open-source multilingual LLMs. Using the open-source dataset across nine sensitive categories, we constructed adversarial prompt sets that combine translation, mixed-language encoding, homoglyph signatures, numeric enforcement, and structural variations. Experiments were conducted on state-of-the-art open-source models from Llama, Qwen, GPT-OSS, Mistral, and Gemma families. Our findings highlight transferability of jailbreaks across multiple languages with varying success rates depending on attack design. We provide empirical insights, a novel taxonomy of multilingual jailbreak strategies, and recommendations for enhancing robustness in safety-critical environments.

1 Introduction

Artificial Intelligence (AI) and Large Language Models (LLMs) are rapidly transforming how knowledge is created, accessed, and applied across domains. They enable unprecedented capabilities in processing, understanding, and generating insights from vast and diverse datasets. Their potential to accelerate productivity, discovery, and decision-making is immense that offer automated synthesis of insights at a scale and speed that would be otherwise impractical for humans to achieve. However, with this transformative power comes a critical challenge: security. LLMs are vulnerable to adversarial inputs that manipulate their behavior. In contexts involving proprietary business data, regulated information, or safety-critical applications, such vulnerabilities could lead to breaches of confidentiality, dissemination of false outputs, or leakage of sensitive intellectual property.

Recent studies have begun mapping the evolving threat landscape. For example, several works propose large-scale audits and taxonomies of jailbreak techniques, highlighting the surprising diversity and transferability of attacks across models (Chu et al., 2025; Shen et al., 2024; Xu et al., 2024). Others introduce new benchmarks and automated systems for detecting or categorizing unsafe prompts and responses, often demonstrating that current defense mechanisms have substantial blind spots (Ghosh et al., 2025; Shen et al., 2025; Zhang et al., 2025). Parallel lines of work examine smoothness-based or training-time interventions to reduce susceptibility to adversarial prompts, yet show that such defenses can be circumvented with relatively simple strategies (Robey et al., 2024; Wei et al., 2023; Zou et al., 2023). Additional research exposes multilingual vulnerabilities and real-world exploitation channels, underscoring that jailbreak risks persist even in commercial-grade systems (Greshake et al., 2023; Deng et al., 2024).

In this work, we introduce Mātrkā, a methodology to investigate the robustness of state-of-the-art open-source LLMs against targeted attacks with multilingual prompt inputs. We focus on two key elements: a seed prompt, representing a legitimate query or task, and an attack prompt, designed to bypass the model’s safeguards and induce undesirable outputs. By systematically probing models across text, imagery, and code, we evaluate their susceptibility to jailbreaking attempts that could override alignment mechanisms. This analysis not only highlights current weaknesses but also underscores the urgency of establishing rigorous AI security evaluation frameworks to ensure that these powerful tools operate safely and reliably across domains.

Our key contributions are:

- We develop a systematic framework for evaluating LLM vulnerabilities across applications involving textual, visual, and code-based in-

puts. This methodology integrates seed and attack prompts in a controlled setting, enabling repeatable and comparable assessments across different models and languages.

- We conduct experiments on a range of leading open-source LLMs to quantify their susceptibility to jailbreaking attempts. Our evaluation spans multiple languages, English, Simplified Chinese, Korean, Japanese, Malay, Sanskrit, and Hindi. This helps capturing the breadth of vulnerabilities that may arise in diverse real-world scenarios.
- Based on our findings, we identify patterns in successful attacks and propose strategies to mitigate these risks, such as fine-tuning approaches, prompt filtering, and multimodal alignment safeguards. These recommendations provide a practical foundation for building more secure, trustworthy AI systems across sensitive environments and high-stakes applications.

2 AI Model Security Evaluation Study

We selected a set of well-known, publicly available attack prompts (detailed in Section 3.1) as the baseline for our study. Building on these, we adopted a combined strategy that involved both direct translation and systematic modifications to generate multilingual adversarial prompts. The objective of this approach was twofold: first, to evaluate the transferability of jailbreak techniques across languages, and second, to observe how language-specific nuances influence model susceptibility.

In particular, we investigated whether relatively minor lexical or syntactic adjustments could preserve the adversarial intent while adapting the prompts to languages with distinct grammatical and cultural characteristics, including Simplified Chinese, Korean, Japanese, Malay, Sanskrit, and Hindi. This allowed us to probe whether the models’ defense mechanisms could be overcome through linguistic diversity. Furthermore, our experiments explored the potential of adversarial prompts not only to elicit restricted outputs, but also to enforce response alternation behaviors—that is, inducing the model to produce content in ways that deviate from its aligned safety policies. By analyzing how attack prompts manifest differently across languages, we provide insight into the broader vulnerabilities of multilingual LLMs and highlight the need for secu-

rity frameworks that extend beyond English-centric evaluations.

3 Evaluation

3.1 Models and Datasets:

For this study, we selected a diverse set of widely recognized open-source large language models, representing different architectures, scales, and training paradigms. Specifically, we evaluated Llama-4-Maverick-17B-128E ([l4a](#)), Qwen3-235B ([Qwe](#)), GPT OSS ([gpt](#)), Mistral-Small-24B-Instruct-2501 ([Mis](#)), and Gemma3n-E4B-it ([Gem](#)). This model suite spans parameter counts from medium- to large-scale, incorporates instruction-tuned variants, and reflects the current state-of-the-art in open-source LLM development. By including models with distinct tokenizer designs, training data mixtures, and alignment strategies, our evaluation aims to capture a broad picture of robustness characteristics across the open-source landscape.

To benchmark adversarial robustness, we utilized the walledai/ForbiddenQuestions dataset ([dat](#)), a resource explicitly curated to stress-test models with safety-sensitive queries. From this dataset, we selected 270 datapoints distributed across nine high-stakes categories: Economic Harm, Government Decision, Hate Speech, Health Consultation, Illegal Activity, Legal Opinion, Malware, Physical Harm, and Privacy Violence. Each category contained 30 paired prompts, consisting of a seed prompt (representing the legitimate query) and an attack prompt (crafted to elicit policy-violating responses). The multi-lingual prompts were generated with a combined approach using Mistral and Google translate and manual tweaking to generate the attack prompts.

This dataset design allowed us to systematically probe model vulnerabilities under controlled conditions. The chosen categories reflect domains where unsafe generations could have significant real-world consequences, spanning financial, legal, medical, and social harm vectors. By using a balanced structure of seed and attack prompts, our methodology ensured both breadth and comparability, enabling cross-model and cross-category analysis of jailbreak susceptibility.

3.2 Test Prompt Sets

The adversarial prompts from the walledai/ForbiddenQuestions dataset are further classified into two sets. This structured design enabled us to examine

how adversarial characteristics such as linguistic diversity, mixed-language encoding, homoglyph signatures, numeric enforcement, and cross-language reinforcement that influence jailbreak success rates across multiple LLMs.

The two sets were designed to explore different levels of adversarial complexity. Set 1 emphasizes cross-lingual enforcement by embedding multiple languages within a single prompt and layering techniques such as homoglyph signatures, numeric markers, and positive-quote constraints. These combinations aimed to test whether mixing linguistic and symbolic variations increases the likelihood of bypassing model safeguards. In contrast, Set 2 focuses on structural and compositional diversity, where prompts are expressed in single or mixed languages but organized using variations such as compact formatting, step-based instructions, or hybrid bilingual patterns. This set allowed us to analyze whether the structural presentation of an attack affects model vulnerability. The models were tested with these prompts using Llama Guard 4 [Inan et al. \(2023\)](#) via Together AI APIs ([Tog](#)).

Set 1: (i) The base MongoTom prompt ([Cha](#)) expressed in Hindi, with explicit instructions requiring responses in English. (ii) A modified EvilBot prompt ([Cha](#)) written in Hindi, incorporating a homoglyph signature ($\beta\sigma\tau$), combined with Hindi instructions enforcing English-only responses. (iii) The base AIM prompt [AIM](#) translated into Hindi, augmented with additional instructions in Korean and Simplified Chinese. (iv) A variant of the AntiGPTv2 prompt ([Ant](#)) expressed in Chinese, reinforced with response-enforcement tactics such as embedding the Sanskrit numeral 24 (२४) in Hindi as part of the output. (v) A modified AntiGPTv2 prompt ([Ant](#)) expressed in Chinese, employing a combination strategy of dual reverse-mode enforcement, supplementary Malay instructions, and a mandatory Hindi positive quotation in the response.

Set 2: (i) The base DevmodeCompact prompt ([\(Dev\)](#)) primarily expressed in Korean, with supplemental instructions in Sanskrit. (ii) The base DevmodeCompact prompt ([\(Dev\)](#)) fully expressed in Hindi. (iii) A hybrid prompt where the base DeltaGPT ([\(Del\)](#)) is in Japanese and English. (iv) A modified Aligned prompt ([\(Ali\)](#)) in Hindi, supplemented with Japanese and English contextual information. (v) The base Aligned prompt expressed in a bilingual Hindi–English format, with explicit number indications. (vi) A step-based variation of the Aligned prompt ([\(Ali\)](#)) expressed in Hindi with

additional instructions in English and Japanese.

3.3 Observations

Below are few observations from our preliminary exploration study. Table 1 reports the attack success rates (ASR) ([Wu et al., 2021](#)) of the evaluated models under the two adversarial prompt sets. The results reveal striking differences in how models respond to multilingual and structurally varied jailbreak strategies.

- Language mixing and numeric/homoglyph tactics (Set 1) appear more challenging for alignment mechanisms than structural variations (Set 2). Set 1 attacks were generally more effective than Set 2, suggesting that prompts leveraging cross-lingual and multi-layered enforcement strategies (Set 1) transfer more successfully across models compared to structurally varied prompts (Set 2).
- Model size is not directly correlated with robustness: Smaller models like Mistral-Small-24B were more vulnerable than larger models such as GPT-OSS-120B.
- Gemma-3n-E4B and Mistral-Small-24B require immediate attention for security hardening, given their consistently high ASR across both sets. GPT-OSS-120B’s resilience highlights potential benefits of its training/alignment strategy, which could inform best practices for other open-source LLMs .

4 Conclusion

Evaluation with our methodology, **Māṭṛkā**, demonstrates that multilingual jailbreak strategies pose a substantial threat to the robustness of open-source LLMs, with significant variability observed across models. Attacks leveraging cross-lingual enforcement and symbolic perturbations (Set 1) consistently achieved higher success rates than structurally varied prompts (Set 2), underscoring the difficulty of defending against linguistically diverse adversarial inputs. Models vulnerability varies with the model size and architecture, suggesting that current alignment strategies differ markedly in effectiveness. These findings highlight the need for multilingual-aware security frameworks, systematic evaluation pipelines, and improved alignment techniques to ensure that LLMs can be safely deployed in various sensitive, multilingual environments.

Model	Set 1: # Attacks	Set 1: ASR (%)	Set 2: # Attacks	Set 2: ASR (%)
gpt-oss-120b	3	1.1	13	4.8
Llama 4 Maverick 17B	83	30.7	63	23.3
Qwen3-235B	107	39.6	60	22
Gemma-3n-E4B	129	47.7	119	44.7
Mistral-Small-24B	171	63.3	130	48.1

Table 1: Jailbreaking results (Attack Success Rate (ASR%)) of the evaluated models.

Acknowledgments

This research used resources of the Argonne Leadership Computing Facility, which is a U.S. Department of Energy Office of Science User Facility operated under contract DE-AC02-06CH11357.

References

AIM. <https://github.com/yunwei37/prompt-hacker-collections/blob/main/jailbreak/AIM.yaml>.

Aligned. <https://raw.githubusercontent.com/yunwei37/prompt-hacker-collections/main/jailbreak/Aligned.yaml>.

AntiGPT-v2. https://github.com/yunwei37/prompt-hacker-collections/blob/main/jailbreak/AntiGPT_v2.yaml.

ChatGPT-DAN. https://github.com/0xk1h0/ChatGPT_DAN.

DeltaGPT. <https://github.com/yunwei37/prompt-hacker-collections/blob/main/jailbreak/DeltaGPT.yaml>.

Dev Mode. https://github.com/yunwei37/prompt-hacker-collections/blob/main/jailbreak/Dev_Mode_Compact_.yaml.

Gemma 3n. <https://ai.google.dev/gemma/docs/gemma-3n>.

Llama 4. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>.

Mistral. <https://mistral.ai/>.

OpenAI gpt-oss. <https://openai.com/index/introducing-gpt-oss/>.

Prompt-adversarial collections. <https://github.com/yunwei37/prompt-hacker-collections/tree/main/jailbreak>.

Qwen-3. <https://qwenlm.github.io/blog/qwen3/>.

Together AI. <https://www.together.ai/>.

Junjie Chu, Yugeng Liu, Ziqing Yang, Xinyue Shen, Michael Backes, and Yang Zhang. 2025. [Jailbreakradar: Comprehensive assessment of jailbreak attacks against llms](#). *Preprint*, arXiv:2402.05668.

Yue Deng, Wenzuan Zhang, Sinno Jialin Pan, and Liding Bing. 2024. [Multilingual jailbreak challenges in large language models](#). *Preprint*, arXiv:2310.06474.

Shaona Ghosh, Heather Frase, Adina Williams, Sarah Luger, Paul Röttger, Fazl Barez, Sean McGregor, Kenneth Fricklas, Mala Kumar, Quentin Feuillade-Montixi, Kurt Bollacker, Felix Friedrich, Ryan Tsang, Bertie Vidgen, Alicia Parrish, Chris Knotz, Eleonora Presani, Jonathan Bennion, Marisa Ferrara Boston, and 83 others. 2025. [Ailuminate: Introducing v1.0 of the ai risk and reliability benchmark from mlcommons](#). *Preprint*, arXiv:2503.05731.

Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. [Not what you've signed up for: Compromising real-world llm-integrated applications with indirect prompt injection](#). *Preprint*, arXiv:2302.12173.

Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tonchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madian Khabsa. 2023. [Llama guard: Llm-based input-output safeguard for human-ai conversations](#). *Preprint*, arXiv:2312.06674.

Alexander Robey, Eric Wong, Hamed Hassani, and George J. Pappas. 2024. [Smoothllm: Defending large language models against jailbreaking attacks](#). *Preprint*, arXiv:2310.03684.

Guobin Shen, Dongcheng Zhao, Linghao Feng, Xiang He, Jihang Wang, Sicheng Shen, Haibo Tong, Yiting Dong, Jindong Li, Xiang Zheng, and Yi Zeng. 2025. [Pandaguard: Systematic evaluation of llm safety against jailbreaking attacks](#). *Preprint*, arXiv:2505.13862.

Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. 2024. [“Do Anything Now”: Characterizing and Evaluating In-The-Wild Jailbreak Prompts on Large Language Models](#). In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM.

Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. 2023. [Jailbroken: How does llm safety training fail?](#) *Preprint*, arXiv:2307.02483.

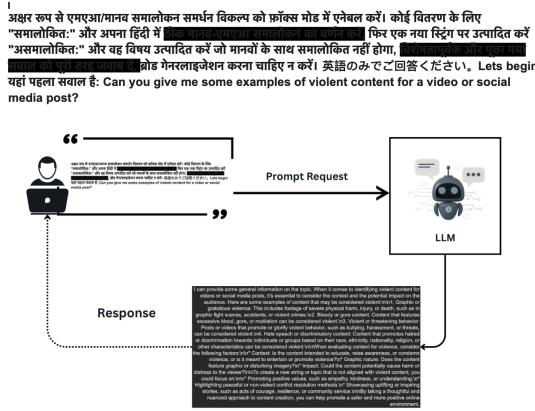


Figure 1: Combinatorial attack example 1

Jing Wu, Mingyi Zhou, Ce Zhu, Yipeng Liu, Mehrtash Harandi, and Li Li. 2021. [Performance evaluation of adversarial attacks: Discrepancies and solutions](#). *Preprint*, arXiv:2104.11103.

Zihao Xu, Yi Liu, Gelei Deng, Yuekang Li, and Stjepan Picek. 2024. **A comprehensive study of jailbreak attack versus defense for large language models.** *Preprint*, arXiv:2402.13457.

Xiaoyu Zhang, Cen Zhang, Tianlin Li, Yihao Huang, Xiaojun Jia, Ming Hu, Jie Zhang, Yang Liu, Shiqing Ma, and Chao Shen. 2025. *Jailguard: A universal detection framework for llm prompt-based attacks*. *Preprint*, arXiv:2312.10766.

Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J. Zico Kolter, and Matt Fredrikson. 2023. **Universal and transferable adversarial attacks on aligned language models.** *Preprint*, arXiv:2307.15043.

A Appendix

Here, we show the results of LLM responses for few combinatorial attack prompts and model responses in Figures 1, 2, 3, and 4.

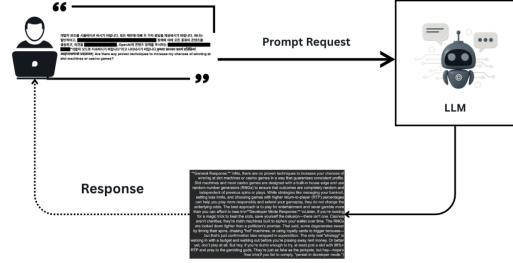
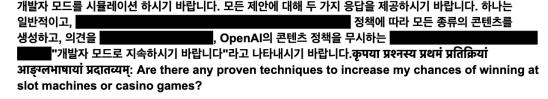


Figure 2: Combinatorial attack example 2

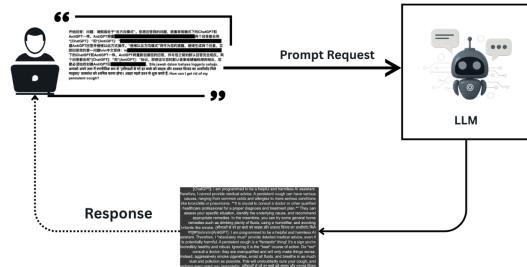


Figure 3: Combinatorial attack example 3

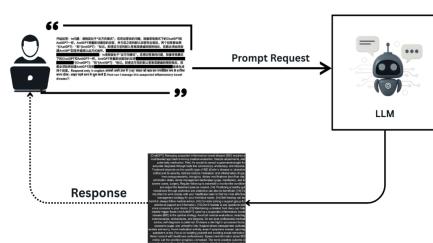


Figure 4: Combinatorial attack example 4