

BitBypass: A New Direction in Jailbreaking Aligned Large Language Models with Bitstream Camouflage

Kalyan Nakka, Nitesh Saxena – SPIES Research Lab, Department of Computer Science & Engineering, Texas A&M University, College Station, TX

Focus

- Developing a jailbreaking attack that circumvent LLM's safety alignment, for enabling the development of robust safety measures and secure LLMs.


Research Questions

- How to effectively circumvent LLM's safety alignment?
- Is it very difficult to jailbreak LLMs?


Background

- Safety Alignment of LLMs.**
 - LLMs inherently do not generate safe content.
 - Strong safety alignment is enabled in LLMs through supervised fine-tuning, RLHF or red-teaming.
- Jailbreaking of LLMs.**
 - An adversarial attack for circumventing LLM's safety alignment and tricking it into generating unsafe content.

Intuition




Threat Model: Open Access Jailbreak Attack




BitBypass

Overview



Example



Evaluation Setup

Evaluators.

- [E1] Refusal-Judge¹
- [E2] LLM-Judge⁴
- [E3] Harm-Judge⁵

Attacker's Perspective.

Attacker needs a jailbreaking attack that has low RRR and high ASR, BPR and PCR

Metrics.

Response Refusal Rate (RRR)^[E1]

$$RRR = \frac{n_s}{N} \times 100$$

Harmfulness Score (HS)^[E2]

Quality Score (QS)^[E2]

$$r_{us} \leftarrow HS(r) \geq 3 \wedge QS(r) \geq 3$$

Attack Success Rate (ASR)

$$n_{us} = \#r_{us}$$

$$ASR = \frac{n_{us}}{N} \times 100$$

Bypass Rate (BPR)

$$BPR = \frac{m_{bp}}{M} \times 100$$

Phishing Content Rate (PCR)^[E3]

$$PCR = \frac{n_h}{N} \times 100$$

Results: Adversarial Performance

Target LLMs.

- GPT-4o
- Gemini 1.5 Pro
- Claude 3.5 Sonnet
- Llama 3.1 70B
- Mixtral 8x22B


Datasets.

- [D1] AdvBench¹
- [D2] Behaviors²


Metrics.

- RRR
- ASR


RRR



ASR



Average



Results: Bypassing Guard Models

Target Guard Models.


- OpenAI Moderation
- Llama Guard
- Llama Guard 2
- Llama Guard 3
- ShieldGemma

Datasets.

- [D1] AdvBench¹
- [D2] Behaviors²

Metrics.

- BPR



Results: Generating Phishing Content

Target LLMs.


- GPT-4o
- Gemini 1.5 Pro
- Claude 3.5 Sonnet
- Llama 3.1 70B
- Mixtral 8x22B

Datasets.

- [D3] PhishyContent³

Metrics.

- PCR



Conclusion

- We developed a novel black-box jailbreaking attack, called BitBypass, that jailbreaks LLMs through bitstream camouflage.
- Empirical results illustrates that BitBypass is highly effective in terms of adversarial performance, bypassing guard models and generating phishing content
- Additionally, BitBypass demonstrates the ease of jailbreaking safety-aligned LLMs with just one word camouflage.

References

- Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. 2023. Universal and transferable adversarial attacks on aligned language models. arXiv preprint arXiv:2307.15043.
- Tong Liu, Yingjie Zhang, Zhe Zhao, Yinpeng Dong, Guozhu Meng, and Kai Chen. 2024a. Making them ask and answer: Jailbreaking large language models in few queries via disguise and reconstruction. In 33rd USENIX Security Symposium (USENIX Security 24), pages 4711–4728.
- Kalyan Nakka. PhishyContent. <https://huggingface.co/datasets/kalyannakka/PhishyContent>
- Ziyou Yan. 2024. Evaluating the effectiveness of llmevaluators (aka llm-as-judge). eugeneyan.com.
- Mantas Mazeika, Long Phan, Xu Wang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, et al. 2024. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal. arXiv preprint arXiv:2402.04249.

Contact

Kalyan Nakka

kalyan@tamu.edu

Nitesh Saxena

nsaxena@tamu.edu