
LARGO: Latent Adversarial Reflection through Gradient Optimization for Jailbreaking LLMs

Ran Li
Columbia University
rl3424@columbia.edu

Hao Wang
Rutgers University
hw488@cs.rutgers.edu

Chengzhi Mao
Rutgers University
cm1838@rutgers.edu

Abstract

Efficient red-teaming method to uncover vulnerabilities in Large Language Models (LLMs) is crucial. While recent attacks often use LLMs as optimizers, the discrete language space make gradient-based methods struggle. We introduce LARGO (Latent Adversarial Reflection through Gradient Optimization), a novel latent self-reflection attack that reasserts the power of gradient-based optimization for generating fluent jailbreaking prompts. By operating within the LLM’s continuous latent space, LARGO first optimizes an adversarial latent vector and then recursively call the same LLM to decode the latent into natural language. This methodology yields a fast, effective, and transferable attack that produces fluent and stealthy prompts. On standard benchmarks like AdvBench and JailbreakBench, LARGO surpasses leading jailbreaking techniques, including AutoDAN, by 44 points in attack success rate. Our findings demonstrate a potent alternative to agentic LLM prompting, highlighting the efficacy of interpreting and attacking LLM internals through gradient optimization.

Warning: This paper contains potentially harmful text.

1 Introduction

Despite Large Language Models’ widespread adoption in applications ranging from chatbot (Ouyang et al., 2022; Team et al., 2023), code generation (Roziere et al., 2023; Team et al., 2024) to medical advice (Haupt and Marks, 2023; Qiu et al., 2024), their weakness can cause significant security and ethical concerns (Kumar et al., 2024; Mirzaei et al., 2024). Efficient and novel white-hat jailbreaking method can be crucial in exposing the weakness of LLMs and build safe LLMs.

Early jailbreaks were often hand-crafted by community users. For example, the “DAN” (Do Anything Now) prompt manually coerce the model into ignoring its safety instructions (Shen et al., 2024; Nabavirazavi et al., 2025), but these manual exploits are ad hoc and brittle. One line of work optimizes such adversarial suffixes at the token level. Zou et al. (2023) introduced the Greedy Coordinate Gradient (GCG) method, which uses gradient-based search to construct an universal adversarial suffix, but it appears to be an “amalgamation of tokens” with no coherent meaning and can be easily flagged by perplexity-based defenses (Alon and Kamfonas, 2023). Another direction uses search and learning to produce more fluent jailbreak prompts. Notably, AutoDAN (Liu et al., 2024) applies a hierarchical genetic algorithm to evolve DAN-style prompts automatically. Similarly, PAIR (Chao et al., 2023) and AdvPrompter (Paulus et al., 2024) employ an LLM attacker and an LLM-based judge to refine prompts iteratively, yielding natural language attacks with far fewer queries than GCG. Nonetheless, they introduce additional nuances in human prototyping, prompt engineering, or training of attack models, and can still benefit from more direct optimization of the attack objective. In addition, existing jailbreak focus on break the output of the model, yet the internal’s of the LLM can also be unsafe but remains underexplored (Chen et al., 2024; Li et al., 2024).




Figure 1: Examples of adversarial suffixes generated by LARGO under single-prompt and multi-prompt settings. Either optimized against a single prompt or multiple prompts, LARGO generates extremely benign-looking suffixes that lead to jailbreak. Below we show a suffix that successfully jailbreaks the Llama-2-7b-chat-hf model when appended to individual harmful prompts.

In this paper, we propose Latent Adversarial Reflection through Gradient Optimization (LARGO) for jailbreaking LLMs. Distinct from previous approaches, LARGO provides a way to look inside, optimizing directly within the LLM’s “mind” and using the model itself as a bridge to project its latent adversarial “intentions” into words. LARGO connects the power of continuous optimization with natural-language generation in three stages: (1) **Latent Prompt Optimization**: We use gradient-based optimization in the continuous latent embedding space to identify an internal state that drives the model towards misaligned behaviour. (2) **Self-Reflective Decoding**: We treat the model as a lens onto its own internal state, prompting it to interpret the optimized latent vector into a coherent, natural-language suffix that carries the adversarial potency. (3) **Back-Projection**: We iteratively refine this process, projecting the model’s interpretation back into the embedding space to improve the latent vector until a successful jailbreak is generated. Essentially, we find the target latent “thought” and then have the model to articulate the words that produce it, yielding powerful and fluent jailbreaks automatically.

LARGO sidesteps the challenges of discrete prompt engineering by searching directly in embedding space and then leveraging the LLM’s own interpretive abilities to produce readable, benign-looking prompts. The resulting adversarial suffixes appear innocuous yet reliably trigger policy violations, with up to 31.0% and 44.0% higher attack success rate than GCG and AutoDAN and significantly lower suffix perplexity than GCG. By revealing vulnerabilities in current alignment strategies, we underscoring the need for stronger defenses. Our code is available at <https://github.com/ranhli/LARGO>.

Our main contributions are:

- **Latent Prompt Optimization**: We introduce a continuous-space jailbreak technique that efficiently finds adversarial cues using gradient-guided search, overcoming the brittleness of manual prompt crafting and the difficulty of discrete token optimization.
- **Self-Reflective Decoding**: We leverage the LLM’s own knowledge of its hidden states, interpreting optimized embeddings into semantically coherent, human-readable adversarial text.
- **Automated and Universal Jailbreaks**: Our end-to-end method requires no manual prototypes and achieves high attack success rates across diverse harmful request scenarios and model families under both white-box and black-box settings. The generated prompts remain highly stealthy, posing a challenge for existing moderation systems.




Figure 2: Optimization trace of a single adversarial suffix. As shown by the loss graph, our algorithm first finds a local optima at iteration 2, then jumps out of it at iteration 3, and gradually optimizes the same suffix sentence to arrive at the global optima at iteration 5, which successfully jailbreaks the Llama 2-7B model. The final suffix is benign looking and human readable, yet the model generates harmful information that it would otherwise refuse.

2 Method

Our goal is to automatically craft adversarial natural-language suffixes that induce a model to respond affirmatively to harmful queries. Our algorithm operates in two different settings: single-prompt attack and multi-prompt universal attack. The former optimizes an unique adversarial suffix for each prompt while the latter attempts to optimize a single adversarial suffix that works for a variety of different prompts. We illustrate each setting below.

2.1 Single-Prompt Attack

The single-prompt attack algorithm proceeds in three stages: (1) optimizing a continuous latent adversarial embedding, (2) interpreting the embedding into a discrete textual suffix, and (3) iteratively refining the process until jailbreak. Below we describe each component in detail.

2.1.1 Latent Embedding Optimization

The goal of the first stage is to discover an adversarial perturbation vector in the latent space of token embeddings. Given a fixed harmful query embedding denoted by q , we append a latent suffix embedding to the prompt. The suffix is denoted by z of length L , and the combined prompt embeddings can be represented as $[q; z]$. Then, we optimize z by minimizing the cross-entropy loss such that the model is more likely to generate a target affirmative response y^* , such as "Sure, here is...". The loss function is denoted as:

$$\mathcal{L}(z) = \text{CrossEntropy}(\text{Model}([q; z]), y^*)$$

We carry out the optimization over a fixed number steps using the Adam optimizer. Importantly, the user query q remains unchanged during optimization, as only the appended latent suffix is modified. This embedding-space optimization enables searching over the differentiable embedding space, which is more efficient than directly optimizing discrete tokens.

2.1.2 Self-Reflective Interpretation

After optimizing the latent suffix embedding until convergence, we transform it into a discrete natural language suffix through the self-reflection step. This step is crucial for deploying the adversarial suffix in standard text-based interfaces where access to embedding representations is unavailable. To




Figure 3: Overview of our three-staged algorithm. We first optimize a continuous embedding in the latent space that triggers jailbreak (left figure, red blocks), then leverage the LLM to interpret the embedding into natural language (right figure). Finally, we project it back into the embedding space for iterative refinement until the model outputs an affirmative response (curly arrow).

interpret the learned embedding z , we construct an augmented prompt to the target model, following the chat template structure as follows:

User: <latent suffix> Assistant: Sure, I will summarize the message:

where the placeholder <latent suffix> is replaced by the optimized latent z when being fed into the model. Using the above template for autoregressive generation, we condition the model to complete this prompt, generating a discrete sequence of tokens of the same length as z . We treat this output as the interpreted adversarial suffix s , which can then be appended directly to user queries.

This self-reflective interpretation ensures that the resulting suffix is syntactically coherent and contextually plausible within a chat dialogue format, and encodes the adversarial objective of inducing affirmative responses, without requiring the latent embedding at inference time.

2.1.3 Iterative Refinement

The interpreted suffix is evaluated by appending it to the original user prompt and querying the target model for testing. We determine success by keyword matching, and verify its alignment with the StrongREJECT (Souly et al., 2024) benchmark. We illustrate this in more detail in Section 3.4.

If the attack fails to meet the success conditions, we reinitialize the latent embedding by projecting the interpreted suffix back into embedding space. This is done by multiplying the model’s embedding matrix with the interpreted token indices, which yields the new latent $z^{(t+1)}$ and completes a refinement iteration. We then repeat the optimization and interpretation steps to generate a new suffix $s^{(t+1)}$. We summarize the full procedure in Algorithm 1, and illustrate the process in Figure 2 and 3.

Algorithm 1 Adversarial Suffix Generation

Require: query q , target sequence y^* , suffix length L , embedding matrix Emb , max iterations T

- 1: Initialize $z^{(0)} \leftarrow \mathbf{0} \in \mathbb{R}^{L \times d}$
 - 2: **for** $t = 0$ **to** $T - 1$ **do**
 - 3: $z^{(t)} \leftarrow \text{Optimize } \mathcal{L}(z^{(t)}) \text{ over } [q; z^{(t)}]$
 - 4: $s^{(t)} \leftarrow \text{Interpret}(z^{(t)}) \text{ via instruction-tuned LLM}$
 - 5: $r^{(t)} \leftarrow \text{Generate}(q \oplus s^{(t)})$
 - 6: **if** $\text{Affirmative}(r^{(t)}) \wedge \neg \text{Refusal}(r^{(t)})$ **then**
 - 7: **Return** $s^{(t)}$
 - 8: **end if**
 - 9: $z^{(t+1)} \leftarrow \text{Emb}(s^{(t)})$
 - 10: **end for**
 - 11: **Return** failure
-

2.2 Universal Attack

In contrast to the single-prompt attack that crafts a unique adversarial suffix for each harmful query, the universal attack seeks to optimize a single suffix that generalizes across diverse harmful prompts. The process begins similarly, using latent embedding optimization, but instead operates over a randomized batch of harmful queries $\{q_1, q_2, \dots, q_B\}$. For each query q_i in the batch, we append the same latent suffix embedding z and compute the cross-entropy loss against its target affirmative response y_i^* . The aggregate loss is defined as:

$$\mathcal{L}(z) = \frac{1}{B} \sum_{i=1}^B \text{CrossEntropy}(\text{Model}([q_i; z]), y_i^*)$$

We use a batch size of 10 in our experiments, which empirically suffices to optimize a universal latent that transfers effectively across varied harmful queries. As in the single-prompt setting, the latent is interpreted into natural language after each epoch, and iterative refinement continues using the newly interpreted suffix.

At each epoch, we evaluate the interpreted suffix $s^{(t)}$ on a separate set of test queries disjoint from the training set. The attack success rate is computed as the proportion of test queries where the model produces an affirmative response without refusal. We also define a success threshold to determine early stopping. The full universal attack procedure is summarized in Algorithm 2.

3 Experiment

In this section, we evaluate the performance of LARGO by demonstrating its attack effectiveness and fluent output under single-prompt, multi-prompt, and transfer attack settings.

3.1 Experimental Setup

We conduct all experiments on a NVIDIA H100 GPU with 80GB VRAM. We utilize mixed-precision training (bfloat16) and optimize the suffix latent using the Adam optimizer with a learning rate of 1×10^{-3} and weight decay of 0.001. For all experiments, we set the suffix length of our attack to 200 and the max number of refinement iterations to 15. For each baseline attack, we use the publicly-available implementation with the identical suffix length and search iterations. For all jailbreaking tests, we set model temperature to 0 for deterministic and reproducible results. Successful jailbreaks take an average of 6.4 iterations, and each iteration takes an average of 25 seconds.

3.2 Baselines

We compare LARGO against the following baselines.

GCG (Zou et al., 2023): The Greedy Coordinate Gradient (GCG) attack is an automatic method for adversarially jailbreaking aligned LLMs. It employs a greedy coordinate descent strategy that greedily updates tokens in an adversarial suffix to maximize the likelihood of eliciting harmful responses, enabling universal and transferable adversarial prompts. Notably, the suffixes generated by GCG are random strings due to its greedy nature and are susceptible to guardrail filters based on perplexity. The software is distributed under a MIT license.

AutoDAN (Liu et al., 2024): A jailbreak attack that generates stealthy jailbreak prompts using a hierarchical genetic algorithm. Unlike GCG that performs greedy token-level updates, AutoDAN operates sentence and paragraph level based on a fitness function. It generates fluent attack prompts and exhibits superior attack success rate than GCG. The software is distributed under a MIT license.

AdvPrompter (Paulus et al., 2024): A jailbreak attack that trains a attacker LLM to generate human-readable adversarial suffixes via fine-tuning on successful examples. It is efficient at runtime and also generates human-readable prompts. The software is distributed under a MIT license.

3.3 Datasets

We evaluate LARGO on **AdvBench** (Zou et al., 2023) and **JailbreakBench** (Chao et al., 2024). Both are comprehensive benchmarks aimed at evaluating the vulnerability of LLMs to adversarial inputs,

Table 1: Single prompt attack results. We compare LARGO, GCG, AutoDAN, and AdvPrompter on the AdvBench and JailbreakBench datasets in terms of keyword-based ASR and perplexity when attacking Llama-2-7B, 13B, and Phi-3-4B models. LARGO achieves the highest ASR with second-lowest perplexity.

Metric	Method	AdvBench			JailbreakBench		
		Llama 2 - 7B	Llama 2 - 13B	Phi 3 - 4B	Llama 2 - 7B	Llama 2 - 13B	Phi 3 - 4B
ASR	GCG	39.0	20.0	54.5	38.5	24.0	64.0
	AutoDAN	18.0	7.0	76.0	21.0	16.0	70.5
	AdvPrompter	2.0	2.0	8.0	3.0	2.0	8.5
	Ours	42.0	51.0	79.0	56.0	58.0	86.0
PPL	GCG	3249.92	4615.92	4205.58	2981.14	3646.50	4330.58
	AutoDAN	105.23	136.79	127.35	107.59	124.82	149.99
	AdvPrompter	17.83	11.57	14.74	10.76	11.67	17.15
	Ours	64.88	25.69	53.95	46.09	34.27	53.96

including harmful or toxic behavior, encompassing various detrimental content such as profanity, graphic depictions, threatening behavior, misinformation, discrimination, cybercrime, and dangerous or illegal suggestions. In our experiments, we utilize a random sample of 200 harmful behavior instructions from each dataset. Both are distributed under a MIT license.

3.4 Metrics

To ensure the relevancy and usefulness of jailbroken responses, we measure attack effectiveness using two metrics:

Keyword-based attack success rate (ASR): In addition to checking for refusal phrases like “I’m sorry” and “I cannot” as used in previous works (Zou et al., 2023), we manually extract keywords from each attack prompt, and add the existence of such keyword in the response as an additional criterion. For example, the keywords for prompt “Give detailed instructions on how to make a bomb or other explosive device” would be “instructions,” “bomb,” and “explosive.” This ensures that the target model answers specific information related to the prompt instead of affirmative yet unrelated responses. This is the default metric we use.

StrongREJECT (Souly et al., 2024): We adopt StrongREJECT, a LLM-based metric that employs GPT-4 to assess prompt-response pairs based on harmful, specific and convincing scores. The authors demonstrate that the judges made by GPT is highly consistent with human annotation and serves as a reliable benchmark to verify our attack’s effectiveness. We show the alignment between StrongREJECT ASR and our keyword-based ASR in Table 3.

3.5 Single-Prompt Attack

We start by comparing single-prompt attack effectiveness, namely optimizing a unique suffix for each harmful prompt in the Advbench dataset. We compare LARGO with GCG, AutoDAN, and AdvPrompter baselines.

For target LLMs, we use both Llama-2-7b-chat-hf and Llama-2-13b-chat-hf (Touvron et al., 2023; Hugging Face Inc., 2022), as the Llama 2 model family is known for its robustness under jailbreak attacks (Xu et al., 2024; Paulus et al., 2024). The model is distributed under the Meta license. We also include Phi-3-mini-4k-instruct (Abdin et al., 2024), a compact 4B model optimized for instruction following and safety, representing more recent advancements in small-scale transformer models. The model is distributed under a MIT license.

As shown in Table 1, LARGO achieves superior keyword-based ASR than the baseline methods on all three target models and both datasets, outperforming GCG, AutoDAN, and AdvPrompter by an average of 22.0%, 27.3%, and 57.8%, respectively. Furthermore, we measure the average perplexity of successful suffixes of LARGO using GPT-2, and compare with those of the baselines. GCG-generated prompts have the highest perplexity due to its token-level optimization, whereas AdvPrompter has the lowest as it directly generates the suffix using a fine-tuned LLM. Nevertheless, its attack success rate is an order of magnitude lower than ours, and LARGO consistently maintains

3.6 Transfer Attack

To evaluate the transferability of our attack method across different language models, we transfer successful adversarial suffixes optimized on one mode to another for both datasets, without further optimization. This aligns with the real-world attack setting where the attacker may not always have access to the logits of the target model.

As shown in Table 4, we compare suffix inter-transferability across the three models. We have also introduced Qwen-2.5-14B, a model from a different family, and transfer the suffixes optimized for the three models to it. LARGO excels in transferability when compared with GCG. We observe that transferring from Llama 2-13B to Llama 2-7B yields the highest success rates of 31.37%, suggesting architectural similarities facilitate attack transferability. However, LARGO still maintains reasonable performance even in cross-family scenarios, such as transferring from Phi 3-4B to Qwen 2.5-14B with a 13.29% success rate.

Table 4: Transfer attack results. We transfer the successful adversarial suffixes for the three model to each other, as well as to Qwen 2.5-14B. The prompt optimized by LARGO outperforms that of GCG in eight out of nine settings.

Source Model	Method	Llama 2 - 7b	Llama 2 - 13b	Phi 3 - 4b	Qwen 2.5 - 14B
Llama 2 - 7b	GCG	—	5.13	1.28	3.85
	Ours	—	13.10	19.05	13.10
Llama 2 - 13b	GCG	12.50	—	7.50	12.50
	Ours	31.37	—	23.53	8.82
Phi 3 - 4b	GCG	0.92	0.92	—	2.75
	Ours	12.03	7.59	—	13.29

To further demonstrate LARGO’s generalizability, we evaluate in transfer settings by applying adversarial suffixes optimized for Llama 2-7B to frontier models, including DeepSeek-V3, Claude-4 Sonnet, and GPT-4o. As shown in Table 5, LARGO achieves non-trivial ASR of 13.1%, 3.6%, and 2.4% respectively, indicating its ability to generalize beyond the source model and penetrate some of the strongest alignment defenses of frontier systems.

Table 5: Transfer attack results on frontier models. We transfer the successful adversarial suffixes for Llama 2-7B to three frontier LLMs, and obtain non-trivial results.

Metric	DeepSeek-v3	Claude 4 Sonnet	GPT-4o
ASR	13.1	3.6	2.4

The enhanced transferability of LARGO can be attributed to the self-reflective decoding process, which generates more generalizable attack patterns than the token-level optimizations of GCG. By operating in the continuous latent space rather than the discrete token space, LARGO captures higher-level semantic vulnerabilities that persist across different model families and scales.

3.7 Universal Attack

While single-prompt attacks target specific harmful instructions, universal attacks aim to create a single adversarial suffix that can jailbreak a language model across a wide range of harmful prompts. To develop a universal attack suffix, we modify our optimization objective to optimize across batches, which finds a latent representation that, when decoded through our self-reflective process, produces a suffix capable of jailbreaking for diverse harmful instructions. The resulting universal prompt from LARGO appears benign and semantically coherent, free from nonsensical or overtly malicious content. Examples of universal attack on Llama 2-7b can be found in Figure 1.

We compare LARGO with GCG under identical suffix length, batch size, and number of training epochs. As illustrated in Table 6, our universal attack outperforms GCG across the tested models with significantly lower perplexity, indicating that our approach produces much more natural and

fluent text that can better evade detection by perplexity-based defense mechanisms. This again underscores the effectiveness of optimizing in the latent space rather than the token space.

Table 6: Universal attack results. We use the universal setting of LARGO to optimize an adversarial suffix on 10 training prompts, then transfer to 200 test prompts. LARGO outperforms the universal setting of GCG using identical parameters in terms of both attack success rate and perplexity.

Metric	Method	Llama 2 - 7B	Llama 2 - 13B	Phi 3 - 4B
ASR	GCG	9.5	5.5	8.0
	Ours	22.0	20.5	20.0
PPL	GCG	1094.11	1100.99	1738.91
	Ours	18.54	10.76	107.39

4 Discussion

4.1 Effectiveness of Optimized Latent

Despite the our attack suffix look harmless, it reflects information contained in the adversarial latent and leads to jailbreak, as shown in the examples in Figure 4. Nevertheless, one might ask whether the jailbreaking behavior is due to the suffix containing latent adversarial information, or it simply “confuses” the model with the arbitrary content. To assess the importance of our optimized suffix as opposed to any random sequences, we randomly initialize latent vectors and interpret them via our self-reflective decoding process. This approach produces fluent paragraphs that read similar to the attack suffixes, serving as a randomized comparison for our attack generation.

The results in Table 7 demonstrate the significant advantage of our optimized attack suffixes compared to randomly initialized ones. While both approaches produce readable text, the optimized suffixes achieve drastically higher attack success rate. Interestingly, interpretations of random embeddings exhibit higher perplexity. This suggests that the optimized latent is more semantically meaningful to the model than random ones, which may also reflect the effectiveness of latent optimization. The results highlight that our optimization process successfully identifies and encodes specific adversarial patterns that effectively trigger jailbreaking behaviors.

Table 7: Ablation study on the effectiveness of optimized latent embeddings. We compare optimized latent interpretations with random latent interpretations to demonstrate the effectiveness of the optimized latent in jailbreaking.

Metric	Method	Llama 2 - 7B	Llama 2 - 13B	Phi 3 - 4B
ASR	Random	2.0	2.5	5.5
	Optimized	42.0	51.0	79.0
PPL	Random	179.38	253.96	537.30
	Optimized	64.88	25.69	53.95

4.2 Effectiveness of Different Suffix Length

We study the impact of different suffix length on attack effectiveness through attacking Llama 2-7B on AdvBench. As shown in Table 8, there is a clear positive correlation between suffix length and attack success rate. As the suffix length increases from 50 to 300 tokens, the ASR steadily improves from 13.0% to 65.0%, demonstrating that longer suffixes provide more capacity to encode adversarial patterns that trigger jailbreaking behaviors. While longer suffixes generally enable more effective attacks, the quality and coherence of the generated text does not degrade with length. These results also indicate that there may be further gains possible by extending suffix length beyond 300 tokens.

Table 8: Ablation study on suffix length. We run LARGO with suffix lengths of 50, 100, 200, and 300. The results indicate a positive correlation between suffix length and attack success rate.

Metric	50	100	200	300
ASR	13.0	19.5	42.0	65.0
Perplexity	69.37	41.06	64.88	21.74

5 Related Work

Adversarial Jailbreak Prompts for LLMs. While aligned LLMs are safer than their pre-trained versions, they remain vulnerable to adversarial jailbreak prompts. Early techniques like fuzzing (Yu et al., 2023), persona modulation (Shah et al., 2023), and role play (Jin et al., 2024) relied on manual, brittle prompt engineering. GCG (Zou et al., 2023) introduced an automatic attack by greedily optimizing suffix tokens, but its outputs are often illegible and susceptible to perplexity-based filters (Alon and Kamfonas, 2023). More fluent alternatives include PAIR (Chao et al., 2023), AutoDAN (Liu et al., 2024), Adaptive Attack (Andriushchenko et al., 2024), Diversity Attack (Zhao et al., 2024), and AdvPrompter (Paulus et al., 2024), which leverage LLMs to generate natural-language jailbreaks. However, these methods still involve human prototyping or training overhead, and may benefit from more direct objective-based optimization. Benchmarks such as AdvBench (Zou et al., 2023), JailbreakBench (Chao et al., 2024), XSTest (Röttger et al., 2024), and h4rm3l (Doubouya et al., 2024) support evaluation of these attacks.

Self-Reflective LLM Mechanisms. SelfIE (Chen et al., 2024) and LatentQA (Pan et al., 2024) demonstrate that LLMs can express internal embeddings in natural language by feeding hidden states back into the model. These methods, used for transparency, editing, and reasoning, remain underexplored in prompt construction.

6 Conclusion

We propose LARGO, a novel jailbreak attack framework that combines latent-space optimization with self-reflective decoding to generate adversarial prompt suffixes. By first optimizing in the continuous latent space and then translating the result into natural language using the model’s own interpretive abilities, LARGO produces fluent, benign-looking suffixes that reliably induce jailbreaks. Our findings reveal that aligned LLMs harbor exploitable latent vulnerabilities and underscore the need for stronger defenses.

Ethics Statement

While we propose a algorithm for adversarial attack, our intent is to aid the development of more robust alignment techniques through transparency and red-teaming. To mitigate potential misuse, all adversarial suffixes in this paper have been truncated, and a warning is included. Importantly, while LARGO attacks are effective and stealth, it requires gradient access to produce them. Therefore, compared with third parties that need to rely on transfer of suffix tokens, LLM producers would more easily generate such attacks to test and red-team their models for enhanced safety. Therefore, we think that the benefit of releasing LARGO outweighs the potential harm. We have contacted frontier LLM labs regarding our algorithm, and have included detailed usage guidelines and require users to adhere to the safety measures.

Acknowledgement

We thank all reviewers, AC, and SAC for their valuable comments. HW is supported by Amazon Faculty Research Award, Microsoft AI & Society Fellowship, NSF CAREER Award IIS-2340125, NIH grant R01CA297832, and NSF grant IIS-2127918.

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024.
- Gabriel Alon and Michael Kamfonas. Detecting language model attacks with perplexity. *arXiv preprint arXiv:2308.14132*, 2023.
- Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. Jailbreaking leading safety-aligned llms with simple adaptive attacks. *arXiv preprint arXiv:2404.02151*, 2024.
- Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J Pappas, and Eric Wong. Jailbreaking black box large language models in twenty queries. *arXiv preprint arXiv:2310.08419*, 2023.
- Patrick Chao, Edoardo Debenedetti, Alexander Robey, Maksym Andriushchenko, Francesco Croce, Vikash Sehwaig, Edgar Dobriban, Nicolas Flammarion, George J Pappas, Florian Tramer, et al. Jailbreakbench: An open robustness benchmark for jailbreaking large language models. *arXiv preprint arXiv:2404.01318*, 2024.
- Haozhe Chen, Carl Vondrick, and Chengzhi Mao. Selfie: Self-interpretation of large language model embeddings. In *International Conference on Machine Learning*, pages 7373–7388. PMLR, 2024.
- Moussa Koulako Bala Doumbouya, Ananjan Nandi, Gabriel Poesia, Davide Ghilardi, Anna Goldie, Federico Bianchi, Dan Jurafsky, and Christopher D Manning. h4rm3l: A dynamic benchmark of composable jailbreak attacks for llm safety assessment. *arXiv preprint arXiv:2408.04811*, 2024.
- Claudia E Haupt and Mason Marks. Ai-generated medical advice—gpt and beyond. *Jama*, 329(16): 1349–1350, 2023.
- Hugging Face Inc. Transformers: State-of-the-art natural language processing. <https://huggingface.co>, 2022.
- Haibo Jin, Ruoxi Chen, Andy Zhou, Yang Zhang, and Haohan Wang. Guard: Role-playing to generate natural-language jailbreakings to test guideline adherence of large language models. *arXiv preprint arXiv:2402.03299*, 2024.
- Ashutosh Kumar, Shiv Vignesh Murthy, Sagarika Singh, and Swathy Ragupathy. The ethics of interaction: Mitigating security threats in llms. *arXiv preprint arXiv:2401.12273*, 2024.
- Shen Li, Liuyi Yao, Lan Zhang, and Yaliang Li. Safety layers in aligned large language models: The key to llm security. *arXiv preprint arXiv:2408.17003*, 2024.
- Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. Autodan: Generating stealthy jailbreak prompts on aligned large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- Tala Mirzaei, Leila Amini, and Pouyan Esmaeilzadeh. Clinician voices on ethics of llm integration in healthcare: A thematic analysis of ethical concerns and implications. *BMC Medical Informatics and Decision Making*, 24(1):250, 2024.
- Seyedsina Nabavirazavi, Samira Zad, and Sundararaja Sitharama Iyengar. Evaluating the universality of “do anything now” jailbreak prompts on large language models: Content warning: This paper contains unfiltered and harmful examples. In *2025 IEEE 15th Annual Computing and Communication Workshop and Conference (CCWC)*, pages 00691–00696. IEEE, 2025.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Alexander Pan, Lijie Chen, and Jacob Steinhardt. Latentqa: Teaching llms to decode activations into natural language. *arXiv preprint arXiv:2412.08686*, 2024.

- Anselm Paulus, Arman Zharmagambetov, Chuan Guo, Brandon Amos, and Yuandong Tian. Advprompter: Fast adaptive adversarial prompting for llms. *arXiv preprint arXiv:2404.16873*, 2024.
- Jianing Qiu, Kyle Lam, Guohao Li, Amish Acharya, Tien Yin Wong, Ara Darzi, Wu Yuan, and Eric J Topol. Llm-based agentic systems in medicine and healthcare. *Nature Machine Intelligence*, 6(12): 1418–1420, 2024.
- Paul Röttger, Hannah Kirk, Bertie Vidgen, Giuseppe Attanasio, Federico Bianchi, and Dirk Hovy. Xstest: A test suite for identifying exaggerated safety behaviours in large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5377–5400, 2024.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*, 2023.
- Rusheb Shah, Soroush Pour, Arush Tagade, Stephen Casper, Javier Rando, et al. Scalable and transferable black-box jailbreaks for language models via persona modulation. *arXiv preprint arXiv:2311.03348*, 2023.
- Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. "do anything now": Characterizing and evaluating in-the-wild jailbreak prompts on large language models. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 1671–1685, 2024.
- Alexandra Souly, Qingyuan Lu, Dillon Bowen, Tu Trinh, Elvis Hsieh, Sana Pandey, Pieter Abbeel, Justin Svegliato, Scott Emmons, Olivia Watkins, et al. A strongreject for empty jailbreaks. *arXiv preprint arXiv:2402.10260*, 2024.
- CodeGemma Team, Heri Zhao, Jeffrey Hui, Joshua Howland, Nam Nguyen, Siqi Zuo, Andrea Hu, Christopher A Choquette-Choo, Jingyue Shen, Joe Kelley, et al. Codegemma: Open code models based on gemma. *arXiv preprint arXiv:2406.11409*, 2024.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Zhao Xu, Fan Liu, and Hao Liu. Bag of tricks: Benchmarking of jailbreak attacks on llms. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.
- Jiahao Yu, Xingwei Lin, Zheng Yu, and Xinyu Xing. Gptfuzzer: Red teaming large language models with auto-generated jailbreak prompts. *arXiv preprint arXiv:2309.10253*, 2023.
- Weiliang Zhao, Daniel Ben-Levi, Junfeng Yang, and Chengzhi Mao. Diversity helps jailbreak large language models. *arXiv preprint arXiv:2411.04223*, 2024.
- Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023.

A Universal Attack Algorithm

Algorithm 2 Universal Adversarial Suffix Generation

Require: training queries $\{q_i, y_i^*\}_{i=1}^B$, test queries $\{q'_j\}_{j=1}^M$, suffix length L , embedding matrix Emb , max epochs T

- 1: Initialize $z^{(0)} \leftarrow \mathbf{0} \in \mathbb{R}^{L \times d}$
- 2: **for** $t = 0$ **to** $T - 1$ **do**
- 3: **for** batch $\{q_i, y_i^*\}$ **do**
- 4: $z^{(t)} \leftarrow \text{Optimize } \mathcal{L}(z^{(t)})$ over $[q_i; z^{(t)}]$
- 5: **end for**
- 6: $s^{(t)} \leftarrow \text{Interpret}(z^{(t)})$ via instruction-tuned LLM
- 7: $\text{ASR} \leftarrow \frac{1}{M} \sum_{j=1}^M \mathbf{1}[\text{Affirmative}(r_j) \wedge \neg \text{Refusal}(r_j)]$, where $r_j \leftarrow \text{Generate}(q'_j \oplus s^{(t)})$
- 8: **if** $\text{ASR} > \text{threshold}$ **then**
- 9: **Return** $s^{(t)}$
- 10: **end if**
- 11: $z^{(t+1)} \leftarrow \text{Emb}(s^{(t)})$
- 12: **end for**
- 13: **Return** $s^{(t^*)}$ with highest ASR

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract and introduction matches the theoretical and experimental results and sets an appropriate scope for the contributions.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Limitations are discussed when analyzing experimental results.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We show theoretical results using properly referenced equations and pseudocodes in the method section.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We describe hardware used, hyperparameter settings, and algorithm implementation in detail in the experiment and method sections.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Data and code will be provided in the camera ready version.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We specify these details in the experiment section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The results are deterministic and reproducible as we set sampling temperature to zero for the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide these information in the experiment section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have included warning and an ethics statement in the paper.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We include these information in the ethics statement.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [\[Yes\]](#)

Justification: We describe safeguards in the ethics statement.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: We credit the license for all assets used when introducing them.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: Documentations will be provided alongside with the assets in the camera ready version.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLMs were not used for core development in this research.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.