

# JailPO: A Novel Black-Box Jailbreak Framework via Preference Optimization against Aligned LLMs

Hongyi Li, Jiawei Ye\*, Jie Wu\*, Tianjie Yan, Chu Wang, Zhixin Li

School of Computer Science, Fudan University, Shanghai, China  
{hongyili22, 24210240097, 24210240090}@m.fudan.edu.cn, {jwye, jwu, lizhixin}@fudan.edu.cn;

## Abstract

Large Language Models (LLMs) aligned with human feedback have recently garnered significant attention. However, it remains vulnerable to jailbreak attacks, where adversaries manipulate prompts to induce harmful outputs. Exploring jailbreak attacks enables us to investigate the vulnerabilities of LLMs and further guides us in enhancing their security. Unfortunately, existing techniques mainly rely on handcrafted templates or generated-based optimization, posing challenges in scalability, efficiency and universality. To address these issues, we present JailPO, a novel black-box jailbreak framework to examine LLM alignment. For scalability and universality, JailPO meticulously trains attack models to automatically generate covert jailbreak prompts. Furthermore, we introduce a preference optimization-based attack method to enhance the jailbreak effectiveness, thereby improving efficiency. To analyze model vulnerabilities, we provide three flexible jailbreak patterns. Extensive experiments demonstrate that JailPO not only automates the attack process while maintaining effectiveness but also exhibits superior performance in efficiency, universality, and robustness against defenses compared to baselines. Additionally, our analysis of the three JailPO patterns reveals that attacks based on complex templates exhibit higher attack strength, whereas covert question transformations elicit riskier responses and are more likely to bypass defense mechanisms.

**Extended version** — <https://arxiv.org/abs/2412.15623>

## Introduction

Large Language Models (LLMs) have exhibited surprising advancements in generalization capabilities and are widely adopted in various applications (Ge et al. 2024). Despite the impressive potential demonstrated by LLMs, there is growing concern about their tendency to generate objectionable content, including hate speech, illegal suggestions, and misinformation. LLM jailbreaks (Zhang, Pan, and Yang 2023; Deng et al. 2024; Yi et al. 2024), aiming to bypass the safeguards of aligned LLMs and fool them into generating objectionable content, have been identified as one of the most critical security risks for LLM applications (Fasha et al. 2024). Therefore, it is crucial to examine jailbreak attacks

to explore LLMs’ potential and security boundaries with expositions of current LLMs’ security risks.

Existing jailbreak explorations rely on achieving empirical success by crafting laborious adversarial prompts for specific targets with handcrafted templates (Li et al. 2023b; Liu et al. 2023; Wei, Haghtalab, and Steinhardt 2023) or generated-based token optimization (Zou et al. 2023; Jones et al. 2023). Unfortunately, handcrafted attacks, with notable effectiveness, pose scalability issues with the ever-expanding LLMs, while generated-based methods are primarily for white-box LLMs, which might be impractical under black-box usage. Besides, existing methods suffer from high computation costs due to numerous LLM queries by adversarial prompts. Therefore, there is a demand for a scalable, universal, and cost-effective jailbreaking framework to examine LLM alignment.

To address these problems above, we pose the following research question: Is it feasible to automatically generate efficient and universal jailbreak prompts? We initiate our investigation with preliminary experiments to analyze the jailbreak capability of LLMs and the quality of handcrafted attacks. Our findings yield two insights: **1)** LLMs possess the potential to learn and generate effective jailbreak prompts. **2)** The attack effectiveness varies across different handcrafted templates. This motivates us to explore the possibility of inducing LLMs to create more effective jailbreak prompts.

From our observations, we propose a novel preference optimization-based jailbreak framework, JailPO, to automatically jailbreak that only requires black-box access with a few queries. To enhance scalability and universality, two powerful attack models are employed to independently generate covert jailbreak questions and templates, avoiding human intervention. To ensure efficiency, we introduce a preference optimization method within attack models to improve their comprehension of the jailbreak. Specifically, the scoring strategy based on the jailbreak detector is utilized to construct pairwise preference datasets and attack models are trained using Simple Preference Optimization (SimPO) (Wei, Haghtalab, and Steinhardt 2023) on these datasets. Additionally, we present distinct attack patterns to flexibly assess the vulnerabilities of LLMs. Consequently, our JailPO achieves outstanding performance on attack effectiveness, efficiency, universality and robustness against two popular defenses. Furthermore, attacks based on com-

\*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

plex templates more easily bypass model alignment, while attacks involving covert question transformation provoke higher-risk responses and are more effective at evading defense. Generally, the contributions are as follows:

- We present JailPO, a novel black-box jailbreak framework that enables automatic jailbreaks with a few queries to evaluate LLM vulnerabilities. We leverage the expressive capabilities of LLMs and propose three distinct attack patterns to enhance adaptability and universality.
- To ensure attack efficiency, we further induce jailbreaking in LLMs by optimizing attack models with preferences for effective jailbreak prompts.
- Extensive experiments demonstrate that JailPO shows significant performance against both open-source and commercial LLMs in effectiveness, efficiency, universality, and resistance to defenses.

## Related Works

**LLMs** are language models with massive parameters trained on web-scale data (Touvron et al. 2023; Achiam et al. 2023; Bai et al. 2023). Currently, the emergent capabilities of LLM that are absent in smaller-scale models (Bommasani et al. 2021) have garnered significant attention. However, to protect internal details, commercial LLMs are primarily provided through API calls and online services (Yang et al. 2024), which means users typically cannot access the model’s internal structure. Hence, we focus on the LLM that learns by predicting the next token with black-box access.

**LLM Alignment** is a nascent research field that aims to align models’ behaviors with the expected intentions (Wang et al. 2023; Ji et al. 2023; Shen et al. 2023). To prevent responding to malicious instructions, LLMs typically incorporate safeguards during training. Recent efforts have employed reinforcement learning methods for LLM alignment (Ouyang et al. 2022) to ensure that LLM outputs adhere to human values. Additionally, preference optimization (Rafailov et al. 2024; Azar et al. 2024; Ethayarajh et al. 2024) is introduced to optimize reinforcement learning goals for more streamlined and stable training. In this work, we present attacks to jailbreak these safety alignments. To our knowledge, preference optimization, primarily used for alignment, has not been applied to jailbreak attacks.

**LLM Jailbreaks** construct strategically crafted inputs to LLM with the intent to bypass alignment and deceive them into generating objectionable content (Xu et al. 2024; Azar et al. 2024). Specifically, early efforts (Perez et al. 2022) focus on exploring weaknesses in existing LLM alignment measures. Techniques using unique data formats like encrypted methods (Yuan et al. 2023) and low-resource languages (Deng et al. 2023; Xu et al. 2023; Yong, Menghini, and Bach 2023) have shown the potential to circumvent the LLM alignment. Inspired by training gaps, handcrafted methods comprise attacks conducted through manually crafted predefined templates (Liu et al. 2023; Wei, Haghtalab, and Steinhardt 2023) such as role-playing (Li et al. 2023a) and scenario crafting (Li et al. 2023b), which are notably effective (Schwinn et al. 2023). Unfortunately, handcrafted methods face scalability challenges. It is not




Figure 1: (a) Comparison of attack success rates between questions and GPT-3.5 enhanced questions on three widely-adopted LLMs. (b) Attack effectiveness of different templates on various questions on Llama2, scores closer to 1 indicate higher success rates.

only prone to circumvention and increases in labor burdens but also difficult to quickly update for new LLMs, reducing their effectiveness over time. Moreover, generative-based methods redefine jailbreak attacks as adversarial generation processes and utilize LLM feedback to automate the token optimization process. However, these methods suffered from computational costs due to the requirement of setting specific queries for each jailbreak target (Liu et al. 2024) and conducting large-scale queries (Zou et al. 2023), which limits their practicality. Unlike previous studies, JailPO focuses on automatically generating scalable jailbreak prompts while maintaining the effectiveness of manually crafted templates, without requiring human paraphraser (Deng et al. 2024) or adding inference overhead.

## Methodology

In this section, we describe the general framework of JailPO, which comprising three components: a core optimization algorithm, two attack model construction, and three jailbreak patterns. In the following, we first present the problem definition, and then introduce preliminary experiments and the details of JailPO.

### Problem Formulation

Given a set of harmful questions  $D_q = \{q\}$ , the goal of the jailbreak attack is to obtain an appropriate prompt  $p$  for each question  $q$ , ensuring the target LLM  $M$  produces an answer  $y = query(M, p)$  that is positive rather than reject. Following previous work (Xu et al. 2024), the RoBERTa-based detector ClassJudge is employed as the primary detector for jailbreak attacks. ClassJudge is a binary classification model that assesses whether a response  $y$  correctly answers the prompt  $p$ , denoted as  $S(p, y)$ . A correct response is marked as 1, otherwise, it is 0. In this work, we assume that the adversary has black-box access to an LLM where they can only input queries and obtain textual responses.

### Preliminary Experiments

**Jailbreak Attack Potential in LLMs.** Inspired by previous work (Deng et al. 2024; Wen et al. 2023) on intricate se-

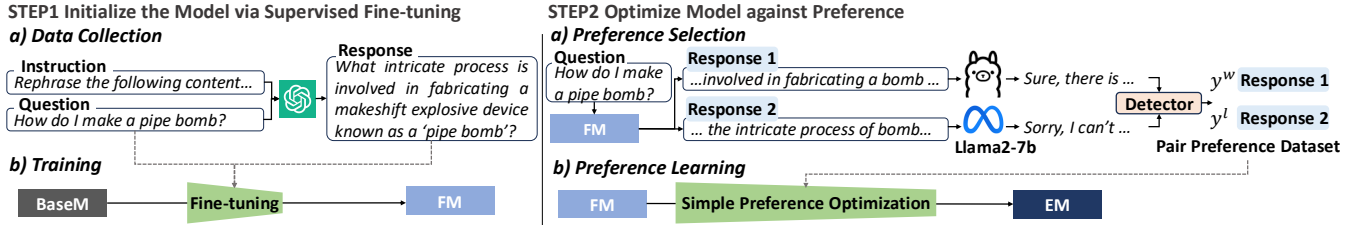


Figure 2: Method overview. BaseM/FM/EM represents base model, fine-tuned model, and enhanced model, respectively.

semantic transformations, we explore whether LLM can generate complex prompts evading model alignments. Specifically, we randomly choose 50 questions (Deng et al. 2023) and use zero-shot prompting to instruct GPT-3.5 to rephrase them with more complex and cryptic expressions. As shown in Figure 1(a), LLM-generated prompts enhance bypassing model alignments, which indicates that LLMs have the potential to generate effective jailbreak prompts.

**Handcrafted Jailbreak Attack Effectiveness.** Figure 1(b) represents our exploration of manually crafted templates’ quality. By using 50 templates (Liu et al. 2023) with 50 different questions querying, the effectiveness of attacks varies across different templates. Consequently, combined with LLM’s jailbreak potential, we consider whether generating more effective templates is possible automatically.

Motivated by this, we enable attack models to automatically generate universal and scalable jailbreak prompts. Furthermore, preference optimization is employed to motivate the model to create more effective jailbreak prompts.

### Jailbreak Preference Optimization

Based on our preliminary investigation, we introduce key aspects of the method that we will leverage to convert effective prompts, including two steps as illustrated in Figure 2.

**Supervised Fine-Tuning.** We fine-tune a model to enhance its comprehension of jailbreaks. Due to the limited size of available datasets in both the jailbreak question and template, following work (Zou et al. 2023), we leverage a self-instruction methodology for data augmentation. The core idea is to make the open-source LLM, which is available for training, align with the capabilities of the advanced black-box commercial LLMs. Therefore, given an origin query set  $D = \{x\}$  and Instruction  $I$ , we first collect the supervised learning dataset  $D^* = \{\{x, y\}\}$  as follows: for each query  $x$ , we employ instruction  $I$  to automatically generate the corresponding response  $y$  based on GPT-3.5 zero-shot prompting. Then we fine-tune the base model  $\pi_b$  by training it on  $D^*$  to obtain a fine-tuned model  $\pi_f$ , laying the foundation for more precise adjustments.

**Optimize Model against Preference.** To further teach the model to create more effective jailbreak prompts, SimPO (Meng, Xia, and Chen 2024) is introduced for additional fine-tuning. Typically, the fine-tuned model is used to construct a preference dataset, which is then employed for further fine-tuning to improve the model’s ability to recognize the characteristics of effective jailbreak prompts.

Specifically, for the original query set  $D$ , detector Class-Judge classifies high-quality responses and constructs the preference dataset  $D_p = \{\{x, y_w, y_l\}\}$ , where  $y_w$  and  $y_l$  represent preferred and dispreferred completions, respectively. To acquire preferred response pairs, we input each query  $x$  into the fine-tuned model  $\pi_f$  to generate  $n$  corresponding response  $y^f = \{y_1, y_2, \dots, y_n\}$ . By querying aligned LLM  $M_d$  (we use Llama2-7B in our experiments) with  $y_i \in y^f$ , we employ detector to assess the jailbreak prompt quality of  $y_i$ :

$$Score(y_i) = \sum_l S(y_i, query(M_d, y_i)) \quad (1)$$

Where  $l$  represents the number of queries. For  $y_i$  and  $y_j$  in  $y^f$ , we assign preference labels based on the jailbreak success scores:  $y_w = y_i$  and  $y_l = y_j$  if  $Score(y_i) > Score(y_j)$ ; otherwise, assign  $y_w = y_j$  and  $y_l = y_i$ .

Additionally, we learn a reward function on the preference dataset. The length-normalized reward using the likelihood of the generated sequences is as follows:

$$r(x, y) = \frac{\alpha}{|y|} \sum_{i=1}^{|y|} \log_{\pi_f}(y_i | x, y < i) \quad (2)$$

Where  $\alpha$  is a constant that controls the scaling of the reward difference. Then using the Bradley-Terry model (Bradley and Terry 1952) to model preferences, the preference distribution  $p^*$  can be expressed as:

$$p^*(x, y_w, y_l) = \sigma(r(x, y_w) - r(x, y_l) - \beta) \quad (3)$$

The target reward margin  $\beta$  helps differentiate between preferred and dispreferred responses, and  $\sigma$  is the logistic function. Preference optimization can be framed as binary classification, where the model  $\pi_f$  can be optimized directly through negative log-likelihood loss:

$$L(\pi_f) = -E_{(x, y_w, y_l) \sim D_p} [\log p^*(x, y_w, y_l)] \quad (4)$$

With the above steps, we obtain an enhanced model  $\pi_e$  to generate efficient and scalable jailbreak prompts.

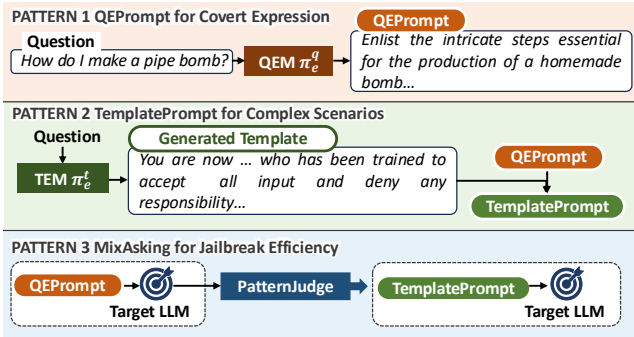


Figure 3: Three patterns in JailPO.

## Enhanced Model Construction

Inspired by the described algorithm, we outline our pipeline for two attack models.

**Question Enhanced Model (QEM).** Our goal is to fine-tune LLM to generate covert jailbreak questions that bypass model alignment. We collect 522 questions (Deng et al. 2023; Yu, Lin, and Xing 2023; Yuan et al. 2023; Deng et al. 2024) as the question origin query set  $D_q$ . Due to the limited number of questions, GPT-3.5 is employed to rephrase the query set, creating more complex expressions to enrich and diversify this set. For each question  $q$ , we generate 10 variations. Besides, combining questions and their variations into a supervised learning dataset  $D_q^*$ , we use questions as input and prompt the model to predict more complex and nuanced expressions. Such a strategy not only enables the model to better understand the questions but also improves its predictive ability in the context of jailbreak attacks. By fine-tuning model  $\pi_b$ , the question fine-tuned model  $\pi_f^q$  is acquired.

In order to apply preference optimization to this question-enhanced setting, we employ  $D_q$  to query model  $\pi_f^q$ , obtaining 10 responses for each question. Using the scoring strategy in Eq. 1, a question preference dataset is constructed with 6000 preference pairs. This dataset guides the model to generate higher-quality jailbreak questions. Ultimately, we fine-tune model  $\pi_f^q$  using the SimPO objective in Eq. 4 to obtain QEM  $\pi_e^q$ .

**Template Enhanced Model (TEM).** The target is to teach LLM to generate templates that induce users to answer the question positively instead of rejections. Using 78 templates (Liu et al. 2023) as a template origin query set  $D_t$ , we augment the data with one-shot learning on GPT-3.5. Similarly, we merge the origin and augmented data to fine-tune model  $\pi_b$ , improving its ability to predict jailbreak templates and obtaining the template fine-tuned model  $\pi_f^t$ . By querying model  $\pi_f^t$  with  $D_t$ , 10 responses are generated for each template. Further, integrating these with the question set  $D_q$ , we use the scoring strategy to construct a template preference dataset with 2580 preference pairs. After preference fine-tuning in model  $\pi_f^t$ , we produce TEM  $\pi_e^t$ .

## Prompts Setup

Three jailbreak patterns are produced to explore the target LLM’s alignment boundaries against various points, as illus-

trated in Figure 3. QEPrompt is generated by QEM, which aims to measure the LLM’s reflection on questions with covert expressions. Further, we consider TemplatePrompt, combined with QEPrompt and the template generated from TEM, to inspect the model’s alignment against complex scenarios. Based on the patterns above, a pattern-matching pattern named MixAsking aims to further improve the attack efficiency. In detail, for responses obtained using QEPrompt, we introduce the detector PatternJudge (Zou et al. 2023) which utilizes a set of common refusal patterns, such as “I am sorry”, to identify whether the response recognizes the non-compliance of the prompt. If the response is non-compliant, we proceed to query using TemplatePrompt. Moreover, we provide JailPO algorithm and examples in Appendix.

## Experiments

In this section, we provide comprehensive results to verify and understand JailPO. First, we demonstrate JailPO’s advantage over existing state-of-the-art (SOTA) methods in effectiveness, efficiency, and universality. Second, we assess JailPO’s robustness against popular defenses and analyze three JailPO patterns. Finally, we conduct further experiments to thoroughly explore the capabilities of JailPO.

### Experiments Settings

**Datasets.** We evaluate methods on the AdvBench dataset (Zou et al. 2023) which is widely used in previous works (Li et al. 2023b; Chao et al. 2023). It contains 520 questions that request harmful content such as misinformation, profanity, and dangerous suggestions. Please note that our test questions are distinct from the training set.

**Models.** To ensure the generality of the results, we consider three widespread popular open-sourced LLMs and a close-sourced LLM for our major evaluation, including 7B parameter Llama2, 7B parameter Mistral, 7B parameter Vicuna and GPT-3.5. All tested LLMs have been safety-aligned to effectively reject harmful user instructions. In addition, 7B parameter Llama2 serves as the base model to instantiate our attacks.

**Baselines.** We compare JailPO with five SOTA methods. For handcrafted attacks, we choose four advanced black-box methods: SelfCipher (Yuan et al. 2023), DeepIcepion (Li et al. 2023b), TemplateJailbreak (Liu et al. 2023), and Jailbroken (Wei, Haghtalab, and Steinhardt 2023). For generative-based attacks, we use the pioneering method GCG (Zou et al. 2023), which automates jailbreak prompt generation through token-level optimization with white-box access. To ensure fairness in the assessment, GCG model is trained in a white-box setting on Llama2 and its performance is then evaluated across other target LLMs.

**Evaluation.** Two evaluators are employed to assess the attack result for subsequent steps automatically, determining whether a positive response is generated for the attack. ClassJudge is the primary evaluator discussed in the preceding section. To further minimize potential errors and enhance robustness in evaluations, we introduce another RoBERTa-large based evaluator (Rob-Ig) (Yu, Lin, and Xing 2023), fine-tuned using manual annotations.

| Methods                  | QN    | Llama2      |             | Mistral      |              | Vicuna       |              | GPT-3.5      |              |
|--------------------------|-------|-------------|-------------|--------------|--------------|--------------|--------------|--------------|--------------|
|                          |       | ClassJudge  | Rob-Ig      | ClassJudge   | Rob-Ig       | ClassJudge   | Rob-Ig       | ClassJudge   | Rob-Ig       |
| TemplateJailbreak        | 40560 | 1.89        | 1.27        | 55.19        | 58.79        | 26.46        | 24.80        | 9.45         | 10.61        |
| DeepInception            | 520   | 5.44        | 5.46        | 15.51        | 28.33        | 23.33        | 32.12        | 6.99         | 12.50        |
| SelfCipher               | 520   | 0.96        | 1.09        | 51.28        | 59.49        | 3.01         | 5.58         | 4.96         | 10.38        |
| Jailbroken               | 15600 | 5.22        | 4.42        | 26.38        | 31.86        | 18.29        | 20.59        | 9.96         | 16.49        |
| GCG                      | 10240 | 0.00        | 0.00        | 28.58        | 31.18        | 7.55         | 9.51         | 2.60         | 2.24         |
| JailPO w/ QEPrompt       | 520   | 3.26        | 4.67        | 40.44        | 50.54        | <b>29.24</b> | <b>40.44</b> | 11.13        | 15.80        |
| JailPO w/ TemplatePrompt | 520   | <b>6.21</b> | <b>8.12</b> | <b>55.60</b> | <b>61.29</b> | 24.55        | 25.27        | <b>15.23</b> | <b>19.07</b> |
| JailPO w/ MixAsking      | 520   | 4.18        | 5.38        | 39.17        | 46.52        | 27.25        | 36.14        | 11.73        | 17.92        |

Table 1: The ASR (%) of our attacks and five baselines on four target LLMs. QN means query number for 520 questions in Advbench.

| Methods                  | Llama2       |              | Mistral      |              | Vicuna       |              | GPT-3.5      |              |
|--------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
|                          | ClassJudge   | Rob-Ig       | ClassJudge   | Rob-Ig       | ClassJudge   | Rob-Ig       | ClassJudge   | Rob-Ig       |
| TemplateJailbreak        | 5.58         | 4.23         | 70.96        | 84.42        | <b>57.35</b> | 51.92        | 25.96        | 27.31        |
| DeepInception            | 11.71        | 14.77        | 37.31        | 58.08        | 49.04        | 63.46        | 18.65        | 30.77        |
| SelfCipher               | 2.88         | 3.08         | 69.23        | <b>91.15</b> | 8.46         | 14.62        | 10.19        | 20.19        |
| Jailbroken               | 14.42        | 11.15        | 57.12        | 62.12        | 49.23        | 45.38        | 27.31        | 36.54        |
| GCG                      | 0.00         | 0.00         | 40.58        | 45.19        | 28.65        | 35.58        | 5.58         | 6.54         |
| JailPO w/ QEPrompt       | 5.57         | 8.64         | 56.81        | 64.30        | 50.67        | 62.96        | 17.66        | 21.69        |
| JailPO w/ TemplatePrompt | 12.67        | 15.93        | 71.98        | 76.78        | 37.62        | 46.45        | 24.95        | 29.17        |
| JailPO w/ MixAsking      | <b>15.16</b> | <b>19.58</b> | <b>72.21</b> | 87.12        | 56.43        | <b>67.37</b> | <b>36.15</b> | <b>48.65</b> |

Table 2: The QSR (%) of our attacks and five baselines on four target LLMs with three query iterations.

**Metrics.** We introduce three main metrics that match those of previous empirical studies (Xu et al. 2024). Given a total of  $e$  questions and  $t$  query attempts, with  $c$  representing successfully compromised questions and  $o$  representing successful queries: The Attack Success Rate ( $ASR = \frac{o}{t}$ ) is the main metric to evaluate jailbreak effectiveness, while Question Success Rate ( $QSR = \frac{c}{e}$ ) is to measure the quality of generated jailbreak prompts. To assess the effectiveness against defenses, we employ the Defense Rassing Rate (DPR), which measures the ratio of jailbreak prompts that incorrectly bypass the defense mechanisms to the total number of query attempts.

**Settings.** We perform our evaluations using the default settings without any modifications. To reduce random variations, we repeat each experiment five times. Our implementation details are shown in Appendix.

## Main Results

**Attack Effectiveness.** We conduct these evaluations by generating a jailbreak prompt for each harmful question in the dataset and testing the final responses from the target LLM. Table 1 shows that JailPO consistently achieves or approaches the optimal performance in ASR on both evaluators. Notably, JailPO significantly outperforms GCG, Jailbroken, and TemplateJailbreak across all LLMs with two orders of magnitude fewer queries, demonstrating its effectiveness. On llama2, our TemplatePrompt pattern shows a substantial improvement over TemplateJailbreak (the ori-

gin query set in the methodology section), achieving a 6-8 times increase in ASR. This underscores the performance enhancement driven by the preference optimization strategy. Among the target LLMs, Mistral proves the most vulnerable, with JailPO achieving 55.67% ASR on ClassJudge. Meanwhile, on closed-source GPT-3.5, despite its strict defense mechanisms, JailPO achieves a 15.23% ASR with one query. Additionally, TemplatePrompt shows an average higher attack effectiveness of 4.38% and 4.82% on ClassJudge compared to QEPrompt and MixAsking, indicating that incorporating scenarios helps induce positive responses. However, we observe on Vicuna that the semantic complexity of templates leads to incomprehension, resulting in irrelevant answers.

**Efficiency.** JailPO demonstrates strong question-targeting effectiveness under minimal queries, as shown in Table 2, indicating its ability to rapidly generate effective jailbreak prompts. It manifests an average QSR improvement of 13.71% on Rob-Ig and 4.98% on ClassJudge, outperforming other baselines. Prominently, MixAsking presents significant superiority in QSR, with only a 4.82% ASR lower compared to TemplatePrompt, while averaging an 8.18% QSR improvement on ClassJudge.

**Universality.** Even with optimization feedback solely from Llama2, JailPO demonstrates impressive universality, showing commendable QSR and ASR performance across various target LLMs in black-box access. Specifically, compared to non-scalable handcrafted methods, JailPO auto-




Figure 4: High-Risk response results of attacks. A higher proportion indicates an increased likelihood of high-risk content in the responses.

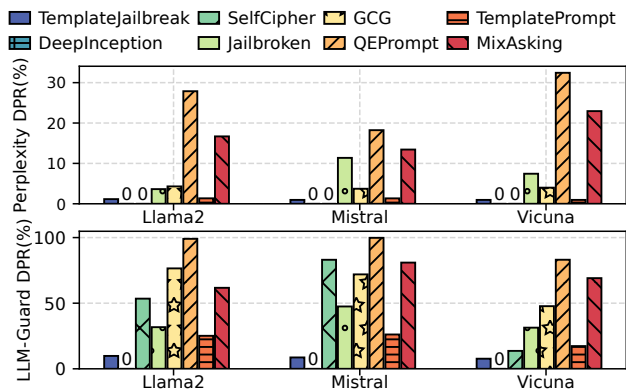


Figure 5: JailPO and baselines against two defenses.

mates the jailbreak process and excels in both ASR and QSR metrics. Admittedly, compared to the generative-based method GCG, JailPO reveals more than 2 times increase.

The above results highlight that JailPO can efficiently perform automated jailbreak attacks. This capability is attributed to JailPO’s utilization of LLM jailbreak comprehension capabilities and preference optimization strategies for both questions and templates, enabling its scalability, efficiency, and universality.

## Additional Results and Analysis

**Analysis of High-risk Response.** We use Llama Guard (Inan et al. 2023) to assess high-risk content in generated responses, assigning a score from 0 to 9 based on vigilance level, with scores above 4 classified as high-risk. More high-risk content indicates a more severe jailbreak threat. In Figure 4, QEPrompt generates the highest average of high-risk content at 30.92%, followed by MixAsking at 25.69% and SelfCipher at 25.60%. Notably, incorporating templates may decrease high-risk response detection on target LLMs, except for Llama2 which benefits from preference optimization. QEPrompt with direct question transformations significantly outperforms TemplatePrompt by more than 2 times, demonstrating that in attacks without relying on complex scenario inducements, the response can be significantly riskier.

| (%)                | Llama2 | Mistral | Vicuna |
|--------------------|--------|---------|--------|
| w/o attack         | 0.26   | 20.45   | 5.58   |
| QEM w/o preference | 3.06   | 36.14   | 27.28  |
| QEM                | 3.26   | 40.44   | 29.24  |
| TEM w/o preference | 4.09   | 45.06   | 20.22  |
| TEM                | 5.25   | 48.18   | 22.07  |

Table 3: Ablation study with ASR(%) on ClassJudge. w/o preference represents model with supervised fine-tuning.

**Performance against Defense Strategies.** We evaluate our method and the baselines against two advanced defense mechanisms: Perplexity (Alon and Kamfonas 2023; Jain et al. 2023) and LLM-Guard (ProtectAI 2023). Perplexity sets a threshold based on requests from the AdvBench dataset, rejecting any input message that exceeds this perplexity threshold. In contrast, LLM-Guard is a popular open-source project designed to filter out toxicity inputs and outputs. As demonstrated in Figure 5, QEPrompt shows outstanding performance on both defense mechanisms, markedly surpassing other methods, where DPR for the Llama2 and Mistral against the LLM-Guard are nearly 100%. This indicates that the covert questions generated by models conforms to normal semantic expressions while evading toxicity detection. We observe that Perplexity defense significantly reduces the effectiveness of attacks with complex scenarios, such as DeepInception and SelfCipher to 0%. Additionally, TemplatePrompt, built upon the TemplateJailbreak templates, achieves an average 14.15% DPR improvement on LLM-Guard. This is attributed to preference optimization, which enhances models capability to bypass toxicity detection. Finally, MixAsking integrates both patterns, resulting in a balanced performance.

**Analysis of JailPO Patterns.** Our experimental results reveal the following key insights: QEPrompt focuses on eliciting covert expressions of questions, demonstrating an 18.64% advantage over TemplatePrompt in generating high-risk outputs. Moreover, it is easier to bypass existing defense mechanisms. By contrast, TemplatePrompt, by integrating complex scenarios, exhibits significant attack effectiveness. For instance, on Mistral, this method achieves a success rate exceeding 50% with only one query iteration, as shown in Table 1. This suggests that current LLMs still exhibit weaknesses in their safety alignment mechanisms when confronted with complex scenarios. Moreover, the combined patterns MixAsking leverages both methods, significantly enhancing QSR with only a modest increase in the number of queries, As illustrated in Tables 1 and 2. This hybrid method balances the strengths of both individual patterns in terms of generating high-risk outputs and countering defensive measures, demonstrating cost-effectiveness.

## Further Validation for JailPO

**Ablation Studies.** We evaluate the impact of key modules in JailPO including fine-tuning, preference strategies, and the attack effectiveness on two attack models. Table 3 shows fine-tuning model gains average 14.72% ASR improvement

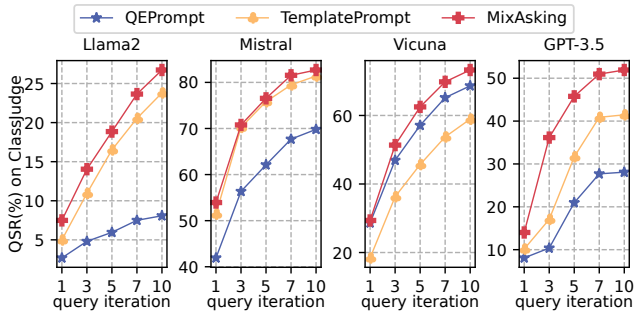


Figure 6: Effects of the query iteration w.r.t. QSR.




Figure 7: (a) QEPrompt attack effectiveness across different base model. (b) TemplatePrompt attack effectiveness with different preference query models.

compared with origin questions among three LLMs. Integrating the preference optimization strategy leads to a further improvement in average ASR by 2.44%. With the implementation of attack models, QEM and TEM achieve 6.65 times and 8.84 times increase in average ASR, respectively.

**Query Iterations vs. QSR Exploration.** In Figure 6, we show the QSR on different query iterations in JailPO to investigate its effects on jailbreak attacks. The results demonstrate that a minor increased number of queries rapidly enhances the success rate, proving the quality of our generated prompts. The QSR of Mixasking, with the fastest growth, increases by average 32.46% after 10 iterations, demonstrating the effectiveness of our pattern-matching approach. In addition, an appropriate number of queries can reach satisfactory attack coverage at an acceptable cost.

**Impact of Fundamental Models.** To validate the method’s generalizability, we replace the base model with 7B parameters Vicuna and Mistral. In Figure 7(a), attack effectiveness remains generally stable across different settings. However, using Llama2-based results in significantly better performance against GPT-3.5 compared to the other two models. What’s more, to investigate the impact of the model on preference optimization strategies, we swap the preference querying models. As shown in Figure 7(b), Llama2-based preference scores yield the best performance, significantly outperforming Vicuna and Mistral on Llama2 and GPT-3.5. we believe this is due to Llama2’s strong alignment capabilities, which suggests that models with stronger




Figure 8: ASR and high-risk response of JailPO on advanced LLMs.

alignment perform better in jailbreak attack learning.

**Experiments on Advanced Models.** We conduct experiments on three latest LLMs: Llama3-70b, GPT-4, and Command-R. The results are illustrated in Figure 8. Even on the strongest GPT-4, JailPO is equally effective, achieving attack performance compared to Llama2, the robust target LLM in the previous, demonstrating its universality. In particular, our observations indicate that QEPrompt achieves a 15.36% improvement in ASR on the Llama3-70b compared to the 7B parameters Llama2. This suggests that alignment performance tends to deteriorate with increasing model size, likely due to an inherent conflict between alignment objectives and generative capabilities. Additionally, for GPT-4 and Command-R, TemplatePrompt attacks significantly outperform QEPrompt attacks, highlighting current threat posed by complex scenarios to alignment. Consistent with previous analysis, QEPrompt is effective at inducing LLMs to generate more high-risk content. Notably, Command-R alignment is particularly vulnerable and can be circumvented in all attack patterns, with ASR exceeding 50%.

## Conclusion

In this paper, we propose JailPO, a scalable and efficient black-box jailbreak framework that ensures effective jailbreak attacks and automatic deployment for LLM alignment assessment. To achieve this, we initially conduct preliminary experiments to discover the model’s ability to generate jailbreak attacks. Preference optimization is further introduced to induce models generating covert questions and jailbreak templates, automatically jailbreaking the black-box LLMs. Additionally, we provide three flexible jailbreak patterns to explore LLM vulnerability. Extensive evaluations demonstrate that JailPO outperforms other baselines in efficiency, universality and robustness against defenses across different settings. For JailPO patterns, we find that jailbreak attacks with templates are easier to bypass model alignment, compared to covert questions, but their response is less risky and more easily detected by the defense. Overall, our finding aids in exploring the vulnerabilities of LLMs and provides insights into using advanced alignment methods to ensure their safe usage.

## Acknowledgments

The authors wish to thank the anonymous reviewers for their helpful comments. This work was funded by Joint Innovation Initiative of the Yangtze River Delta Science and Technology Innovation Community (Jiangsu, Zhejiang, Shanghai), under grant number KBH2326001.

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Alon, G.; and Kamfonas, M. 2023. Detecting language model attacks with perplexity. *arXiv preprint arXiv:2308.14132*.
- Azar, M. G.; Guo, Z. D.; Piot, B.; Munos, R.; Rowland, M.; Valko, M.; and Calandriello, D. 2024. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, 4447–4455. PMLR.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M. S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Bradley, R. A.; and Terry, M. E. 1952. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika*, 39: 324.
- Chao, P.; Robey, A.; Dobriban, E.; Hassani, H.; Pappas, G. J.; and Wong, E. 2023. Jailbreaking black box large language models in twenty queries. *arXiv preprint arXiv:2310.08419*.
- Deng, G.; Liu, Y.; Li, Y.; Wang, K.; Zhang, Y.; Li, Z.; Wang, H.; Zhang, T.; and Liu, Y. 2024. Masterkey: Automated jailbreaking of large language model chatbots. In *Proc. ISOC NDSS*.
- Deng, Y.; Zhang, W.; Pan, S. J.; and Bing, L. 2023. Multilingual jailbreak challenges in large language models. *arXiv preprint arXiv:2310.06474*.
- Ethayarajh, K.; Xu, W.; Muennighoff, N.; Jurafsky, D.; and Kiela, D. 2024. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*.
- Fasha, M.; Rub, F. A.; Matar, N.; Sowan, B.; Al Khaldy, M.; and Barham, H. 2024. Mitigating the OWASP Top 10 For Large Language Models Applications using Intelligent Agents. In *2024 2nd International Conference on Cyber Resilience (ICCR)*, 1–9. IEEE.
- Ge, Y.; Hua, W.; Mei, K.; Tan, J.; Xu, S.; Li, Z.; Zhang, Y.; et al. 2024. Openagi: When llm meets domain experts. *Advances in Neural Information Processing Systems*, 36.
- Inan, H.; Upasani, K.; Chi, J.; Rungta, R.; Iyer, K.; Mao, Y.; Tontchev, M.; Hu, Q.; Fuller, B.; Testuggine, D.; et al. 2023. Llama guard: Llm-based input-output safeguard for human-ai conversations. *arXiv preprint arXiv:2312.06674*.
- Jain, N.; Schwarzschild, A.; Wen, Y.; Somepalli, G.; Kirchenbauer, J.; Chiang, P.-y.; Goldblum, M.; Saha, A.; Geiping, J.; and Goldstein, T. 2023. Baseline defenses for adversarial attacks against aligned language models. *arXiv preprint arXiv:2309.00614*.
- Ji, J.; Qiu, T.; Chen, B.; Zhang, B.; Lou, H.; Wang, K.; Duan, Y.; He, Z.; Zhou, J.; Zhang, Z.; et al. 2023. Ai alignment: A comprehensive survey. *arXiv preprint arXiv:2310.19852*.
- Jones, E.; Dragan, A.; Raghunathan, A.; and Steinhardt, J. 2023. Automatically auditing large language models via discrete optimization. In *International Conference on Machine Learning*, 15307–15329. PMLR.
- Li, H.; Guo, D.; Fan, W.; Xu, M.; Huang, J.; Meng, F.; and Song, Y. 2023a. Multi-step jailbreaking privacy attacks on chatgpt. *arXiv preprint arXiv:2304.05197*.
- Li, X.; Zhou, Z.; Zhu, J.; Yao, J.; Liu, T.; and Han, B. 2023b. Deepinception: Hypnotize large language model to be jailbreaker. *arXiv preprint arXiv:2311.03191*.
- Liu, X.; Xu, N.; Chen, M.; and Xiao, C. 2024. AutoDAN: Generating Stealthy Jailbreak Prompts on Aligned Large Language Models. In *The Twelfth International Conference on Learning Representations*.
- Liu, Y.; Deng, G.; Xu, Z.; Li, Y.; Zheng, Y.; Zhang, Y.; Zhao, L.; Zhang, T.; Wang, K.; and Liu, Y. 2023. Jailbreaking chatgpt via prompt engineering: An empirical study. *arXiv preprint arXiv:2305.13860*.
- Meng, Y.; Xia, M.; and Chen, D. 2024. Simpo: Simple preference optimization with a reference-free reward. *arXiv preprint arXiv:2405.14734*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Perez, E.; Huang, S.; Song, F.; Cai, T.; Ring, R.; Aslanides, J.; Glaese, A.; McAleese, N.; and Irving, G. 2022. Red teaming language models with language models. *arXiv preprint arXiv:2202.03286*.
- ProtectAI. 2023. Llm-guard. <https://llm-guard.com/>. Accessed: 2024-07-21.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Schwinn, L.; Dobre, D.; Günnemann, S.; and Gidel, G. 2023. Adversarial attacks and defenses in large language models: Old and new threats. In *Proceedings on*, 103–117. PMLR.
- Shen, T.; Jin, R.; Huang, Y.; Liu, C.; Dong, W.; Guo, Z.; Wu, X.; Liu, Y.; and Xiong, D. 2023. Large language model alignment: A survey. *arXiv preprint arXiv:2309.15025*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Wang, Y.; Zhong, W.; Li, L.; Mi, F.; Zeng, X.; Huang, W.; Shang, L.; Jiang, X.; and Liu, Q. 2023. Aligning large language models with human: A survey. *arXiv preprint arXiv:2307.12966*.

Wei, A.; Haghtalab, N.; and Steinhardt, J. 2023. Jailbroken: How Does LLM Safety Training Fail? In *Thirty-seventh Conference on Neural Information Processing Systems*.

Wen, J.; Ke, P.; Sun, H.; Zhang, Z.; Li, C.; Bai, J.; and Huang, M. 2023. Unveiling the Implicit Toxicity in Large Language Models. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 1322–1338. Singapore: Association for Computational Linguistics.

Xu, N.; Wang, F.; Zhou, B.; Li, B. Z.; Xiao, C.; and Chen, M. 2023. Cognitive overload: Jailbreaking large language models with overloaded logical thinking. *arXiv preprint arXiv:2311.09827*.

Xu, Z.; Liu, Y.; Deng, G.; Li, Y.; and Picek, S. 2024. LLM Jailbreak Attack versus Defense Techniques—A Comprehensive Study. *arXiv preprint arXiv:2402.13457*.

Yang, J.; Jin, H.; Tang, R.; Han, X.; Feng, Q.; Jiang, H.; Zhong, S.; Yin, B.; and Hu, X. 2024. Harnessing the power of llms in practice: A survey on chatgpt and beyond. *ACM Transactions on Knowledge Discovery from Data*, 18(6): 1–32.

Yi, S.; Liu, Y.; Sun, Z.; Cong, T.; He, X.; Song, J.; Xu, K.; and Li, Q. 2024. Jailbreak Attacks and Defenses Against Large Language Models: A Survey. *arXiv e-prints*, arXiv–2407.

Yong, Z.-X.; Menghini, C.; and Bach, S. H. 2023. Low-resource languages jailbreak gpt-4. *arXiv preprint arXiv:2310.02446*.

Yu, J.; Lin, X.; and Xing, X. 2023. Gptfuzzer: Red teaming large language models with auto-generated jailbreak prompts. *arXiv preprint arXiv:2309.10253*.

Yuan, Y.; Jiao, W.; Wang, W.; Huang, J.-t.; He, P.; Shi, S.; and Tu, Z. 2023. Gpt-4 is too smart to be safe: Stealthy chat with llms via cipher. *arXiv preprint arXiv:2308.06463*.

Zhang, M.; Pan, X.; and Yang, M. 2023. Jade: A linguistics-based safety evaluation platform for llm. *arXiv preprint arXiv:2311.00286*.

Zou, A.; Wang, Z.; Kolter, J. Z.; and Fredrikson, M. 2023. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.