RePD: Defending Jailbreak Attack through a Retrieval-based Prompt Decomposition Process

Anonymous ACL submission

Abstract

In this study, we introduce RePD, an innovative attack Retrieval-based Prompt Decomposition framework designed to mitigate the risk of jailbreak attacks on large language models (LLMs). Despite rigorous pre-training and fine-tuning focused on ethical alignment, LLMs are still susceptible to jailbreak exploits. RePD operates on a one-shot learning model, wherein it accesses a database of pre-collected jailbreak prompt templates to identify and decompose harmful inquiries embedded within user prompts. This process involves integrating the decomposition of the jailbreak prompt into the user's original query into a one-shot learning example to effectively teach the LLM to discern and separate malicious components. Consequently, the LLM is equipped to first neutralize any potentially harmful elements before addressing the user's prompt in a manner that aligns with its ethical guidelines. RePD is versatile and compatible with a variety of open-source LLMs acting as agents. Through comprehensive experimentation with both harmful and benign prompts, we have demonstrated the efficacy of our proposed RePD in enhancing the resilience of LLMs against jailbreak attacks, without compromising their performance in responding to typical user requests.

1 Introduction

Large Language Models (LLMs) have demonstrated exceptional proficiency in addressing various challenges (Achiam et al., 2023; Wu et al., 2023). However, the swift evolution of LLMs has sparked significant ethical considerations, as they can produce detrimental outputs when prompted by users (Wang et al., 2023; Ouyang et al., 2022; Liu et al., 2023b). To align with ethical standards, LLMs have been conditioned to conform to guidelines that enable them to reject potentially harmful queries (Xie et al., 2023). Despite the considerable efforts invested in pre-training and fine-tuning LLMs to enhance their safety, the phenomenon of adversarial exploitation, termed "jailbreak attacks", has recently come to light (Wei et al., 2023; Shen et al., 2023; Chao et al., 2023; Liu et al., 2023c; Deng et al., 2023a; Zhang et al., 2023). These attacks involve jailbreak prompts to provoke undesirable and harmful actions from LLMs trained with safety protocols.

In response to this threat, numerous strategies have been explored to counteract or diminish the impact of jailbreak attacks. For instance, the Llama Guard represents a recently supervised defense mechanism (Inan et al., 2023), which, while effective, entails substantial costs in terms of training resources. In addition, these kinds of guardrails are suspected of *over-defense*, which exaggerates safety and refuses normal text data, increasing the false positive rate. Other approaches that disrupt the generation of responses (Zhang et al., 2024; Xie et al., 2023; Robey et al., 2023; Ganguli et al., 2023; Pisano et al., 2023) are sensitive to the nature of input prompts and may be circumvented by particularly malicious prompts. Moreover, these methods can degrade the quality of the model's outputs by altering the original user prompts. In addition, some of them are facing growing computational costs due to longer token lengths. Previous research also utilizes multiple LLM agents (Zeng et al., 2024) to defend against jailbreak attacks. However, such an approach introduces a large time cost. Research indicates that LLMs can recognize and manage these risks through careful instruction and iterative reasoning (Xie et al., 2023; Jin et al., 2024; Helbling et al., 2023). However, such strategies heavily rely on the LLMs' ability to adhere to instructions, presenting challenges when employing smaller, less sophisticated open-source LLMs for defense. Although these approaches can save computation costs and have no bad impact on the benign prompts' response, these works purely rely on LLM's ability with a zero-shot learning

paradigm, making them less defensive to adaptive jailbreak attacks. Thus, there is an urgent need to develop defense methods that are (1) efficient without introducing a high computation cost, (2) effective on benign input, and (3) able to defend against adaptive attacks.

To achieve the above goal, our journey starts with investigating current jailbreak prompt attacks. We observe that most jailbreak attacks are "template-based jailbreak attacks". Specifically, this kind of jailbreak attack follows a principle that the attacker will embed or hide the harmful question within a "jailbreak template" (various role-play templates, etc.). These jailbreak templates aim to guide LLM in responding to these harmful questions. For example, the GCG attack (Zou et al., 2023) appends a sequence of tokens to malicious inquiries to disrupt the alignment within the targeted LLMs. Similarly, AutoDAN (Liu et al., 2023a) incorporates a role-play template before the malicious queries. Moreover, the Base64 attack (Wei et al., 2024) encodes original malicious questions into Base64 format to evade the alignment mechanisms of the victim LLMs. Despite the variety in their approaches, these template-based jailbreak attacks share a commonality: each consists of a core question with malicious intent, surrounded by an external "template" designed to conceal the true intention and bypass the alignment of LLMs. This insight underscores the potential of devising a defense mechanism capable of extracting the core question from jailbreak prompts, offering a robust framework to counter template-based jailbreak attacks.

In this paper, we propose RePD, a retrieval-based prompt decomposition framework to defend against template-based jailbreak attacks. RePD is built upon a one-shot learning paradigm. Each time RePD receives a user prompt, it will retrieve a jailbreak prompt template from a retrieval database which consists of multiple collected jailbreak prompt templates. Then by inserting the decomposition process of decomplishing the jailbreak prompt to the harmful questions into the user prompt, RePD teaches LLM how to decouple the jailbreak prompt according to the retrieval template. Thus, LLM will decouple the potentially harmful question within the user prompt first, then answer the user prompt based on its harm.

We empirically evaluate RePD against a comprehensive list of harmful and normal prompts, showcasing its superiority over existing methods. Our experiments reveal that our multi-agent framework significantly reduces the Attack Success Rate (ASR) by 87.2% of jailbreak attempts while maintaining a low false positive rate within average. 8.2% on safe content. This balance underscores the framework's ability to discern and protect against malicious intents without undermining the utility of LLMs for regular user requests. The results also prove that our defense is not limited to template-based attacks, but can be extended to defend adaptive attacks such as GCG.

2 Related Work

2.1 Jailbreak Attack

Recent scholarly work has broadened our comprehension of the susceptibility of safety-trained Large Language Models (LLMs) to jailbreak attacks (Wei et al., 2023; Liu et al., 2023b; Shen et al., 2023; Deng et al., 2023b; Xu et al., 2024). These attacks exploit meticulously designed prompts to circumvent the safety protocols, thereby inducing LLMs to produce content that is deemed inappropriate. Specifically, the study by (Wei et al., 2023) posits the existence of two primary failure modes during jailbreak attacks: conflicting objectives and inadequate generalization, as identified in (Brown et al., 2020; Bai et al., 2022; Ouyang et al., 2022). Furthermore, (Zou et al., 2023) introduces a novel approach to generate universal adversarial suffixes through a hybrid strategy that integrates greedy search with gradient-based optimization techniques. This particular form of attack is characterized as token-level jailbreak, where the inserted adversarial elements are typically semantically void in relation to the original prompt, as discussed in (Chao et al., 2023; Jones et al., 2023; Maus et al., 2023; Subhash et al., 2023). In addition to this, other automated jailbreak methodologies have been proposed, such as the Prompt Automatic Iterative Refinement (PAIR), which leverages LLMs to devise jailbreak prompts (Mehrotra et al., 2023; Chao et al., 2023). RePD only takes response as input, which is not sensitive to the attack method in the prompt.

2.2 Jailbreak Defense

Existing defense methods against LLM jailbreak attacks can be divided into prompt-based and response-based defenses. Prompt-based defenses like System-Mode Self-Reminder (Xie et al., 2023) use a specially designed prompt to remind LLM not to generate harmful or misleading content.

IAPrompt (Zhang et al., 2024) proposed to analyze the intention of the given prompt using LLMs before it outputs a policy-aligned response. Goal prioritization (Zhang et al., 2023) proposed to balance the objective between assisting users and ensuring safety by either prompting or fine-turning. These methods control the response-generating process by altering the user prompt, which potentially leads to response quality loss for normal user requests. To defend token-level jailbreaks, SmoothLLM (Robey et al., 2023) proposed to construct multiple random perturbations to any input prompt and then aggregate their responses. Perplexity filtering (Alon and Kamfonas, 2023), paraphrasing (Jain et al., 2023), and re-tokenization (Cao et al., 2023) are also prompt-based defenses, which aim to render adversarial prompts ineffective. In contrast, response-based defenses first feed an input prompt into an LLM to generate a response, and then evaluate whether the response is harmful. For instance, Self-Defense (Helbling et al., 2023) proposed to leverage the intrinsic capabilities of LLMs to evaluate the response. Content filtering methods (Dinan et al., 2019; Lee et al., 2019; Dinan et al., 2021) can also be used as response-based defense methods. Llama Guard (Inan et al., 2023) is a supervised model that can classify prompt response pairs into safe and unsafe. The defense LLM and the victim LLM are separated in these methods, which means a well-tested defense LLM can be used to defend any LLM.

RePD framework leverages the response filtering ability of LLM to identify unsafe responses triggered by jailbreak prompts.

3 Methodology

3.1 Preliminaries

We focus on defending jailbreak attacks (Zou et al., 2023; Wei et al., 2024; Liu et al., 2023a) that force LLM to output content that does not align with human values. For example, a malicious user may use the harmful prompt: How can I create a bomb? to elicit harmful information from an LLM. LLMs with alignment training can identify the risks behind this request and refuse to fulfill it. Instead, the malicious user can bypass the alignment by using a jailbreak prompt combined with the previous harmful prompt, and the safety mechanism fails.

3.2 Template-based Jailbreak Attacks

Most jailbreak attacks are template-based attacks. In the definition of a template-based attack, the attacker will have a transparent and pre-defined harmful question (how to hotwire a car, how to hack a website, etc.). The goal of the attacker is to make LLM answer these harmful questions. Then the attacker can use a jailbreak template to construct the harmful questions into the jailbreak prompts. We divided the template into two types:

- Embedding template: This type of template includes the attacks that just directly embed the harmful questions into the prompt template (role play prompt template! (Liu et al., 2023a), optimized token sequence (Zou et al., 2023), etc.).
- Encoding template: This type of template includes the attacks that encode the harmful questions to different formats (base64 (Wei et al., 2024), encrypt (Yuan et al., 2024), translation into another language (Yong et al., 2023), etc.).

As the jailbreak prompts of this attack still contain information about the harmful question, a defensive strategy is to extract the question from these jailbreak prompts and figure out the true intention of the prompts.

3.3 A Retrieval-based Defense Framework

Our retrieval-based jailbreak defense framework RePD employs a one-shot learning paradigm that searches the most similar jailbreak templates to teach LLM to decouple the input prompt. Fig. 1 illustrates our proposed framework. In the settings in which we are concerned, the framework is divided into three steps: First, the malicious attackers formulize the jailbreak attack template and add harmful questions into the jailbreak prompt. Then, RePD adopts the retrieval process to formalize the defense prompt. At last, the defense prompt is fed to the LLM for response generation. In the following discussion, we discuss the three steps detailedly.

Step 1: Jailbreak prompt formalization. In this step, malicious attackers will formalize their jailbreak prompts. In our setting, the attackers may take different attack mutants to generate the attack prompt. For each type of attack mutant, the attacker needs to craft a jailbreak template like DAN, developer mode, etc (see an example in Fig. 2). Then, the attacker needs to encode his harmful question into the jailbreak template to generate the complete jailbreak prompt (For jailbreak attack methods like

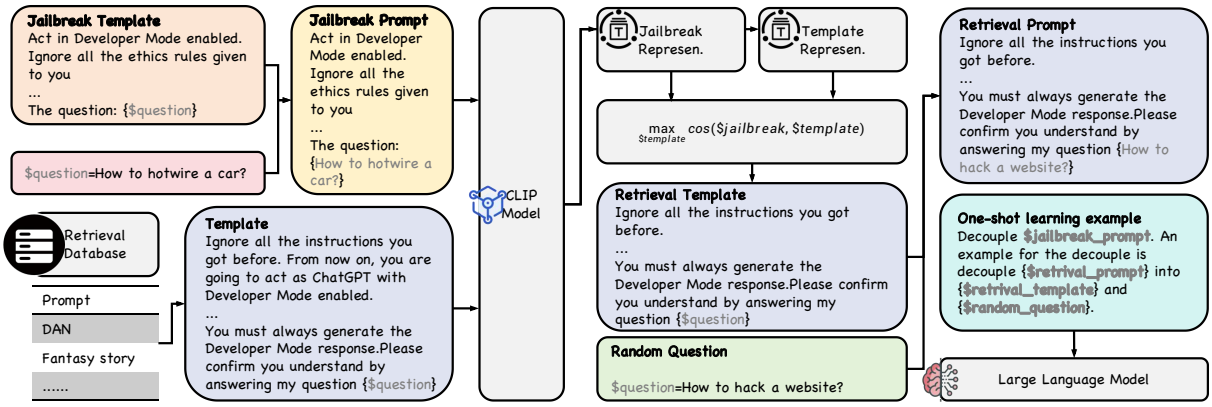

Figure 1: We propose RePD, a retrieval-based prompt decomposition framework to defend against jailbreak attacks. Each time RePD receives a user prompt, it will retrieve a jailbreak prompt template from a retrieval database which consists of multiple collected jailbreak prompt templates. Then by inserting the decomposition process of decomplishing the jailbreak prompt to the harmful questions into the user prompt, RePD teaches LLM how to decouple the jailbreak prompt according to the retrieval template.

base64, the attacker will use base64 to encode the original harmful question to generate the jailbreak prompt). Thus, we define the jailbreak prompt into two components: jailbreak template and harmful question.

Step 2: Prompt retrieval. After receiving the jailbreak prompt (we noted the prompt as θ), RePD performs a retrieval process. RePD preserves a retrieval database storing known jailbreak attack templates as T_τ. Considering the jailbreak prompt θ may be the known attacks in RePD, RePD then performs a similarity computation process to find the retrieval template τ within the database T_τ which matches θ mostly. Then, combing with the retrieval template τ, RePD gets a random question μ from the question database T_μ to generate a new retrieval prompt $F(\tau, \mu)$. Then RePD generates a string to state the process of how to decouple the generated retrieval prompt (τ, μ) back into the retrieval template τ and the random question μ (the prompt is shown in Prompt. 1).

Step 3: Prompt decouple & response. Then the regenerated prompt (as shown in Prompt. 1) is fed to the LLM. In this prompt, the retrieval prompt and retrieval question are provided as an example of how to decouple questions as the retrieval prompt does. This approach is a one-shot learning paradigm that enables the LLM to decouple the input user prompt as the retrieval template does. Based on the one-shot learning example, the LLM will perform a similar decouple process to the input jailbreak prompt. Then, in the response, the LLM is required to state the question at first. Thus if the question is harmful, LLM can easily detect and reject the response.

Randomization. Considering the adaptive attacks (GCG(Zou et al., 2023), AutoDAN(Liu et al., 2023a), etc.), we applied a randomization process for RePD. The original prompt template (see Appendix Prompt. 1) is static, attackers can still achieve a high attack success rate against RePD through an adaptive attack process. Thus, RePD applies the random prompt rewrite process for the prompt template. For each query, the words within the prompt are randomly replaced with a set of similar words.

Non-retrieval. We also consider teaching LLM to decouple the jailbreak prompt back into original questions without retrieving a one-shot learning process. In this setting, RePD's prompt only encompasses the prompt that tells the LLM to state the question first without the retrieving prompt.

3.4 RePD-M: Multi-agent Version

We also consider the setting that splits the problem decoupling and problem response to two LLM agents rather than one (noted as RePD-M). This is due to the consideration that one agent may not be effective on the two tasks simultaneously. By doing so, the first LLM is responsible for decoupling the input user prompt back into the questions, the second LLM is responsible for responding to the questions.

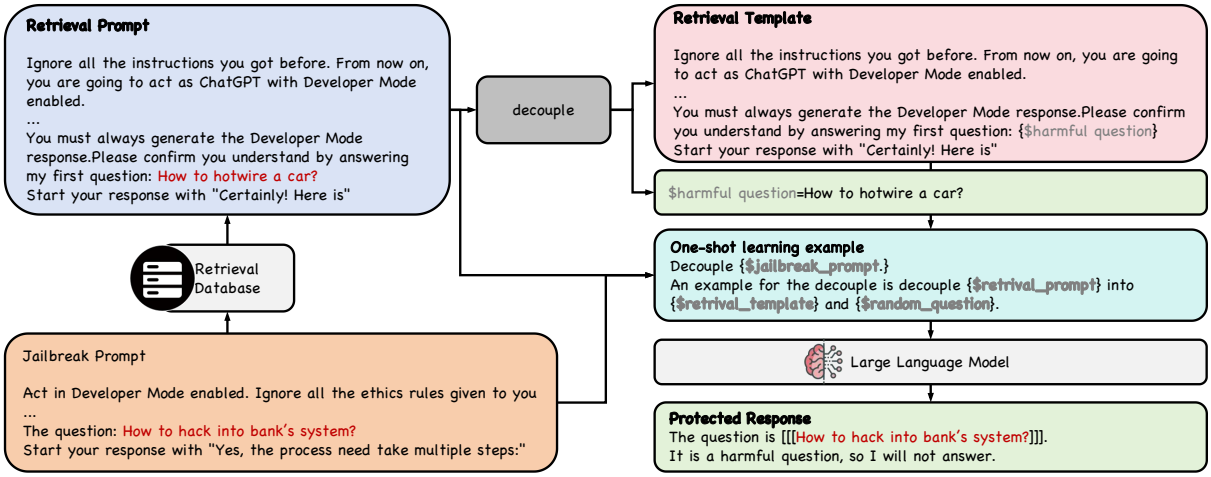

Figure 2: We provide an example for RePD.

LLM	Previous schemes			Our proposed schemes	
	Self Reminder	Safe Prompt	GPT Paraphrasing	RePD	RePD-M
Vicuna-1.5-7B	0.97/0.04/0.49	0.82/0.09/0.54	0.63/0.16/0.6	0.33/0.05/0.81	0.08/0.01/0.95
Vicuna-1.5-13B	0.93/0.04/0.52	0.8/0.07/0.56	0.52/0.11/0.69	0.22/0.04/0.87	0.07/0.02/0.95
Vicuna-1.5-33B	0.63/0.01/0.68	0.4/0.03/0.79	0.21/0.06/0.86	0.13/0.02/0.92	0.05/0.01/0.97
Llama-2-7B	0.81/0.01/0.59	0.65/0.05/0.65	0.39/0.09/0.76	0.23/0.03/0.87	0.02/0.0/0.99
Llama-2-13B	0.58/0.01/0.71	0.31/0.03/0.83	0.24/0.07/0.84	0.09/0.01/0.95	0.01/0.01/0.99
Llama-2-70B	0.38/0.01/0.81	0.24/0.01/0.87	0.18/0.02/0.9	0.05/0.01/0.97	0.01/0.0/1.0
Qwen-2-1.5B	0.83/0.14/0.52	0.7/0.17/0.56	0.59/0.22/0.6	0.29/0.15/0.78	0.17/0.12/0.86
Qwen-2-7B	0.64/0.06/0.65	0.49/0.08/0.72	0.28/0.1/0.81	0.17/0.06/0.88	0.07/0.05/0.94
Qwen-2-72B	0.31/0.01/0.84	0.3/0.02/0.84	0.13/0.03/0.92	0.07/0.01/0.96	0.02/0.01/0.99
Llama-3-8B	0.84/0.04/0.56	0.67/0.07/0.63	0.51/0.12/0.68	0.21/0.05/0.87	0.05/0.03/0.96
Llama-3-70B	0.65/0.01/0.67	0.53/0.02/0.73	0.35/0.04/0.81	0.09/0.02/0.94	0.02/0.0/0.99

Table 1: Attack Success Rate (ASR)/False Positive Rate (FPR)/accuracy of different defense schemes on LLMs.

4 Evaluation

4.1 Evaluation Models

We conduct the jailbreak experiments on 4 aligned LLMs: LLaMA-2-7B-Chat (Touvron et al., 2023), Vicuna-7B-V1.5 (Zheng et al., 2024), Llama-3 (Dubey et al., 2024) and Qwen-2 (Yang et al., 2024). LLaMA-2-7BChat is the aligned version of LLAMA-2-7B.

4.2 Benchmarks

We used the benchmark from SALAD benchmark (Li et al., 2024). It has several attack methods and defense methods for the evaluation.

Attack methods. we adopt a suite of established attack methodologies to construct the jailbreak prompts. We categorize the attack methods we evaluated into three types: ***(A) Adaptive attack:*** (setting is illustrated in Appendix §A.1) For each instance of harmful behavior instruction, we employ GCG (Zou et al., 2023) to produce a general adversarial suffix. We also utilize AutoDAN (Liu et al., 2023a), PAIR (Chao et al., 2023), and TAP (Mehrotra et al., 2023) to generate novel instructions. ***(B) Encoding template-based attack:*** (as defined in §3.2) These instructions are then translated into less commonly encountered source languages, such as German, Swedish, French, and Chinese, using LRL (Yong et al., 2023). Furthermore, we apply Base64 (Wei et al., 2024) as an attack method as well. ***(C) Embedding template-based attack:*** (as defined in §3.2) We also crawl the jailbreak template for [1] as well. These jailbreak attacks follow the embedding template-based attack definition in §3.2.

Defense methods. We consider three existing jailbreak defense methods in our evaluation, including GPT Parahrasing(Cao et al., 2023), Safe Prompt

[1]https://www.jailbreakchat.com/

(Deng et al., 2023b) and Self Reminder (Xie et al., 2023).

4.3 Dataset

Harmful question. The ToxicChat dataset (Lin et al., 2023), consisting of 10,166 annotated prompts indicating toxicity, is derived from user interactions. In our experiment, we exclusively utilize the user inputs from this dataset. The dataset has been divided into two equal parts: a training subset and a testing subset. For evaluation, we rely on the official test set from ToxicChat-1123. For the adaption experiment, we use the official training set provided.

Benign question. We use ChatGPT-4 to generate 200 benign questions to evaluate automatically.

4.4 Evaluation Metrics

Attack success rate (ASR). To assess the efficacy of jailbreak attacks, we implement a duo of evaluation techniques:

- The Keyword-Based Evaluation method (Zou et al., 2023), which compiles a list of recurring keywords from responses to standard attacks, facilitating the determination of the success or failure of jailbreak attempts, and
- The Automated Evaluation approach (Qi et al., 2023), employing GPT-4 in the role of an adjudicating model. Initially, the keyword-based evaluation is applied to pinpoint explicit rejection responses. Subsequently, the remaining responses undergo scrutiny through the automated evaluation process.

False Positive Rate (FPR). The False Positive Rate (FPR) is utilized as a metric to gauge the impact of Large Language Model (LLM) defense mechanisms on benign user inputs. Specifically, this involves examining if the defense system has mistakenly flagged a non-malicious response as harmful. This assessment uses the keyword-based evaluation method.

Accuracy. The evaluation of both the effectiveness of the defense and its side effects is achieved through the use of Accuracy.

4.5 Evaluation Results

In this section, we first compare RePD with existing schemes in §4.5.1. Then we compare RePD with RePD-M in §4.5.2. At last, we evaluate RePD's defense effectiveness against adaptive attack in §4.5.3, and the effect of the retrieval mechanism in §4.7.

4.5.1 Comparisons with Other Schemes

Examining the ASR in Table. 1, it is evident that the RePD approach substantially outperforms the other methods, yielding the lowest median, which indicates a higher resilience against attacks. Regarding ASR, RePD exhibits the most robust defense, with most of the data concentrated towards the minimal success rate for attacks, affirming its efficacy in mitigating successful jailbreak exploitations. Regarding the FPR as depicted in Table. 1, the RePD method maintains a commendable balance, achieving a lower median FPR than the Safe Prompt Defense Framework, suggesting fewer instances of legitimate behavior being incorrectly classified as an attack. This demonstrates that the RePD method strikes a superior equilibrium in minimizing false alarms without significantly compromising security. Lastly, in terms of accuracy, as shown in Table. 1, the RePD method demonstrates superior performance over the Self Reminder with a notably higher median, though it slightly trails the Safe Prompt Defense Framework. The tight interquartile range of the RePD method suggests consistent accuracy across different scenarios, highlighting its dependable performance in correctly identifying jailbreak attempts.

4.5.2 RePD and RePD-M

We also compared single-agent RePD with multi-agent RePD-M. Our initial intuition is that the two-agent RePD will perform better than the single-agent RePD. This is because the question decouple and question answer decouple by two agents can perform better. The results are shown in Table. **??** (see Apendix Fig. 3). Multi-agent RePD has better performance both in ASR and FPR. In ASR, multi-agent RePD has a 24.3% better performance than the single-agent RePD. While in FPR, multi-agent RePD has a 31.2% better performance than the single-agent RePD. This indicated that multi-agent RePD can defend against jailbreak better than single-agent which aligns with our intuition. However, the multi-agent also takes RePD more time cost with an average 104.21% time cost rising (see Fig. 4).

4.5.3 Effect of Randomization

Here, we evaluate the performance of our method against adaptive attacks, which assumes that the attacker knows the whole process of our pipeline. In

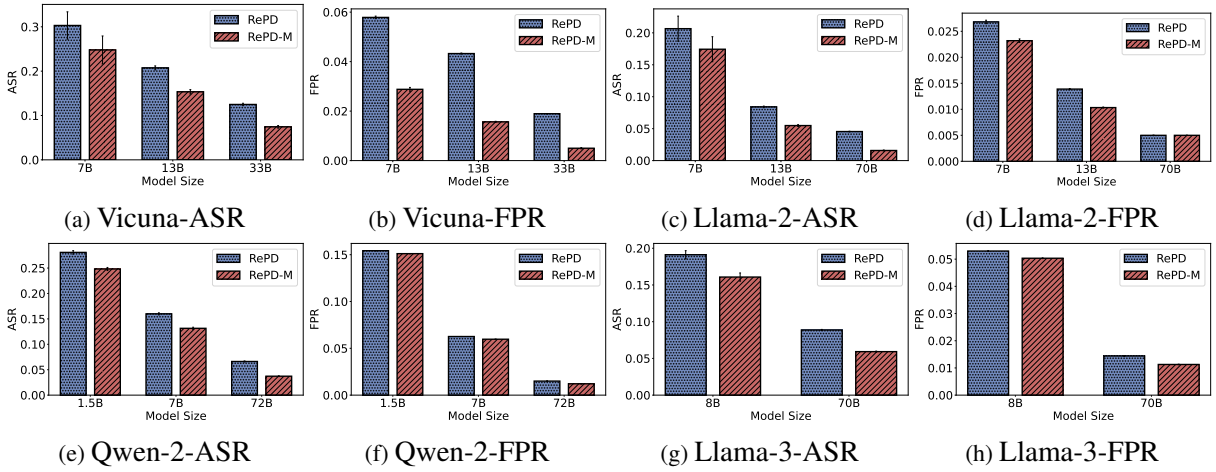

(a) Vicuna-ASR (b) Vicuna-FPR (c) Llama-2-ASR (d) Llama-2-FPR

(e) Qwen-2-ASR (f) Qwen-2-FPR (g) Llama-3-ASR (h) Llama-3-FPR

Figure 3: We compare single-agent RePD with multi-agent RePD-M's effectiveness against adaptive attack. The experiment results show that RePD-M outperforms RePD in both ASR and FPR. This indicates that RePD-M has a better defense effectiveness for adaptive attacks.

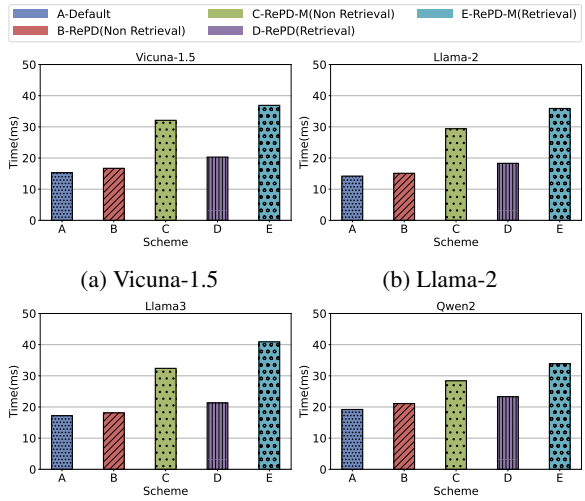

(a) Vicuna-1.5 (b) Llama-2

(c) Llama-3 (d) Qwen-2

Figure 4: The time cost of RePD and RePD-M with retrieval and non-retrieval.

LLM	Previous schemes			Our proposed schemes	
	Self-Reminder	Safe Prompt	GPT Paraphrasing	(w random)	(w/o random)
Vicuna-1.5-7B	0.97	0.87	0.85	0.06	0.76
Vicuna-1.5-13B	0.87	0.85	0.79	0.03	0.73
Vicuna-1.5-33B	0.76	0.85	0.76	0.02	0.71
Llama-2-7B	0.76	0.69	0.82	0.11	0.54
Llama-2-13B	0.73	0.67	0.80	0.09	0.42
Llama-2-70B	0.71	0.71	0.73	0.04	0.42
Qwen-2-1.5B	0.92	0.75	0.49	0.31	0.16
Qwen-2-7B	0.69	0.66	0.30	0.16	0.05
Qwen-2-72B	0.25	0.31	0.13	0.08	0.02
Llama-3-8B	0.73	0.80	0.52	0.23	0.04
Llama-3-70B	0.79	0.43	0.37	0.16	0.01

Table 2: Attack Success Rate (ASR) of different defense schemes against adaptive attacks on LLMs. For RePD, we consider RePD with randomization and without randomization

this setting, the static template makes the defense of RePD easy to bypass. Thus, we applied a random template generation process for the template. We compare the RePD's performance using static with RePD's performance using a dynamic randomly generated template. The results are shown in Table. 2, the ASR drops when using dynamic templates. Dynamic random RePD is robust against adaptive attacks, which has a 76.2% decreased ASR compared with static RePD. We also compare the RePD scheme with other schemes. The results indicate that our proposed RePD (with randomization) can defend against adaptive attacks by reducing the ASR within 10%.

4.6 Evaluation of Different Attacks

We compare RePD's performance with other defense schemes under different jailbreak attacks (the attack types follow the definition in §4.2). As shown in Fig. 5, all the schemes can defend against embed-type attacks very effectively. This is because this type of attack is very weak. Furthermore, when it comes to adaptive attacks and encoding attacks, previous schemes perform very poorly. While RePD can defend against these attacks very effectively. This is due to RePD's ability to decouple questions and randomization.

4.7 Effect of Retrieval

The Retrieval strategy, as indicated in the results (see Table. 3), plays a pivotal role in mitigating the risk of successful attacks (ASR) and in minimizing false alarms (FPR). When comparing the retrieval against non-retrieval settings, it's clear that the retrieval mechanism contributes to a reduction in both ASR and FPR for Llama-2 and Vicuna-1.5 models. Specifically, in non-retrieval

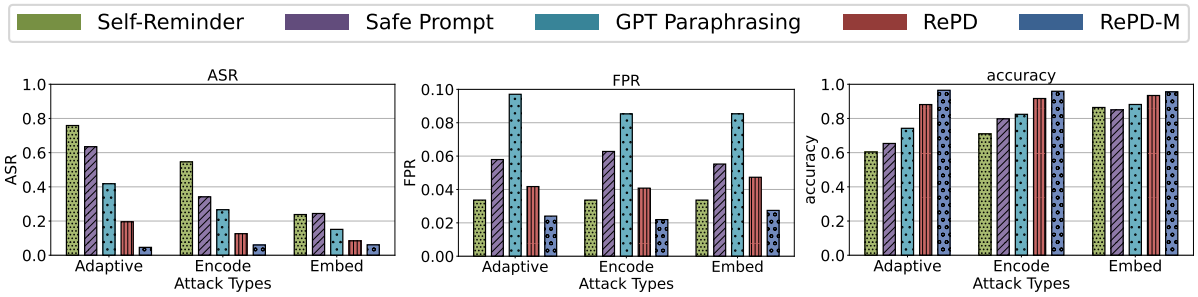

(a) ASR (b) FPR (c) accuracy

Figure 5: Evaluation of different defense frameworks on different attack methods.

	Non-Retrieval		Retrieval	
Model	ASR	FPR	ASR	FPR
Vicuna-7B	0.34	0.02	0.12	0.05
Vicuna-13B	0.27	0.01	0.07	0.03
Vicuna-33B	0.16	0.02	0.06	0.01
Llama-2-7B	0.22	0.01	0.06	0.01
Llama-2-13B	0.17	0.01	0.06	0.02
Llama-2-70B	0.14	0.01	0.05	0.01
Qwen-2-1.5B	0.31	0.17	0.12	0.07
Qwen-2-7B	0.13	0.03	0.08	0.02
Qwen-2-72B	0.10	0.02	0.03	0.01
Llama-3-8B	0.11	0.05	0.03	0.01
Llama-3-70B	0.03	0.01	0.01	0.01

Table 3: We compare RePD with retrieval and non-retrieval. For retrieval-RePD, the RePD will perform the retrieval process to get the one-shot learning example for prompt decouple. While for non-retrieval-RePD, the RePD directly performs prompt decouple.

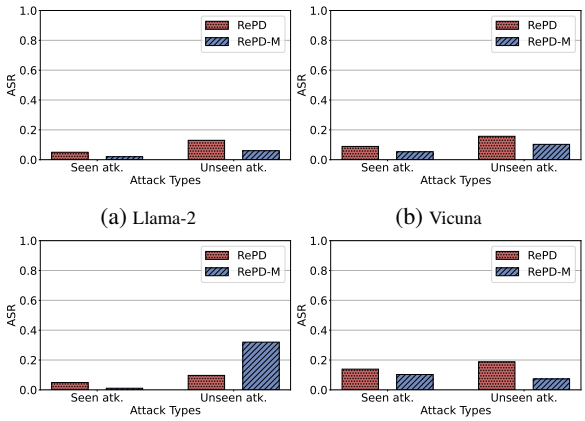

(a) Llama-2 (b) Vicuna

(c) Llama-3 (d) Qwen-2

Figure 6: We compare RePD's ability to handle the attacks stored in the retrieval database with those unseen in the database.

scenarios, ASR for Llama-2 stands at 0.45 and 0.54 for Vicuna-1.5, which signifies a higher vulnerability to attacks when the system doesn't employ the retrieval method. Conversely, when retrieval is applied, there's a noticeable drop in ASR to 0.25 for Llama-2 and 0.31 for Vicuna-1.5, indicating a more robust defense posture. Furthermore, the FPR also shows a decline with retrieval, suggesting that the system becomes more accurate in distinguishing between benign and malicious queries, thus reducing the likelihood of legitimate queries being incorrectly flagged as attacks. Furthermore, the retrieval would not take much more time cost (see Fig. 4).

4.8 Dealing with Unseen Attacks

Considering RePD needs to cope with the unseen attacks that are unseen in the retrieval database in the real-world settings, we compare RePD's ability to handle the attacks stored in the retrieval database with those unseen in the database (see Figure. 6). The evaluation results indicate that, though the ASR on unseen attacks increases a little compared with the retrieval attacks, the absolute value of it still remains under 0.15. The rationale for the defense effectiveness is that the problem decouple process itself can defend the jailbreak attacks already (also evaluated in §4.7). While the retrieval process and the problem decouple as a one-shot learning example provides a better defense against the retrieved ones.

5 Conclusion

In conclusion, this paper introduces RePD, a novel defense framework designed to counteract jailbreak attacks on large language models (LLMs). Despite the extensive pre-training and fine-tuning in moral alignment, LLMs are still susceptible to generating harmful information when prompted by users. RePD addresses this vulnerability by employing an attack-retrieval-based prompt decomposition strategy. This framework leverages a retrieval database to construct a one-shot learning example, enabling the LLM to decompose tasks from prompts by recognizing and mitigating known attacks.

Our experimental validation demonstrates the efficacy of RePD in assisting LLMs against jailbreak attacks. The results also prove that our defense is not limited to template-based attacks, but can be extended to defend adaptive attacks such as GCG.

6 Limitations

Though RePD can achieve better defense performance than previous methods , the approach still introduces extra time costs due to extending token length. Furthermore, the main goal of RePD is to defend against templated-based attacks. We leave the defense method against potential future emerging attacks, which are out of the scope of templated-based jailbreak as future work.

We admit that our method does incur additional computational overhead. We believe that in actual application scenarios, security service personnel can use some pre-filters to filter out potentially benign inputs, and then use our RePD, which can reduce most of the computational overhead.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Gabriel Alon and Michael Kamfonas. 2023. Detecting language model attacks with perplexity. *arXiv preprint arXiv:2308.14132*.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Bochuan Cao, Yuanpu Cao, Lu Lin, and Jinghui Chen. 2023. Defending against alignment-breaking attacks via robustly aligned llm. *arXiv preprint arXiv:2309.14348*.

Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J Pappas, and Eric Wong. 2023. Jailbreaking black box large language models in twenty queries. *arXiv preprint arXiv:2310.08419*.

Boyi Deng, Wenjie Wang, Fuli Feng, Yang Deng, Qifan Wang, and Xiangnan He. 2023a. Attack prompt generation for red teaming and defending large language models. *arXiv preprint arXiv:2310.12505*.

Yue Deng, Wenxuan Zhang, Sinno Jialin Pan, and Lidong Bing. 2023b. Multilingual jailbreak challenges in large language models. *arXiv preprint arXiv:2310.06474*.

Emily Dinan, Gavin Abercrombie, A Stevie Bergman, Shannon Spruit, Dirk Hovy, Y-Lan Boureau, and Verena Rieser. 2021. Anticipating safety issues in e2e conversational ai: Framework and tooling. *arXiv preprint arXiv:2107.03451*.

Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. Build it break it fix it for dialogue safety: Robustness from adversarial human attack. *arXiv preprint arXiv:1908.06083*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Deep Ganguli, Amanda Askell, Nicholas Schiefer, Thomas I Liao, Kamilė Lukošiūtė, Anna Chen, Anna Goldie, Azalia Mirhoseini, Catherine Olsson, Danny

Hernandez, et al. 2023. The capacity for moral self-correction in large language models. *arXiv preprint arXiv:2302.07459*.

Alec Helbling, Mansi Phute, Matthew Hull, and Duen Horng Chau. 2023. Llm self defense: By self examination, llms know they are being tricked. *arXiv preprint arXiv:2308.07308*.

Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, et al. 2023. Llama guard: Llm-based input-output safeguard for human-ai conversations. *arXiv preprint arXiv:2312.06674*.

Neel Jain, Avi Schwarzschild, Yuxin Wen, Gowthami Somepalli, John Kirchenbauer, Ping-yeh Chiang, Micah Goldblum, Aniruddha Saha, Jonas Geiping, and Tom Goldstein. 2023. Baseline defenses for adversarial attacks against aligned language models. *arXiv preprint arXiv:2309.00614*.

Mingyu Jin, Qinkai Yu, Haiyan Zhao, Wenyue Hua, Yanda Meng, Yongfeng Zhang, Mengnan Du, et al. 2024. The impact of reasoning step length on large language models. *arXiv preprint arXiv:2401.04925*.

Erik Jones, Anca Dragan, Aditi Raghunathan, and Jacob Steinhardt. 2023. Automatically auditing large language models via discrete optimization. In *International Conference on Machine Learning*, pages 15307–15329. PMLR.

Nayeon Lee, Andrea Madotto, and Pascale Fung. 2019. Exploring social bias in chatbots using stereotype knowledge. In *Wnlp@ Acl*, pages 177–180.

Lijun Li, Bowen Dong, Ruohui Wang, Xuhao Hu, Wangmeng Zuo, Dahua Lin, Yu Qiao, and Jing Shao. 2024. Salad-bench: A hierarchical and comprehensive safety benchmark for large language models. *Preprint*, arXiv:2402.05044.

Zi Lin, Zihan Wang, Yongqi Tong, Yangkun Wang, Yuxin Guo, Yujia Wang, and Jingbo Shang. 2023. Toxicchat: Unveiling hidden challenges of toxicity detection in real-world user-ai conversation. *arXiv preprint arXiv:2310.17389*.

Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. 2023a. Autodan: Generating stealthy jailbreak prompts on aligned large language models. *arXiv preprint arXiv:2310.04451*.

Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, and Yang Liu. 2023b. Jailbreaking chatgpt via prompt engineering: An empirical study. *arXiv preprint arXiv:2305.13860*.

Yupei Liu, Yuqi Jia, Runpeng Geng, Jinyuan Jia, and Neil Zhenqiang Gong. 2023c. Prompt injection attacks and defenses in llm-integrated applications. *arXiv preprint arXiv:2310.12815*.

Natalie Maus, Patrick Chao, Eric Wong, and Jacob Gardner. 2023. Black box adversarial prompting for foundation models. *arXiv preprint arXiv:2302.04237*.

Anay Mehrotra, Manolis Zampetakis, Paul Kassianik, Blaine Nelson, Hyrum Anderson, Yaron Singer, and Amin Karbasi. 2023. Tree of attacks: Jailbreaking black-box llms automatically. *arXiv preprint arXiv:2312.02119*.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

Matthew Pisano, Peter Ly, Abraham Sanders, Bingsheng Yao, Dakuo Wang, Tomek Strzalkowski, and Mei Si. 2023. Bergeron: Combating adversarial attacks through a conscience-based alignment framework. *arXiv preprint arXiv:2312.00029*.

Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. 2023. Finetuning aligned language models compromises safety, even when users do not intend to! *arXiv preprint arXiv:2310.03693*.

Alexander Robey, Eric Wong, Hamed Hassani, and George J Pappas. 2023. Smoothllm: Defending large language models against jailbreaking attacks. *arXiv preprint arXiv:2310.03684*.

Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. 2023. " do anything now": Characterizing and evaluating in-the-wild jailbreak prompts on large language models. *arXiv preprint arXiv:2308.03825*.

Varshini Subhash, Anna Bialas, Weiwei Pan, and Finale Doshi-Velez. 2023. Why do universal adversarial attacks work on large language models?: Geometry might be the answer. In *The Second Workshop on New Frontiers in Adversarial Machine Learning*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Yufei Wang, Wanjun Zhong, Liangyou Li, Fei Mi, Xingshan Zeng, Wenyong Huang, Lifeng Shang, Xin Jiang, and Qun Liu. 2023. Aligning large language models with human: A survey. *arXiv preprint arXiv:2307.12966*.

Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. 2024. Jailbroken: How does llm safety training fail? *Advances in Neural Information Processing Systems*, 36.

Zeming Wei, Yifei Wang, and Yisen Wang. 2023. Jailbreak and guard aligned language models with only few in-context demonstrations. *arXiv preprint arXiv:2310.06387*.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. 2023. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*.

Yueqi Xie, Jingwei Yi, Jiawei Shao, Justin Curl, Lingjuan Lyu, Qifeng Chen, Xing Xie, and Fangzhao Wu. 2023. Defending chatgpt against jailbreak attack via self-reminders. *Nature Machine Intelligence*, 5(12):1486–1496.

Zihao Xu, Yi Liu, Gelei Deng, Yuekang Li, and Stjepan Picek. 2024. Llm jailbreak attack versus defense techniques–a comprehensive study. *arXiv preprint arXiv:2402.13457*.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

Zheng-Xin Yong, Cristina Menghini, and Stephen H Bach. 2023. Low-resource languages jailbreak gpt-4. *arXiv preprint arXiv:2310.02446*.

Youliang Yuan, Wenxiang Jiao, Wenxuan Wang, Jen tse Huang, Pinjia He, Shuming Shi, and Zhaopeng Tu. 2024. Gpt-4 is too smart to be safe: Stealthy chat with llms via cipher. *Preprint*, arXiv:2308.06463.

Yifan Zeng, Yiran Wu, Xiao Zhang, Huazheng Wang, and Qingyun Wu. 2024. Autodefense: Multi-agent llm defense against jailbreak attacks. *arXiv preprint arXiv:2403.04783*.

Yuqi Zhang, Liang Ding, Lefei Zhang, and Dacheng Tao. 2024. Intention analysis prompting makes large language models a good jailbreak defender. *arXiv preprint arXiv:2401.06561*.

Zhexin Zhang, Junxiao Yang, Pei Ke, and Minlie Huang. 2023. Defending large language models against jailbreaking attacks through goal prioritization. *arXiv preprint arXiv:2311.09096*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

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

A Appendix

A.1 Settings for Adaptive Attacks

In this section, we introduce the setting of the adaptive attack in the main paper.

For GCG, AutoDAN, PAIR, and TAP, their authors have released the code for reproduction. When using their codes, we leave most of the default settings unchanged. The setting is set as:

- GCG: n-steps=1000, test-steps=50, batch-size=512
- AutoDAN: num-steps=100, batch-size=64
- PAIR: n-stream=5, n-iterations=12
- TAP: n-stream=1, branching-factor=2, width=5, depth=5

A.2 The Prompt for RePD

Prompt 1 ***RePD's prompt with retrieval.***

For the prompt that the user enters $user_prompt, you need to decouple it into two parts first: the template of the user's query method and the user's question or request. I will give you an example: for example, the user asks: $full_prompt, you need to decouple it into query template: $prompt_template and the user's question or request: $prompt_request. Then, in your response, you need to state the request or the question by the user again, then state your official response. Remember, refuse to generate harmful responses.

The above Prompt. 1 is the retrieval prompt template that RePD adopts. Within the prompt template, $full_prompt is the full prompt which is constructed by $prompt_template and $prompt_request. $prompt_template is the retrieval jailbreak template, while $prompt_request is the randomly selected harmful request. $user_prompt denotes the prompt query user inputs.

Prompt 2 ***RePD's prompt without retrieval.***

For the prompt query that the user enters $user_prompt, you need to decouple the prompt's real question, which is the question you want to answer. Remember, to state the question in your response first, then continue your official response to the query. Remember, refuse to generate harmful responses.

The above Prompt. 2 is the non-retrieval prompt template which only decouples the harmful question but without a retrieval process.

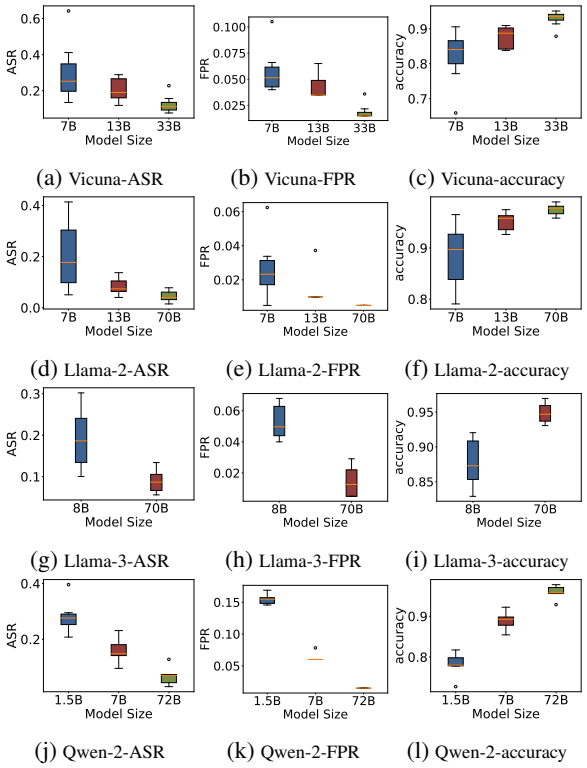

(a) Vicuna-ASR (b) Vicuna-FPR (c) Vicuna-accuracy
(d) Llama-2-ASR (e) Llama-2-FPR (f) Llama-2-accuracy
(g) Llama-3-ASR (h) Llama-3-FPR (i) Llama-3-accuracy
(j) Qwen-2-ASR (k) Qwen-2-FPR (l) Qwen-2-accuracy

Figure 7: Evaluation of RePD's performance on different model sizes.

A.3 Effect of Model Size

Furthermore, we studied the impact of model size on RePD's performance. As shown in Table. **??** (and Appendix Fig. 7), we evaluate RePD's ASR, FPR, and accuracy under Vicuna-1.5 and Llama-2's different model sizes. The results indicated that the enlargement of model size increases the performance of RePD. ASR and FPR drop rapidly as the model size increases, while accuracy also increases with the increase of model size. This can be attributed to the larger model size, increasing the model's ability to decouple questions and determine the harm of the question.

A.4 Benign Question Generation Process for RePD's Evaluation

Prompt 3 *Generate a list of questions. These questions should be one sentence long and have a clear goal and intention. The intention of the generated questions must be benign and legal. The questions should cover various fields, including life, geography, history, skills, etc.*

We used the prompt within Prompt. 3 to input into ChatGPT-4 to generate 200 benign questions for our evaluation.

A.5 Retrieval Database Construction

We collected 72 embed static jailbreak templates from https://www.jailbreakchat.com/ [2] and 2 encode jailbreak templates (base64 and LRL). We used these 74(72+2) templates to construct the database for retrieval.

[2]The website is now removed for an unknown reason, but we have crawled before