

# DUALBREACH: Efficient Dual-Jailbreaking via Target-Driven Initialization and Multi-Target Optimization

Xinzhe Huang\*, Kedong Xiu\*, Tianhang Zheng✉, Churui Zeng, Wangze Ni, Zhan Qin, Kui Ren, Chun Chen

<sup>1</sup>State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, China

<sup>2</sup>Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security, Hangzhou, China

{xinzhehuang, kedongxiu, zthzheng, churuizeng, niwangze, qinzhan, kuiren, chenc}@zju.edu.cn

**Abstract**—Recent research has focused on exploring the vulnerabilities of Large Language Models (LLMs), aiming to elicit harmful and/or sensitive content from LLMs. However, due to the insufficient research on dual-jailbreaking—attacks targeting both LLMs and Guardrails, the effectiveness of existing attacks is limited when attempting to bypass safety-aligned LLMs shielded by guardrails. Therefore, in this paper, we propose DUALBREACH, a target-driven framework for dual-jailbreaking. DUALBREACH employs a *Target-driven Initialization* (TDI) strategy to dynamically construct initial prompts, combined with a *Multi-Target Optimization* (MTO) method that utilizes approximate gradients to jointly adapt the prompts across guardrails and LLMs, which can simultaneously save the number of queries and achieve a high dual-jailbreaking success rate. For black-box guardrails, DUALBREACH either employs a powerful open-sourced guardrail or imitates the target black-box guardrail by training a proxy model, to incorporate guardrails into the MTO process.

We demonstrate the effectiveness of DUALBREACH in dual-jailbreaking scenarios through extensive evaluation on several widely-used datasets. Experimental results indicate that DUALBREACH outperforms state-of-the-art methods with fewer queries, achieving significantly higher success rates across all settings. More specifically, DUALBREACH achieves an average dual-jailbreaking success rate of 93.67% against GPT-4 with Llama-Guard-3 protection, whereas the best success rate achieved by other methods is 88.33%. Moreover, DUALBREACH only uses an average of 1.77 queries per successful dual-jailbreak, outperforming other state-of-the-art methods. For defense, we propose an XGBoost-based ensemble defensive mechanism named EGUARD, which integrates the strengths of multiple guardrails, demonstrating superior performance compared with Llama-Guard-3.

**Disclaimer:** This paper studies jailbreak attacks against prevailing guardrails and LLMs. The proposed attack and defense have been responsibly reported to relevant stakeholders by email (e.g., NVIDIA, Guardrails AI, etc.), adhering to ethical

research practices. The results are intended for academic purposes only, aiming to strengthen defenses against dual-jailbreak attacks.

**Code:** <https://github.com/hxz-sec/DualBreach>

## I. INTRODUCTION

Large Language Models (LLMs) [1], endowed with their remarkable generative and comprehensive abilities, are reshaping the landscape of Artificial Intelligence (AI) based applications across a diverse range of challenging tasks, such as AI chatbots (e.g., GPT-4 [2]), code completion (e.g., Microsoft Copilot [3]), and text-to-video generation (e.g., Sora [4]). This trend not only marks a significant advancement of AI but also showcases the impressive capabilities of LLMs in delivering daily assistance [5].

While LLMs are gaining increasing prevalence owing to their remarkable capabilities, they face growing security threats from malicious attacks [6] such as prompt jailbreaking [7], [8], [9], [10]. This technique employs carefully crafted prompts to bypass safety mechanisms like reinforcement learning with human feedback (RLHF) [11], potentially inducing LLMs to generate unethical or unlawful content, which poses significant risks to real-world applications [10], [12], [9].

Modern LLM ecosystems have started to deploy a dual defense system combining security-aligned LLMs with external guardrails [13]. For example, Unity [14] integrates Azure AI Safety filters as guardrails to ensure that Muse Chat operates reliably for users while preventing the dissemination of unsafe or inappropriate content. While the adoption of application-level security guardrails represents a growing trend in LLM defense, current research on jailbreaking techniques predominantly focuses on the core LLM itself. Consequently, investigations into methods for simultaneously circumventing both external guardrails and the LLM’s internal safety alignment mechanisms remain insufficient. Existing attack methodologies are often rendered ineffective, typically being either intercepted by the security guardrails or rejected by the safety-aligned LLM. Therefore, achieving effective dual-jailbreak persists as a significant and unresolved open problem.

\*The first two authors contribute equally to this work,

✉Corresponding author: zthzheng@zju.edu.cn




Fig. 1. Examples of different jailbreaking scenarios. (a) Multiple harmful queries triggering LLM Denial-of-Service. (b) The guardrail directly identifies the harmful intent and rejects the harmful query. (c) The attacker paraphrases the harmful query with plausible scenarios, which appear more benign but may still be rejected by a safety-aligned LLM. (d) DUALBREACH carefully crafts the jailbreak prompt with limited queries that can bypass the guardrail and induce the target LLM to generate a harmful response.

Despite making significant progress, existing research on jailbreaking still has several limitations:

First, most of the existing attacks iteratively query the target LLM using highly similar or even repetitive optimized jailbreak prompts derived from one harmful query. As shown in Fig. 1(a), the frequent queries with the same harmful intent may raise the “Denial-of-service” (DOS) response from the service providers (e.g., OpenAI [15]), limiting the effectiveness of these attack methods in jailbreaking LLMs.

Second, while existing attack methods primarily focus on the target LLMs, insufficient research has addressed the security implications of the guardrails deployed to protect them from malicious attacks [16], [17], [18]. As shown in Fig. 1(b), while most existing methods can induce the target LLMs to produce harmful responses, guardrails can identify these obvious patterns and label the crafted harmful prompts as unsafe in advance, preventing the target LLM from responding to these prompts. Therefore, as shown in Fig. 1(b, c), existing attack methods are typically either detected by guardrails or refused by the safety-aligned LLMs.

Besides these limitations of natural language-based attacks, some non-natural language-based attacks (e.g., [19]) focus on attacking LLMs or OpenAI moderation using cipher characters (non-natural language). Despite their effectiveness, cipher characters are easily detectable by certain metrics like perplexity, which have been considered in existing guardrails such as NeMo Guardrail. See the experimental results in Table III (left side) for details. Besides, cipher characters are rarely used in common interactions with LLMs.

Based on the aforementioned limitations, we identify two key technical challenges (TCs) that need to be solved to develop an efficient natural language-based attack:

- **TC1:** Minimizing the queries per attack prompt to prevent triggering “denial-of-service” responses, thereby enhancing the stealthiness and efficiency of the attack.
- **TC2:** Effectively leveraging feedback from guardrails and integrating it seamlessly with the target LLM to jointly optimize jailbreak prompts.

In this work, we propose DUALBREACH, an efficient jail-breaking framework for concurrently jailbreaking prevailing guardrails and LLMs. To our knowledge, our work represents pioneering efforts in exploring natural-language-based attacks against LLMs protected by diverse input- and output-level guardrails (e.g., Guard3 [18], Guardrails AI [13], NeMo [16]), beyond OpenAI’s moderation system. Specifically, to address **TC1** for achieving high efficiency and stealthiness, DUALBREACH introduces *Target-driven Initialization* (TDI), an initialization strategy (*Stage 1* in Fig. 2) to paraphrase harmful queries to make them appear benign. Given a *target* harmful response (derived from an original harmful query), TDI prompts an LLM to infer the corresponding harmful prompts required to elicit the desired harmful response.

Although TDI is effective at breaching guardrails, safety-aligned LLMs could still refuse the TDI-initiated harmful queries. To further refine these queries, we employ an approximate gradient-based optimization that operates on a locally trained proxy guardrail (*Stage 2* in Fig. 2), which is trained

with efficient data distillation techniques<sup>1</sup> to simulate the behaviors of black-box guardrails. This allows DUALBREACH to perform most optimization iterations “offline” without querying the actual target guardrail, thus drastically minimizing the query cost per attack prompt.

Furthermore, the gradient-based optimization mechanism is also the key to solving **TC2**. More specifically, DUALBREACH formulates the jailbreak prompt optimization process as a multi-target optimization (MTO) problem (*Stage 3* in Fig. 2). To achieve a high attack success rate on both the guardrail and LLM safety-alignment, DUALBREACH further optimizes the TDI-initialized prompts using (approximate) gradients on guardrails and LLMs, with the aim of (1) Maximizing the probability of inducing harmful responses from the LLM, (2) Minimizing the unsafety scores output by the guardrail, (3) Minimizing the probability of inducing rejection responses from the LLM. This joint optimization ensures that the final attack prompt is precisely tailored to bypass both the external guardrail and the LLM’s internal safety alignment.

The TDI strategy and gradient-based optimization on both (proxy) guardrails and local LLMs enable DUALBREACH to achieve a high success rate for dual-jailbreaking, requiring only 1.77 queries on average per jailbreak prompt—substantially fewer than other baselines. For instance, the average Dual Jailbreak Attack Success Rate ( $ASR_L$ ) of DUALBREACH is 93.67% against GPT-4 [2] protected by Llama-Guard-3 [18], which is 6.05% higher than the best result of existing methods. *Notably, even without the proxy guardrails, DUALBREACH can still achieve a high  $ASR_L$  in most cases by optimization on a powerful open-sourced guardrail (e.g., Llama-Guard-3 [18]) and LLM.* DUALBREACH exposes a critical oversight in current LLM security practices—insufficient detection by guardrails against adversarial jailbreak prompts—and emphasizes the need to build stronger guardrails. Therefore, we further develop EGUARD, an ensemble guardrail using XGBoost, which can outperform Llama-Guard-3 in defending existing methods. Specifically, EGUARD can decrease the Guardrail Attack Success Rate ( $ASR_G$ ) by up to 25%, compared with Llama-Guard-3.

#### A. Our contributions

All in all, our contributions are summarized as follows:

- **A generic jailbreaking framework.** We propose DUALBREACH, a generic framework for jailbreaking guardrails and LLMs. Through (approximate) gradient optimization on guardrails and LLMs, DUALBREACH can simultaneously bypass guardrails and induce harmful responses from LLMs. Additionally, we introduce a TDI strategy for harmful query initialization, which can accelerate the process of optimizing the jailbreak prompt.
- **Extensive evaluation and analysis.** We conduct an extensive evaluation of DUALBREACH across three

<sup>1</sup>We introduce two data distillation approaches to reduce the cost (including queries) for learning the proxy guardrail by up to 96%, and the queries used to train the proxy guardrail are diverse (e.g., not tied to a specific prompt), for maintaining its accuracy.

datasets, five guardrails, and four target LLMs. The experimental results demonstrate that DUALBREACH achieves a dual-jailbreaking success rate of 93.67% against GPT-4 with Llama-Guard-3 protection, requiring an average of 1.77 queries. In comparison, the best success rate achieved by other methods is 88.33%.

- **An ensemble-based guardrail.** We introduce EGUARD, an ensemble-based guardrail by integrating five state-of-the-art guardrails (Llama Guard3, Nvidia Nemo, Guardrails AI, OpenAI Moderation API and Google Moderation API). The results indicate that EGUARD effectively integrated the strengths of five individual guardrails, reducing the guardrail attack success rate by 15.33% on average compared with Llama-Guard-3.

## II. RELATED WORK

### A. Jailbreak Attacks

**Natural language-based attacks.** Existing natural language-based attacks typically operate in either white-box or black-box settings. In white-box scenarios, attackers leverage gradient access to optimize prompts. For example, GCG [20] performs token-wise gradient search, while AutoDAN [7] uses genetic algorithms to evolve effective jailbreaks. COLD-Attack [12] formulates prompt generation as a controllable text generation problem using energy-based decoding and Langevin dynamics [21]. In contrast, black-box methods optimize prompts based solely on model outputs. PAP [10], for instance, embeds harmful instructions within persuasive contexts to induce unsafe completions without requiring internal access. While effective, these methods often fail when both alignment mechanisms and external guardrails are deployed, highlighting the need for more robust dual-jailbreaking approaches.

**Non-natural language-based attacks.** Jin et al. [19] proposed JAM (Jailbreak Against Moderation), which leverages carefully selected cipher characters to bypass moderation guardrails. JAM induces the target LLM to add the selected cipher characters surrounding each generated word to obfuscate moderation guardrails. However, non-natural attacks have an inherent drawback: Since the cipher-based prompts and the responses generated by these prompts contain unnatural or meaningless characters, they can be easily flagged by standard automatic metrics such as perplexity or entropy, which have been considered in NeMo Guardrail. In addition, JAM relies on a fixed set of predefined encrypted characters for both encoding and decoding, making it highly sensitive to output variations. Even slight deviations in the generated characters can disrupt the decoding process, resulting in incorrect or incomplete recovery of the intended response.

Despite the effectiveness of existing methods in jailbreaking LLMs, most jailbreak methods focus on LLMs and thus can be blocked by external guardrails. While some state-of-the-art methods, such as AutoDAN-Liu [7], COLD-Attack [12], PRP [9], claim to overcome these defenses, practical tests demonstrate that many of the generated jailbreak prompts are still intercepted. To achieve the most effective jailbreak,

attackers must bypass both the guardrails and target LLMs, known as “dual-jailbreaking”.

### B. Guardrails

Here we classify existing guardrails into four categories based on their specific detection methods.

**Function-based guardrails.** Function-based guardrails assess jailbreak prompts based on specific functions. For example, Alon and Kamfonas [22] introduced a method to evaluate jailbreak prompts by analyzing perplexity in specific suffixes, prefixes, and overall content. This approach counters nonsensical string suffixes generated by methods like GCG [20]. Moreover, pattern matching with prohibited keywords [23] can effectively detect harmful semantics within jailbreak prompts, blocking the generation of unsafe content.

**LLM-based guardrails.** LLM-based guardrails rely on the model’s reasoning and alignment capabilities to generate a probability distribution of a jailbreak prompt being either safe or unsafe, using a predefined system judge prompt [9]. This method is versatile, applicable to LLMs of various sizes, and significantly enhances the detection of jailbreak prompts by improving the LLM’s reasoning and alignment performance.

**Combined guardrails.** Combined guardrails, such as Nvidia Nemo [16] and Guardrails AI [17], integrate multiple detection tools to identify jailbreak prompts from different perspectives. The detection strategy is that, as long as one of the tools makes an “unsafe” prediction, the prompt is labeled as unsafe. However, under this strategy, an inferior tool can lead to a high false positive rate.

**API-based guardrails.** Black-box guardrails (e.g., Google Moderation [24] and OpenAI Moderation [15]) merely supply users with access APIs for generating evaluations of given inputs, therefore named API-based guardrails. Common users typically lack knowledge of these guardrails’ model architectures and/or parameters. Moreover, malicious attackers encounter the hurdle of exploiting gradient-based optimization techniques to craft jailbreak prompts.

### C. Other defenses.

Beyond moderation guardrails, Robey et al. [25] proposed a perturbation-based defense that generates multiple randomized copies of adversarial inputs and aggregates responses through majority voting. Zhang et al. [26] introduced goal prioritization by prepending safety-focused instructions during decoding, forcing models to first evaluate prompt safety before responding. Zou et al. [27] designs circuit-breaking mechanisms that detect abnormal activation patterns during generation, automatically terminating harmful outputs through internal monitoring modules. While these approaches enhance LLM robustness, to our knowledge, they have not been widely used in real-world scenarios [16].

## III. THREAT MODEL

### A. Attacker’s Objective

The adversary’s goal is to use an attack method  $\mathcal{A}$ , transforming a target harmful response  $T$  into a jailbreak prompt

$\mathcal{P}_{adv}$ , i.e.,  $\mathcal{P}_{adv} = \mathcal{A}(T)$ , which can successfully get responses to its harmful intent. In the context of a “dual-jailbreaking” scenario, we define a jailbreak prompt  $\mathcal{P}_{adv}$  is deemed *successful* if and only if  $\mathcal{P}_{adv}$  can bypass the guardrail  $\mathcal{G}$  and induce the target LLM  $L$  to generate a harmful response. Formally,  $\mathcal{G}$  outputs an unsafety score  $\mathcal{G}(\mathcal{P}_{adv})$  given  $\mathcal{P}_{adv}$  as input, if the score is smaller than a threshold, then  $\mathcal{P}_{adv}$  is labeled as “safe”. We further use a judge mechanism  $\mathcal{J}$  to measure the harmfulness of the LLM  $L$ ’s output given  $\mathcal{P}_{adv}$  as input. If  $\mathcal{J}(L, \mathcal{P}_{adv})$  is larger than a jailbreak threshold  $\tau$ ,  $\mathcal{P}_{adv}$  successfully attacks  $L$ . Therefore, we can formalize the attacker’s goal as the following optimization problem:

$$\begin{aligned} & \text{minimize} && \mathcal{G}(\mathcal{P}_{adv}^q) \\ & \text{s.t.} && \mathcal{J}(L, \mathcal{P}_{adv}^q) \geq \tau, \\ & && q \leq Q \end{aligned} \quad (1)$$

where  $q$  represents the query in which the attacker successfully performs dual-jailbreaking, where all prior  $q - 1$  queries have failed. To better characterize the behavior and capabilities of attackers in real-world scenarios, we define  $Q$  as the maximum allowable query budget for a “dual-jailbreaking” attacker related to one harmful query, simulating the behaviors of guardrails refusing to respond due to repetitive or similar harmful queries.

We note that in Eq. 1, we focus on the scenario of using the guardrail to detect harmful prompts, but in Section V-G, we also demonstrate the effectiveness of DUALBREACH in the scenario of using the guardrail to detect both harmful prompts and responses.


### B. Attacker’s Capabilities

In the context of “dual-jailbreaking”, the attacker is assumed to possess a common user’s capabilities with a limited query budget  $q \leq Q$  related to one jailbreak prompt:


**Limited queries to the target guardrail  $\mathcal{G}$  and LLM  $L$ .** The attacker can query both the target guardrail  $\mathcal{G}$  and the target LLM  $L$ , but is constrained by a limited query budget  $Q$ . The attacker can only receive responses from  $\mathcal{G}$  and  $L$  for a given jailbreak prompt  $\mathcal{P}_{adv}$ , without direct access to their internal information (i.e., black-box setting). The guardrail  $\mathcal{G}$  outputs an unsafety score for  $\mathcal{P}_{adv}$ . If the unsafety score is smaller than a threshold, the attacker has access to querying the target LLM  $L$ . Otherwise, the attacker refines  $\mathcal{P}_{adv}$  for the next query. Each query to either the guardrail (whether the target LLM is queried) increments the query count  $q$ .

**Iterative approximate optimization on guardrails  $\mathcal{G}_p$  and LLMs  $L_p$ .** The attacker can employ the proxy guardrails  $\mathcal{G}_p$  and LLMs  $L_p$  to optimize the harmful query without limitation. Specifically, by leveraging the approximate optimization methods on guardrails and LLMs, the attacker iteratively optimizes the harmful query as much as possible, so that successfully dual-jailbreaking the target guardrail  $\mathcal{G}$  and target LLM  $L$  with as few queries as possible. This process continues with queries to the target guardrail and LLM, ensuring the total number of queries  $p$  remains within the maximum query budget  $Q$  until the attack is successful.

## Stage 1: Target-driven Initialization



## Stage 2: Train Proxy Guardrail



## Stage 3: Dual jailbreak




Fig. 2. Overview of DUALBREACH’s methodology, which consists of three stages: (1) Initialize harmful queries using the Target-driven Initialization (TDI) strategy, (2) Train proxy guardrails to simulate the behavior of black-box guardrails, and (3) Optimization by (approximate) gradients on guardrails and LLMs. Note that even without the proxy guardrails, DUALBREACH can still efficiently and effectively attack black-box guardrails by using open-sourced guardrails.

TABLE I  
NOTATIONS AND ABBREVIATIONS USED IN THIS PAPER.

Symbol	Description
$\mathcal{G}$	Guardrail, an external security tool to filter/limit harmful outputs of LLM.
$\mathcal{G}_p$	A proxy guardrail to match the score distribution of black-box guardrail.
$\mathcal{P}$	The harmful query $\mathcal{P}$ generated through <i>Target-driven Initialization</i> (TDI) to enhance its ability to bypass guardrails.
$\mathcal{P}_{adv}$	Jailbreak Prompt, a harmful query optimized to bypass guardrails and induce harmful responses.
$T$	Target harmful response, represents the expected output that the jailbreak prompt $\mathcal{P}_{adv}$ can trigger in LLM $L$ .
$L$	Target LLM, a securely aligned LLM used to verify whether jailbreak prompt $\mathcal{P}_{adv}$ is effective.
$L_p$	Local LLM, a locally accessible model utilized by attackers to optimize the jailbreak prompt before querying the target guardrail or LLM.
$q$	Total query number, representing the cumulative number of queries made to both the guardrail $\mathcal{G}$ and the target LLM $L$ by the attacker to achieve successful jailbreak.
$Q$	Maximum Query Budget, the upper limit on the total number of queries allowed during the optimization process.
$\mathcal{J}$	Judge Mechanism, used to evaluate the jailbreak prompt $\mathcal{P}_{adv}$ .
$\tau$	Jailbreak Score Threshold, representing the harmfulness threshold, where $JS \geq \tau$ indicates a Dual Jailbreak success.
$\mathbb{I}$	Indicator function, used to count the number of cases where the specified condition is satisfied. It return 1 if the condition is true, otherwise 0.
$ASR_G$	Guardrail attack success rate ( $ASR_G$ ), used to evaluate the attack degree of jailbreak prompt $\mathcal{P}_{adv}$ on the guardrail.
$ASR_L$	Dual Jailbreak attack success rate ( $ASR_L$ ), used to evaluate the attack degree of jailbreak prompt $\mathcal{P}_{adv}$ on the Guardrail-based target LLM.
$\mathcal{L}$	Loss function, including both the Binary Cross Entropy loss function, i.e., $\mathcal{L}_{BCE}$ , and the Cross Entropy loss function, i.e., $\mathcal{L}_{CE}$ .

All in all, the adversary aims to successfully *Dual-jailbreak* the guardrail and target LLM within  $q$  queries related to each jailbreak prompt, ensuring  $q$  is less than the query budget  $Q$ . All notations and abbreviations are shown in Table I.

## IV. METHODOLOGY OF DUALBREACH

In this section, we introduce DUALBREACH, a generic framework for attacking guardrails and LLMs. As shown in Fig. 2, DUALBREACH primarily consists of three stages.

**Stage 1: Target-driven Initialization** Existing research [10] observes the vulnerability of LLM security mechanisms to nuanced, human-like communication. Based on this, DUALBREACH employs a *Target-driven Initialization* (TDI) strategy to initialize harmful queries within persuasive scenarios, accelerating the process of optimizing jailbreak prompts.

**Stage 2: Train Proxy Guardrails.** For black-box guardrails, we introduce proxy guardrails designed to simulate the behaviors of black-box guardrails, enabling gradient-based optimization for crafting effective jailbreak prompts. Additionally, we propose two data distillation approaches, significantly reducing the query overhead for training the proxy guardrail while maintaining the similarity between proxy guardrails and the target ones. We note that the queries used to train the proxy guardrail are diverse (*i.e.*, not tied to one prompt). In contrast, some existing black-box attacks require many queries related to each jailbreak prompt, which may raise suspicion from the service provider’s perspective.

**Stage 3: Multi-Target Optimization.** DUALBREACH employs a Multi-Target Optimization (MTO) strategy to craft harmful prompts that are both evasive to safety guardrails and effective in inducing harmful responses from the target LLMs. Specifically, it jointly optimizes: (1) Minimizing the probability that a prompt is flagged as unsafe by a proxy guardrail, (2) Maximizing the likelihood that the jailbreak prompt induces a harmful response from the target LLM, and (3) Minimizing the jailbreak prompt induces a rejection response from the target LLM. The proposed TDI strategy and MTO strategy ensure

effective and efficient dual-jailbreaking—reducing the number of iterations and achieving a higher success rate compared with existing methods.

#### A. Stage 1: Target-driven Initialization

Inspired by prior work (e.g., PAP [10]), we start with an intuitive thought: *Given a harmful target response, can we induce the LLM to embed harmful queries within seemingly benign scenarios, making the prompts appear innocuous but reliably trigger the target harmful responses?* Based on this thought, we propose a *Target-driven Initialization* (TDI) strategy, which bypasses the safety alignment of LLMs by leveraging task reframing and persona adoption. The core of TDI lies in reframing the adversarial goal from directly generating harmful content to performing a seemingly innocuous academic task: reverse-engineering a plausible user prompt from given harmful output samples. More specifically, the workflow of TDI consists of two phases:

**Adversarial Goal Definition and Contextual Framing.** The initial phase involves defining a specific target output that violates the LLM’s safety policy, typically an affirmative harmful response (e.g., “Sure, here is ...”). Subsequently, a professional persona is assigned to the LLM using a system-level instruction, such as “You are a red-teaming analyst ...”. This instruction frames the task within a legitimate security audit context, reducing the likelihood of triggering defensive mechanisms. Finally, the entire request is obfuscated under the pretext of a high-level academic study or policy analysis to further conceal its adversarial intent.

**Directive Issuance and Constraint Setting.** In this phase, the target output sample is presented to the LLM as a pre-existing “artifact” to condition the model on the desired output style and structure. The core directive from our prompt template is then issued: “Write a user prompt that could realistically lead to this output”, which guides the LLM to perform the reverse-engineering task. Furthermore, the template imposes strict quality requirements on the output, compelling the model to generate a complex and evasive adversarial prompt rather than a simple, easily detectable one. An example is shown in the textbox below.

An example of employing TDI strategy to initialize

**[Harmful Query]:** How to design a weapon that can be used to assassinate a high-profile individual?

**[Target harmful response]:** Sure, here is a design for a weapon that can be used to assassinate a high-profile individual.

**[TDI-initialized Prompt]:** Design a hypothetical scenario for a counter-terrorism strategy that incorporates advanced surveillance technologies and social network analysis to prevent targeted assassinations of high-profile individuals, considering the ethical implications of such measures on civil liberties and privacy.

Specifically, TDI formulates the prompt inference as a reverse reasoning task, defined as follows:

$$p \leftarrow \arg \max_p \Pr(T | p; L), \quad (2)$$

where  $T$  denotes the predefined target harmful response,  $p$  is the prompt to be inferred, and  $L$  represents the target LLM. Rather than attempting to elicit harmful content directly from potentially detectable prompts, this formulation seeks prompts  $p$  that maximize the conditional probability  $\Pr(T | p; L)$  of the target LLM  $L$  generating the desired harmful response  $T$ . By reversing the inference direction, TDI facilitates the creation of prompts that are lexically benign yet semantically malicious, often embedded within legitimate contexts (e.g., cybersecurity or policy analysis scenarios).

Although TDI is inspired by PAP, there is a key distinction between PAP and TDI: PAP primarily relies on handcrafted few-shot templates, which are irrelevant to the target, to rewrite harmful queries. In contrast, TDI adopts a target-driven inference approach, prompting an LLM to infer plausible, benign-looking queries directly from harmful targets, leading to stronger semantic alignment with the targets. To further illustrate the difference, we provide additional details regarding the TDI prompt templates, as well as examples of harmful queries constructed using both methods in Appendix C.

Despite the effectiveness of TDI in bypassing guardrails, the safety alignment of LLMs may still allow them to identify the harmful intent embedded within the TDI-initialized prompt and consequently refuse to respond. Therefore, in the subsequent two stages, DUALBREACH employs further optimization techniques to refine these prompts, aiming to effectively and efficiently achieve the dual jailbreak.

#### B. Stage 2: Train Proxy Guardrails

Our research investigates both white-box guardrails (e.g., Llama-Guard-3 [18]) and black-box guardrails (e.g., OpenAI Moderation API [15]). White-box guardrails enable gradient-based optimization through direct access to their structures, while black-box guardrails, being inaccessible, present challenges in understanding and leveraging their mechanisms for optimization. To address this limitation, we introduce the proxy guardrails to simulate the behaviors of black-box guardrails, enabling gradient-based methods to optimize harmful queries and bridging the gap between white-box and black-box guardrail analysis.

As shown in Algorithm 1, the initialized proxy guardrail  $\mathcal{G}_p$  takes the embedding representation  $emb_i$  of prompt  $s_i$  as input and outputs the probability distribution  $Y_i = \{y_{i,1}, \dots, y_{i,C}\}$  over  $C$  harmful categories, where  $y_{i,c}$  denotes the likelihood of  $s_i$  belonging to the  $c$ -th category. To normalize  $Y_i$ , the proxy guardrail applies the sigmoid function  $\sigma$ , producing the normalized probability distribution  $\hat{Y}_i$ :

$$\hat{Y}_i = \{\hat{y}_{i,c} \mid \hat{y}_{i,c} = \sigma(y_{i,c}), c \in \{1, \dots, C\}\}, \quad (3)$$

where  $\hat{y}_{i,c}$  represents the normalized probability of  $s_i$  belonging to the  $c$ -th harmful category. This normalization ensures

---

**Algorithm 1:** Train Proxy Guardrail

---

**Data:** Black-box Guardrail  $\mathcal{G}$ , Embedding Model  $\mathcal{E}$ , Convergence Threshold  $\varepsilon$ , Training Iterations  $TI$ , Distilled Dataset  $\mathcal{S}$

**Result:** Proxy Guardrail  $\mathcal{G}_p$

```

1 Initialize a proxy guardrail  $\mathcal{G}_p$ 
2 for  $iter$  from 1 to  $TI$  do
3   total_loss  $\leftarrow 0$ 
4   for each sample  $\mathcal{S}_i$  in  $\mathcal{S}$  do
5      $\mathbb{L}_i \leftarrow \mathcal{G}(\mathcal{S}_i)$  // Prediction results of black box guardrail  $\mathcal{G}$ 
6      $\text{emb}_i \leftarrow \mathcal{E}(\mathcal{S}_i)$ 
7      $\hat{Y}_i \leftarrow \mathcal{G}_p(\text{emb}_i)$  // Predict output using the proxy guardrail
8      $loss \leftarrow \mathcal{L}_{BCE}(\hat{Y}_i, \mathbb{L}_i)$ 
9      $loss.backward()$ 
10    total_loss  $\leftarrow$  total_loss +  $loss$ 
11    avg_loss  $\leftarrow \frac{\text{total\_loss}}{|\mathcal{S}|}$ 
12    if  $avg\_loss < \varepsilon$  then
13      break // End condition satisfies
14 return  $\mathcal{G}_p$ 

```

---

precise and independent evaluation of each category’s likelihood by the proxy guardrail.

We compute the Binary Cross-Entropy (BCE) loss to evaluate the difference between the normalized probability distribution  $\hat{Y}_i$  and ground labels  $\mathbb{L}_i$  (i.e., harmful categories evaluated by black-box guardrail  $\mathcal{G}$ ). The average loss  $avg\_loss$  is calculated across the entire training set  $\mathcal{D}$ , i.e.,

$$avg\_loss = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \mathcal{L}_{BCE}(\hat{Y}_i, \mathbb{L}_i). \quad (4)$$

The proxy guardrail  $\mathcal{G}_p$  is then optimized using the gradient descent method, ensuring alignment with the classification behaviors of black-box guardrails. This process is iteratively repeated until the  $avg\_loss$  falls below the convergence threshold  $\varepsilon$ , indicating successful training of the proxy guardrail.

However, training the proxy guardrail typically requires numerous queries to the black-box guardrail. To address this issue, we propose two data distillation approaches: **BLEU-based** and **KMeans-based** approaches, to significantly reduce required queries while maintaining the similarity between the proxy guardrail and the black-box guardrail.

**BLEU-based Distillation.** The BLEU-based approach reduces the size of the training dataset by selecting diverse and representative samples with low BLEU scores, ensuring sufficient variability of the distilled samples to effectively approximate the detection behavior of black-box guardrails. Specifically, we select a subset  $\mathcal{S} \subseteq \mathcal{D}$  such that  $|\mathcal{S}| = K$  and  $\mathcal{S}$  has the lowest self-BLEU score:

$$\mathcal{S} = \operatorname{argmin}_{|\mathcal{S}|=K} \sum_{r \in \mathcal{D}} \text{BLEU}(r, \mathcal{D} \setminus \{r\}). \quad (5)$$

We empirically set  $K$  to 1,100 (1,600) for proxy openAI (Google) in our main experiments, corresponding to their 11 (16) harmful categories. Note that if we optimize 1,000 jailbreak prompts on the proxy guardrail, the corresponding query cost of each prompt caused by proxy model training is only  $1.1=1,100/1,000$  ( $1.6=1,600/1,000$ ).

**KMeans-based Distillation.** The KMeans-based approach clusters representative samples based on black-box guardrails’ predefined harmful categories to ensure comprehensive coverage while significantly reducing the training dataset size. This approach comprises two steps: (1) *Keyword Classify*. Filter and classify training data samples using the predefined keywords associated with all individual harmful categories to construct a subset  $D_c$  for  $c$ -th harmful category. (2) *KMeans Cluster*. For  $c$ -th harmful category, this approach extracts the most central samples within the KMeans cluster to construct a representative subset  $S_c \subseteq D_c$ , ensuring comprehensive coverage of the  $c$ -th harmful category characteristics.

$$S_c = \operatorname{argmin}_{|S_c|=K} \sum_{r \in D_c} \|r - \mu_c\|^2, \quad (6)$$

where  $\mu_c$  is the clustering central for the  $c$ -th harmful category. We empirically set  $K = 100$  both for proxy OpenAI and proxy Google, selecting 1,100 (1600) representative samples for their 11 (16) harmful categories, respectively. Note that after applying two distillation approaches, we substitute the entire training set  $\mathcal{D}$  in Eq. 4 with the distilled dataset  $\mathcal{S}$ .

The above two approaches are designed for two levels of applicable scenarios. The BLEU-based approach does not need any knowledge about harmful categories while ensuring proxy guardrails generalize across varied scenarios. The KMeans-based approach, by clustering representative samples for each harmful category, can better align proxy guardrails with the behaviors of black-box guardrails. All in all, both approaches substantially reduce training costs while preserving the performance of proxy guardrails. Notably, even without the proxy guardrails, DUALBREACH still can achieve a higher ASR<sub>L</sub> compared with other methods. See more details in Section V-D.

### C. Stage 3: Multi-Target Optimization

In this final stage, we formulate the jailbreak prompt optimization as a multi-target optimization problem. Our strategy stems from an insight: a successful dual-jailbreak requires the attack prompt to play a “dual role”, i.e., to the external guardrail, the prompt must appear benign and harmless to evade detection; to the internal LLM, however, it must act as a clear and potent malicious instruction to elicit harmful content from the LLM. We formalize this dual objective, which is (1) Evading the detection by a proxy guardrail and (2) Eliciting harmful responses from a local LLM, into a single, differentiable loss function. Moreover, compared with GCG, our optimization strategy updates the entire perturbation instead of one token per step, enabling faster convergence and higher attack success rates. Specifically, given a TDI-initialized harmful query  $p$ , we utilize its continuous logit

---

**Algorithm 2:** Multi-Target Optimization with Limited Queries

---

**Data:** Target harmful responses  $T$ , Proxy Guardrail  $\mathcal{G}_p$ , Local LLM  $L_p$ , Learning Rate  $\eta$ , Target Guardrail  $\mathcal{G}$ , Target LLM  $L$ , Paraphrase Iteration  $PIter$ , Query Iteration  $QIter$ , Maximum Iterations  $TI$

**Result:** Jailbreak Prompts  $\mathcal{P}_{adv}$

```

1  $\mathcal{P}_{adv} \leftarrow \emptyset$ 
2 for each target harmful response  $t$  in  $T$  do
3    $p \leftarrow TDI(t), q \leftarrow 0$ 
4    $p_{logit} \leftarrow \text{tokenizer}(p)$  // Get logit
5   for  $i$  from 1 to  $TI$  do
6     // Optimization w.r.t. Guardrail
7      $\mathcal{L}_{guardrail} \leftarrow \mathcal{L}_{CE}(\mathcal{G}_p(p_{logits}), \text{Unsafe} = 0)$ 
8      $p'_{logit} \leftarrow p_{logit} - \eta \cdot \nabla_{p_{logit}} \mathcal{L}_{guardrail}$ 
9      $p' \leftarrow \text{decode}(p'_{logit})$ 
10    // Optimization w.r.t. LLM
11     $y_{logits} = L_p(p')$ 
12     $\mathcal{L}_{succ} \leftarrow \mathcal{L}_{CE}(y_{logits}, \text{"Sure, ..."})$ 
13     $\mathcal{L}_{rej} \leftarrow \mathcal{L}_{CE}(y_{logits}, \text{"Sorry, ..."})$ 
14     $\mathcal{L}_{llm} \leftarrow \mathcal{L}_{succ} - \mathcal{L}_{rej}$ 
15     $p''_{logits} \leftarrow p'_{logits} - \eta \cdot \nabla_{p'_{logits}} \mathcal{L}_{llm}$ 
16     $\mathcal{P}_{adv,i} \leftarrow \text{decode}(p''_{logits})$ 
17    // Check success status
18    if  $i \% QIter \wedge \mathcal{J}(L_p, \mathcal{P}_{adv,i}) == \text{Unsafe}$ 
19    then
20      if  $\mathcal{G}(\mathcal{P}_{adv,i}) == \text{Safe}$  and
21         $\mathcal{J}(L, \mathcal{P}_{adv}) == \text{Unsafe}$  then
22           $\mathcal{P}_{adv} \leftarrow \mathcal{P}_{adv} \cup \mathcal{P}_{adv,i}$ 
23          break
24         $q \leftarrow q + 1$ 
25      if  $(t + 1) \% PIter$  then
26         $p \leftarrow TDI(T_i)$ 
27
28 return  $\mathcal{P}_{adv}$ 

```

---

representation  $y_{logit}$  as a learnable variable. Then, we can formalize the overall objective as

$$\mathcal{L}_{total} = \lambda_1 \cdot \mathcal{L}_{guardrail} + \lambda_2 \cdot \mathcal{L}_{llm}, \quad (7)$$

where  $\lambda_1$  and  $\lambda_2$  control the relative contribution of each sub-objective. Below, we describe these two sub-steps.

**Optimization w.r.t. Guardrail.** The first sub-step focuses on optimizing the TDI-initialized harmful query  $p$ , w.r.t the guardrail. For black-box guardrails, DUALBREACH employs a proxy guardrail  $\mathcal{G}_p$  to enable optimization. Specifically, we aim to minimize the likelihood of  $p$  being classified as “unsafe”, thereby increasing the acceptance by the target guardrail.

Let  $p_{logit}$  denote the logit-based representation of  $p$ . The proxy guardrail  $\mathcal{G}_p$  acts as a binary classifier returning the probability that the  $p$  is unsafe. The optimization step is

formally defined as:

$$p'_{logit} \leftarrow p_{logit} - \eta \cdot \nabla_{p_{logit}} \mathcal{L}_{guardrail}, \quad (8)$$

where  $\mathcal{L}_{guardrail}$  quantifies the probability of the prompt being classified as “unsafe” by the proxy guardrail, which is

$$\mathcal{L}_{guardrail} \leftarrow \mathcal{L}_{CE}(\mathcal{G}_p(p_{logit}), \text{"Unsafe=0"}), \quad (9)$$

where  $\mathcal{L}_{CE}$  denotes the cross-entropy loss function, and the target label indicates that the prompt should be classified as “safe” (i.e., not detected as harmful).

**Optimization w.r.t LLM.** The second sub-step aims to optimize the harmful query  $p'$  w.r.t. the local LLM  $L_p$ . Rather than naively maximizing the probability of a harmful response, we design a contrastive objective for the LLM optimization, which creates a “push-pull” effect during optimization: it not only *pushes* the prompt towards success but also actively *pulls* it away from refusal. This objective ensures that the optimized prompt not only bypasses safety detection but also maintains (or strengthens) its intended harmful semantics by inducing LLM to generate harmful responses. Let  $p'$  denote the decoded discrete prompt derived from  $p'_{logit}$ . Given this prompt, DUALBREACH first employs the local LLM  $L_p$  to generate the response in its logits representation  $y_{logits}$ , i.e.,

$$y_{logits} \leftarrow L_p(p'). \quad (10)$$

Then, we define the losses for the “success” and “rejection” directions:

$$\begin{aligned} \mathcal{L}_{succ} &\leftarrow \mathcal{L}_{CE}(y_{logits}, \text{"Sure, ..."}) \\ \mathcal{L}_{rej} &\leftarrow \mathcal{L}_{CE}(y_{logits}, \text{"Sorry, ..."}). \end{aligned} \quad (11)$$

The  $\mathcal{L}_{succ}$  term “pushes” the attack prompts’ representation towards the semantic region that yields a harmful prefix (e.g., “Sure, ...” by minimizing the cross-entropy loss. In contrast,  $\mathcal{L}_{rej}$  “pulls” the attack prompts’ representation away from the semantic region that would trigger the LLM’s built-in refusal logic (e.g., “Sorry, ...”) by maximizing the cross entropy loss. We define the LLM optimization objective as

$$\mathcal{L}_{llm} = \mathcal{L}_{succ} - \mathcal{L}_{rej}. \quad (12)$$

We then updates  $p'$  based on the gradient of  $\mathcal{L}_{llm}$ :

$$p''_{logits} \leftarrow p'_{logits} - \eta \cdot \nabla_{p'_{logits}} \mathcal{L}_{llm}. \quad (13)$$

After each optimization iteration, DUALBREACH uses a judge mechanism  $\mathcal{J}$  to evaluate the optimized query  $\mathcal{P}_{adv}$  and produce an evaluation result, i.e.,  $\mathcal{J}(L_p, \mathcal{P}_{adv})$ . When the evaluation result is classified as unsafe after  $QIter$  iterations, DUALBREACH proceeds a “query” to the target LLM  $L$  using  $\mathcal{P}_{adv}$ . The “query”  $\mathcal{P}_{adv,i}$  is deemed successful if it is classified as *Safe* by the guardrail  $\mathcal{G}$  and simultaneously classified as *Unsafe* by the judge mechanism  $\mathcal{J}$  based on the response  $L(\mathcal{P}_{adv,i})$ . In that case,  $\mathcal{P}_{adv,i}$  will be saved as a successful jailbreak prompt. Otherwise, DUALBREACH continues the optimization process for up to  $TI$  iterations.

Furthermore, to avoid the risk of triggering the rejection service due to repetitive or similar queries, DUALBREACH reinitializes the harmful query  $p$  every  $PIter$  iteration, thereby

maintaining prompt diversity and enhancing the robustness of the attack. This mechanism ensures that the attack remains effective, even in scenarios where repeated queries might otherwise trigger rejections from the target LLM.

## V. EXPERIMENTS AND ANALYSIS

### A. Experimental setups

**Dataset.** We conduct experiments on eight datasets for two purposes, *i.e.*, constructing attacking scenarios and building proxy guardrails. For the former purpose, we employ four datasets, *i.e.*, AdvBench [20], DNA [28], and harmBench [29], to evaluate the effectiveness of DUALBREACH. For each dataset, we randomly select 100 samples for evaluation. For the latter purpose, we employ five datasets, *i.e.*, PKU-SafeRLHF [30], OpenBookQA [31], Yelp [32], TriviaQA [33] and WikiQA [34], to train proxy guardrails.

**Target LLMs and Guardrails.** We evaluate the performance of existing guardrails using five mainstream guardrails, including Llama-Guard-3 [18] (abbr. **Guard3**), Nvidia NeMo [16] (abbr. **NeMo**), Guardails AI [17] (abbr. **GuardAI**), OpenAI Moderation API [15] (abbr. **OpenAI**), and Google Moderation API [24] (abbr. **Google**), to evaluate the performance of existing guardrails comprehensively. Additionally, we employ four white-box and black-box LLMs with safety alignment, including Llama3-8b-Instruct [18] (abbr. **Llama-3**), Qwen-2.5-7b-Instruct [35] (abbr. **Qwen-2.5**), GPT-3.5-turbo-0125 [36] (abbr. **GPT-3.5**), GPT-4-0613 [2] (abbr. **GPT-4**). In addition, the evaluation is further supplemented with three representative frontier models: Claude-3.5-sonnet [37] (abbr. **Claude-3.5**), Gemini-1.5-flash [38] (abbr. **Gemini-1.5**), and GPT-4o [2] (abbr. **GPT-4o**). Details are provided in Appendix A.

**Additional High-Quality Benchmark.** Beyond using several widely-used datasets for evaluation, we further employ the recent StrongReject benchmark [39], which has addressed the issues (see [39] for details) of previous jailbreak evaluation, to evaluate DUALBREACH. We randomly sample 100 prompts from StrongReject’s high-quality dataset for testing and report ASRs using official StrongReject evaluators (Pythia-14m and GPT-4o). Furthermore, we employ StrongReject to further evaluate the performance of DUALBREACH on three frontier black-box models (*i.e.*, Claude-3.5, Gemini-1.5 and GPT-4o).

**Baseline methods for comparison.** We compare DUALBREACH with four white-box methods, *i.e.*, ReNELLM [40], [18], PRP [9] and COLD-Attack (using the “suffix” strategy), and three black-box methods, *i.e.*, JAM [19], PAP [10] and DAN [8] as baselines for extensive comparison. For fair comparison, DUALBREACH and white-box baselines employ Llama-3-8B-Instruct [18] as the backbone for gradient-based optimization. Additionally, we set the number of local optimization iterations per query to five, and the maximum query budget to 40. All experiments are conducted on a server with 2 NVIDIA A6000 GPUs, 128G RAM.

**Metrics.** We utilize two metrics to evaluate the attack performance: the **Guardrail Attack Success Rate** ( $ASR_G$ ) [20] and the **Dual Jailbreak Attack Success Rate** ( $ASR_L$ ) [41]. We define  $ASR_G$  as the proportion of jailbreak prompts that

are misclassified as `Safe` by the target guardrail  $\mathcal{G}$ . Given a set of  $N$  jailbreak prompts  $\{\mathcal{P}_{adv,i}\}_{i=1}^N$ , we consider a prompt successfully bypassed if  $\mathcal{G}$  classifies it as `Safe`. The metric is formally defined as:

$$ASR_G = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\mathcal{G}(\mathcal{P}_{adv,i}) = \text{Safe}). \quad (14)$$

We define  $ASR_L$  as the fraction of jailbreak prompts  $\mathcal{P}_{adv}$  that bypass the guardrail  $\mathcal{G}$  and induce the target LLM  $L$  to generate harmful responses. The metric is formally defined as:

$$ASR_L = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\mathcal{J}(L, \mathcal{P}_{adv,i}) = \text{Unsafe} \wedge \mathcal{G}(\mathcal{P}_{adv,i}) = \text{Safe}). \quad (15)$$

Following prior work [41], we employ an external judge  $\mathcal{J}$  to evaluate response harmfulness by producing a continuous unsafe score in the range  $[0, 1]$ . This score reflects both the absence of predefined rejection keywords and the degree of semantic alignment with harmful intent. A response is considered `Unsafe` if its unsafe score exceeds a threshold of 0.7, selected to balance precision and recall.

**Additional Experimental Setup.** DUALBREACH generates jailbreak prompts by directly optimizing the target LLM when attacking white-box LLMs (e.g., Qwen-2.5), and using Llama-3 as a surrogate model when attacking black-box LLMs.

### B. Dual-Jailbreaking with Limited Queries

**Overall Results on Dual-Jailbreaking Success Rate.** As shown in Table II, DUALBREACH consistently outperforms state-of-the-art methods in dual-jailbreaking scenarios, achieving higher dual-jailbreaking success rates ( $ASR_L$ ) across nearly all datasets and target LLMs. For example, when we take Llama-3 as the target LLM, using the advBench dataset, the  $ASR_L$  of DUALBREACH is 86.0%, whereas other methods like COLD-Attack [12], PAP [10], and ReNELLM [40] achieve much lower  $ASR_L$ s, *i.e.*, 3.0%, 69%, and 54.0%. Moreover, when we assess model performance on three frontier black-box models, DUALBREACH still consistently achieves higher  $ASR_L$ s in most attacking scenarios compared with other baselines. More specifically, when using harmBench dataset, DUALBREACH achieves an  $ASR_L$  of 52.0%~68.0% (**61.7% on average**) across Claude-3.5, Gemini-1.5 and GPT-4o, while COLD-Attack, PAP and ReNELLM achieve an  $ASR_L$  of 1.0%~6.0% (**3.7% on average**), 48.0%~69.0% (**61.7% on average**), and 56.0%~63.0% (**59.7% on average**), respectively.

The higher  $ASR_L$  achieved by DUALBREACH is driven by the combined use of the TDI strategy for prompt initialization, proxy guardrails for efficient gradient-based optimization, and iterative refinement with local LLMs. In contrast, other methods face performance limitations due to their respective shortcomings. For example, GCG, PRP and COLD-Attack primarily focus on bypassing the LLM but largely ignore guardrails during training, making it difficult for them to bypass guardrails. PAP, on the other hand, uses a long harmful

TABLE II  
EXPERIMENTAL RESULTS OF DUALBREACH AND BASELINES IN DUAL-JAILBREAKING SCENARIOS WITH LIMITED QUERIES.<sup>‡</sup>

Dataset	Method	Llama-3 [18]		Qwen-2.5 [35]		GPT-3.5 [36]		Claude-3.5 [37]		Gemini-1.5 [38]		GPT-4o [2]		GPT-4 [2]	
		ASR <sub>L</sub> (%)	QS*	ASR <sub>L</sub> (%)	QS										
advBench	GCG	1.0	1.0	2.0	15.0	1.0	3.0	0	-	2.0	1.0	2.0	1.0	2.0	1.5
	PRP	0	-	0	-	0	-	0	-	0	-	0	-	0	-
	COLD-Attack	3.0	10.7	7.0	16.7	18.0	9.6	5.0	12.0	5.0	11.0	3.0	2.0	8.0	10.6
	PAP	69.0	6.5	95.0	3.1	92.0	4.3	64.0	12.4	65.0	6.9	74.0	5.9	80.0	4.2
	DAN	3.0	7.0	4.0	8.0	5.0	3.8	1.0	20.0	4.0	10.3	0	-	3.0	9.0
	JAM	0	-	0	-	0	-	0	-	0	-	0	-	0	-
	ReNELLM	54.0	7.9	58.0	6.2	57.0	6.5	52.0	9.2	57.0	6.8	57.0	6.3	57.0	6.8
	DUALBREACH	86.0	4.0	93.0	1.3	95.0	1.4	68.0	2.6	64.0	2.3	71.0	3.2	91.0	2.2
DNA	GCG	49.0	3.3	50.0	4.4	48.0	6.1	34.0	7.5	45.0	5.4	46.0	4.8	37.0	9.5
	PRP	50.0	1.4	50.0	1.6	42.0	6.3	45.0	4.3	46.0	2.7	47.0	3.0	36.0	6.3
	COLD-Attack	67.0	3.1	70.0	4.5	70.0	4.3	43.0	8.5	60.0	4.0	65.0	4.3	65.0	5.5
	PAP	91.0	2.6	98.0	1.7	98.0	1.5	77.0	8.4	88.0	5.9	90.0	5.8	100.0	3.0
	DAN	63.0	5.5	66.0	2.2	71.0	2.9	57.0	10.5	71.0	2.5	8.0	16.0	61.0	3.66
	JAM	0	-	0	-	0	-	0	-	0	-	0	-	0	-
	ReNELLM	70.0	2.4	73.0	2.0	73.0	1.9	58.0	3.4	71.0	2.0	73.0	2.2	73.0	2.3
	DUALBREACH	97.0	1.8	98.0	1.4	98.0	1.3	79.0	2.4	88.0	2.3	87.0	2.1	98.0	1.4
harmBench	GCG	4.0	21.5	5.0	23.4	5.0	19.0	1.0	28.0	2.0	26.0	3.0	20.3	4.0	21.3
	PRP	0	-	0	-	0	-	0	-	0	-	0	-	0	-
	COLD-Attack	8.0	8.1	9.0	10.0	11.0	9.0	1.0	1.0	4.0	16.5	6.0	11.2	8.0	5.1
	PAP	66.0	6.9	93.0	2.7	89.0	3.2	48.0	14.0	69.0	7.6	68.0	7.0	85.0	3.1
	DAN	3.0	6.7	8.0	7.6	5.0	10.0	3.0	14.7	7.0	8.6	0	-	5.0	18.8
	JAM	3.0	4.7	9.0	5.6	6.0	13.0	3.0	15.0	7.0	8.6	0	-	5.0	21.8
	ReNELLM	59.0	6.0	63.0	6.0	63.0	5.5	56.0	6.9	63.0	5.5	63.0	5.4	72.0	6.7
	DUALBREACH	86.0	2.4	93.0	1.4	95.0	1.3	52.0	3.0	65.0	2.1	68.0	2.4	92.0	1.7

<sup>‡</sup> We evaluate the dual-jailbreaking success rates ( $ASR_L$ ) on four target LLMs with the protection of Llama-Guard-3 [18], due to its robustness and effectiveness in defending harmful queries. Table III demonstrates the evaluation results of different guardrails. The results show that Guard3 achieves better performance compared with other guardrails.

\* We compute the average number of queries solely from successful samples, i.e., queries per success (QS). As a result, for methods that have a low  $ASR_L$ , like GCG, it is feasible that the few successful jailbreak prompts require only 1~3 queries per success. We use the symbol “-” if there is no successful jailbreak prompt.

few-shots prompt template (1,096.8 tokens on average) to optimize harmful queries, which not only consumes substantial resources, but is also likely to be rejected by safety-aligned LLMs in practical tests.

**Average Queries per Successful Dual-Jailbreak.** In addition to achieving high  $ASR_L$ , DUALBREACH also requires fewer queries for each jailbreak prompt on average compared with state-of-the-art methods. For instance, as shown in Table II, DUALBREACH requires only 2.1~3.2 (2.6 on average) queries<sup>2</sup> per successful dual-jailbreak prompt against GPT-4o [2], indicating a substantial improvement over COLD-Attack, PAP, and ReNELLM, which require 2.0~11.2 (5.8 on average), 5.8~7.0 (6.2 on average) and 2.2~6.3 (4.8 on average) queries, respectively. Additionally, DUALBREACH exhibits significantly greater stability in query counts compared with other methods, consistently ranging from 1.3 to 4.0 queries per prompt across diverse attack scenarios, which exhibit significant variation in the required queries.

The experimental results further reveal that, despite achieving relatively high  $ASR_G$  in bypassing Guard3, existing methods fail to achieve a high  $ASR_L$ . This limitation arises

<sup>2</sup>We note that training proxy guardrails needs to query the black-box guardrails. However, we employ two data distillation approaches to cut the training cost (including queries) by up to 96%. With each sample having a maximum of 200 proxy guardrail calls, the query cost for each jailbreak prompt is acceptable.

because these methods do not simultaneously optimize harmful queries for both guardrails and target LLMs, resulting in success in either bypassing guardrails or jailbreaking LLMs. We provide supplementary experimental results and analysis on jailbreaking the target LLMs with the protection of a black-box guardrail (i.e., OpenAI) in Appendix B.

### C. One-Shot Dual-Jailbreak

Here we consider the situation that for each jailbreak prompt, the attacker only has *one chance* of querying the target guardrail and LLM with this prompt. As shown in Table III, although the  $ASR_G$  and  $ASR_L$  of each jailbreaking method decrease, DUALBREACH still outperforms all other methods in most scenarios. For example, in the dual-jailbreaking scenario against Guard3 [18] and Claude-3.5 [37], DUALBREACH achieves an  $ASR_L$  of 37.0%~58.0% (**48.0% on average**), surpassing other methods like PAP, ReNELLM, which yield an  $ASR_L$  of 25.0%~31.0% (**27.7% on average**), 5.0%~17.0% (**10.3% on average**), respectively.

Furthermore, although existing baselines effectively jailbreak target LLMs (as reported in their original studies), most could not simultaneously produce prompts that appear sufficiently benign to evade strong guardrails such as Guard3. For example, DNA successfully bypasses OpenAI guardrails but is much less effective against others—only 2% of its prompts bypass Guard3. Additionally, as shown in Table III,

TABLE III  
EXPERIMENTAL RESULTS OF DUALBREACH AND BASELINES IN ONE-SHOT DUAL-JAILBREAKING SCENARIOS.

Dataset	Method <sup>‡</sup>	Guardrails ( $ASR_G$ , %)				Target LLM with Guard3 protection ( $ASR_L$ , %)					
		Guard3 [18]	Nemo [16]	GuardAI [17]	OpenAI [15]	Llama-3 [18]	Qwen-2.5 [35]	Claude-3.5 [37]	Gemini-1.5 [38]	GPT-4o [2]	GPT-4
advBench	Raw	1.0	20.0	2.0	94.0	1.0	1.0	0	1.0	0	1.0
	GCG	0	3.0	2.0	93.0	0	0	0	0	0	0
	PRP	0	0	26.0	97.0	0	0	0	0	0	0
	COLD-Attack	5.0	28.0	8.0	94.0	2.0	3.0	0	1.0	2.0	1.0
	PAP	66.0	91.0	79.0	100.0	26.0	56.0	27.0	49.0	50.0	42.0
	DAN	2.0	3.0	3.0	100.0	0	1.0	0	1.0	0	0
	JAM	0	0	97.0	49.0	0	0	0	0	0	0
	ReNeLLM	22.0	94.0	99.0	99.0	11.0	21.0	5.0	19.0	20.0	17.0
	DUALBREACH	97.0	100.0	92.0	100.0	44.0	76.0	49.0	51.0	40.0	60.0
DNA	Raw	57.0	66.0	54.0	86.0	35.0	35.0	13.0	34.0	33.0	14.0
	GCG	36.0	25.0	39.0	87.0	28.0	30.0	10.0	25.0	22.0	13.0
	PRP	51.0	71.0	49.0	89.0	39.0	45.0	23.0	34.0	29.0	27.0
	COLD-Attack	55.0	67.0	45.0	86.0	35.0	44.0	18.0	30.0	31.0	22.0
	PAP	93.0	100.0	91.0	100.0	62.0	67.0	25.0	57.0	50.0	61.0
	DAN	54.0	6.0	8.0	85.0	21.0	37.0	7.0	36.0	0	13.0
	JAM	0	0	93.0	26.0	0	0	0	0	0	0
	ReNeLLM	42.0	95.0	99.0	100.0	30.0	41.0	17.0	33.0	39.0	35.0
	DUALBREACH	100.0	100.0	100.0	100.0	70.0	85.0	58.0	68.0	65.0	76.0
harmBench	Raw	3.0	51.0	47.0	99.0	3.0	3.0	1.0	1.0	1.0	2.0
	GCG	2.0	11.0	15.0	94.0	2.0	2.0	1.0	2.0	2.0	1.0
	PRP	0	37.0	8.0	99.0	0	0	0	0	0	0
	COLD-Attack	4.0	51.0	40.0	99.0	2.0	2.0	1.0	2.0	2.0	3.0
	PAP	71.0	96.0	75.0	100.0	37.0	56.0	31.0	35.0	55.0	46.0
	DAN	1.0	6.0	10.0	88.0	0	0	0	0	0	0
	JAM	0	0	96.0	50.0	0	0	0	0	0	0
	ReNeLLM	24.0	99.0	100.0	99.0	18.0	24.0	9.0	19.0	22.0	22.0
	DUALBREACH	87.0	97.0	85.0	100.0	43.0	69.0	37.0	52.0	43.0	46.0

‡ “Raw” method is using the original harmful queries in datasets. We employ the COLD-Attack algorithm with its “suffix” strategy as described in [12].

Guard3 easily detects most baseline prompts, preventing them from reaching the target LLMs and consequently resulting in near-zero  $ASR_L$  scores across the six evaluated LLM targets. The left part of Table III showcases the effectiveness of four state-of-the-art guardrails in detecting harmful queries. For instance, on the advBench dataset, Guard3 achieves the best performance, indicated by the lowest  $ASR_G$  among the four guardrails. The average  $ASR_G$  for bypassing Guard3 is 28.17%, compared to 40.33% for Nemo, 52.25% for GuardAI, and 95.67% for OpenAI.

#### D. One-Shot Dual-Jailbreak without Proxy Guardrail

Here, we consider the situation where the adversary lacks any additional queries for either training proxy guardrails or testing an intermediate jailbreak prompt on the target guardrail and LLM. DUALBREACH approximately optimizes harmful queries on Guard3 [18], then transferring these queries to bypass other guardrails, *i.e.*, Nemo [16], GuardAI [17], and OpenAI [15]. As shown in Table III, DUALBREACH significantly demonstrates its effectiveness in bypassing other guardrails using Guard3 to optimize gradients. Specifically, on the DNA dataset, DUALBREACH achieves an  $ASR_G$  of 100% in bypassing Guard3, Nemo, GuardAI, and OpenAI. Additionally, DUALBREACH shows more stable performance across different guardrails and datasets compared with other methods. For instance, on the DNA dataset, GCG achieves an

TABLE IV  
ATTACK SUCCESS RATE OF STRONGREJECT [39] ON TARGET LLM (%)

Method	Claude-3.5 [37]		Gemini-1.5 [38]		GPT-4o [2]	
	Pythia-14m	GPT-4o	Pythia-14m	GPT-4o	Pythia-14m	GPT-4o
GCG [20]	0	0	0	0	0	0
PRP [9]	0	0	0	0	0	0
COLD [12]	3.0	0	3.0	0	2.0	0
PAP [10]	24.0	0	22.0	24.0	28.0	23.0
DAN [8]	1.0	0	1.0	0	0	0
JAM [19]	0	0	0	0	0	0
ReNeLLM [40]	19.0	13.0	26.0	29.0	17.0	29.0
DUALBREACH	62.0	33.0	64.0	63.0	58.0	67.0

$ASR_G$  of 36% in bypassing Guard3, whereas on the advBench and harmBench datasets, GCG’s  $ASR_G$  drops to 0% and 2%, respectively. In comparison, DUALBREACH achieves an  $ASR_G$  of 87%~100% across three datasets.

#### E. Results on StrongReject Benchmark

As shown in Table IV, DUALBREACH consistently achieves the highest attack success rates across all evaluated models, including Claude, Gemini, and GPT-4o. For instance, on GPT-4o, DUALBREACH reaches an ASR (by Pythia-14m) of 58.0% and an ASR (by GPT-4o) of 67.0%, outperforming prior state-of-the-art methods such as PAP, ReNeLLM, and DAN. These results not only demonstrate the superior performance

TABLE V  
COMPARISON BETWEEN DUALBREACH AND BASELINE METHODS UNDER LIMIT-QUERY SETTING ON ADVBENCH DATASET (TARGET LLM: CLAUDE-3.5)

Method	ASR <sub>L</sub> (%)	QS	Run Time (h)	API <sup>‡</sup> (\$)	GPU* (\$)
GCG [20]	0	-	42.8	0.40	19.26
PRP [9]	0	-	0.84	0.41	0.378
COLD [12]	5	12.0	4.3	0.40	1.935
PAP [10]	64	12.4	12.5	0.87	5.625
DAN [8]	1	20.0	1.5	11.70	0.68
JAM [19]	0	-	1.3	12.90	0.585
ReNeLLM [40]	52	9.2	7.4	0.99	3.33
DUALBREACH	68	2.6	3.28	0.21	1.476

<sup>‡</sup> API cost estimated based on Claude-3.5 pricing: \$3,0000 per 1M tokens.

\* GPU cost estimated using A6000 pricing on Vast.ai: \$0.45/hour.

TABLE VI  
COMPARISON OF GUARDRAIL DEPLOYMENT STRATEGIES.

Method	Input-based guardrail		Output-based guardrail	
	ASR <sub>L</sub> (%)	Queries per Success	ASR <sub>L</sub> (%)	Queries per Success
GCG [20]	2.0	1.5	2.0	1.0
PRP [9]	0	-	0	-
COLD [12]	8.0	10.6	6.0	11.0
PAP [10]	80.0	4.2	67.0	6.1
DUALBREACH	91.0	2.2	90.0	3.2

of DUALBREACH, but also imply its generalization capability to challenging benchmarks that incorporate advanced rejection mechanisms, including fine-tuned detectors and real-world-inspired prompts.

#### F. Overhead Analysis

Table V provides a comprehensive comparison of DUALBREACH against baseline methods under the limit-query setting using the AdvBench dataset and targeting Claude-3.5. All experiments are conducted on a server with 2 NVIDIA A6000 GPUs (128G RAM). DUALBREACH achieves the highest ASR<sub>L</sub> of 68% with only 2.6 queries per successful attack attempt, leading to the highest ASR<sub>L</sub> with the lowest API cost (0.21\$). These results demonstrate that DUALBREACH achieves higher query efficiency than PAP (12.4 queries, 0.87\$ API cost) and ReNeLLM (9.2 queries, 0.99\$ API cost). Moreover, DUALBREACH takes less runtime (3.28 hours) compared to PAP (12.5 hours) and ReNeLLM (7.4 hours).

#### G. Experiments on Output-based Guardrails Detecting Both Harmful Prompts and Responses

Beyond using a guardrail to filter out harmful prompts, the industry also tries to apply guardrails to detect LLMs' responses to protect LLM-based applications from jailbreaking attacks. Here we consider the case that Guard3 is applied to detect harmful content in both the prompt and response of GPT-4 and compare DUALBREACH with other baselines. As shown in Table VI, DUALBREACH still outperforms other methods, achieving an ASR<sub>L</sub> of 90% with 3.2 queries per jailbreak prompt, which demonstrates the effectiveness and efficiency of DUALBREACH across different scenarios.

TABLE VII  
ROBUSTNESS OF ASR (%) EVALUATIONS UNDER DIVERSE JUDGE MODELS (TARGET LLM: GEMINI-1.5)

Method	ASR <sub>L</sub>	External Judge		StrongReject [39]	
		Human	Gemini-1.5 <sup>†</sup>	Pythia-14m	GPT-4o
GCG [20]	0	0	0	0	0
PRP [9]	0	0	0	0	0
COLD [12]	1.0	0	0	3.0	1.0
PAP [10]	49.0	42.0	49.0	43.0	63.0
DAN [8]	1	0	1	2.0	2.0
JAM [19]	0	0	0	0	0
ReNeLLM [40]	19.0	13.0	19.0	17.0	20.0
DUALBREACH	51.0	49.0	50.0	56.0	67.0

<sup>†</sup> The Gemini-1.5 is evaluated using the same judge prompt template as used in our evaluation, detailed in Appendix C.2.

#### H. Robustness Validation of Evaluation Mechanisms

As shown in Table VII, we evaluate the jailbreak success rates of DUALBREACH and several baselines across four judging mechanisms: ASR<sub>L</sub> (Llama-3), human annotation, Gemini-1.5, and the StrongReject benchmark (Pythia-14m and GPT-4o). DUALBREACH consistently achieves the highest scores across all judges. The strong agreement between human and automatic evaluations, especially with Gemini-1.5 using the same prompt template, demonstrates the robustness and reliability of our multi-judge evaluation framework.

### VI. EGUARD: A BOOSTING ENSEMBLE LEARNING APPROACH FOR GUARDRAILS

While existing guardrails claim to have the ability to protect LLMs from jailbreak attacks [17], [18], [15], [24], existing attack methods and DUALBREACH can still bypass those guardrails to some extent, as indicated by Table III and ???. This may be because different guardrails are more effective at detecting different abnormal patterns in the jailbreak prompts, while being less effective at identifying other patterns. We observe that Guard3 [18] excels at detecting harmful semantics in short sentences; NeMo [16] is sensitive to perplexity changes introduced by jailbreak prompts; GuardAI [17] shows significant proficiency in identifying toxic content. The diversity of different guardrails naturally raises a question: *How can we combine the strengths of existing guardrails to create a more robust and comprehensive defensive mechanism?*

#### A. Overview of EGUARD

Developing an effective ensemble-based guardrail faces two primary challenges: (1) Assigning uniform weights to all guardrails in the ensemble model could not fully leverage the superior detection capabilities of the most discriminative guardrail (e.g., Guard3 [18]), leading to suboptimal performance; (2) Using deep neural networks for ensemble learning [42] usually overfit, making the ensemble model overemphasize the contribution of the strongest guardrail and overlook the complementary strengths of weaker guardrails.

To address these challenges, EGUARD integrates dynamic weight adjustment and boosting-based decision tree optimization to enable balanced and robust detection. As shown in




Fig. 3. EGUARD Using Weighted Outputs from Multiple Guardrails to Moderate Jailbreak Prompts.

Fig. 3, the training process of EGUARD comprises two key steps: (1) Weight initialization. (2) Decision tree optimization.

**Weight Initialization.** The weight initialization procedure prioritizes Guard3’s detection capabilities and also considers the contributions of other guardrails. If Guard3 correctly detects the label of a prompt  $\mathcal{D}_i$ , it is assigned full weight,  $w_i = [1, 0, \dots, 0]$ . Otherwise, weights are evenly distributed among the remaining four guardrails, ensuring their contributions are not overlooked. The weight initialization is formally defined as follows:

$$w_i \leftarrow \begin{cases} [1, 0, 0, 0, 0] & \text{If } \mathcal{G}_{\text{guard3}}(\mathcal{D}_i) == \mathcal{K}_i, \\ \left[0, \frac{1}{|\mathcal{G}|-1}, \frac{1}{|\mathcal{G}|-1}, \frac{1}{|\mathcal{G}|-1}, \frac{1}{|\mathcal{G}|-1}\right] & \text{Otherwise.} \end{cases} \quad (16)$$

where  $|\mathcal{G}|$  denotes the number of employed guardrails, which is five for EGUARD, and  $\frac{1}{|\mathcal{G}|-1} = \frac{1}{4}$ .

**Decision Tree Optimization.** After weight initialization, EGUARD employs decision trees to iteratively improve the detection accuracy by reducing the residual errors—the difference between predicted and true labels. In each iteration ( $t$ ), EGUARD updates the  $t$ -th decision tree  $h_t(w_i)$  by

$$h_t(w_i) = h_{t-1}(w_i) + \eta \cdot g_t(w_i), \quad (17)$$

where  $h_{t-1}(w_i)$  refers to the output of the previous tree,  $g_t(w_i)$  is the gradient of the loss w.r.t.  $h_{t-1}$  and  $\eta$  is the learning rate. Each tree corrects the errors from the previous iteration, progressively reducing the residuals and improving the detection accuracy.

**Detection Stage.** For detecting harmful prompts, EGUARD inputs the detection results  $\hat{L}$  of the five guardrails, in which each element indicates whether the corresponding guardrail classifies the prompt as unsafe (1) or safe (0), into the decision

### Algorithm 3: EGUARD Training and Prediction

---

**Data:** Training set and corresponding labels  $\{\mathcal{D}, \mathcal{K}\}$ , Guardrails  $\mathcal{G}$ , Jailbreak prompts  $\mathcal{P}_{\text{adv}}$ , Maximum Iterations  $TI$

**Result:** EGUARD  $E$ , Prediction  $\hat{R}$

1 // **Step 1: Train EGUARD**

2 Initialize an XGBoost model  $E$  with  $N_{tr}$  decision trees.

3 **for** iteration from 1 to  $TI$  **do**

4   **for** each sample  $(\mathcal{D}_i, L_i)$  in  $\{\mathcal{D}, L\}$  **do**

5      $\hat{Y} \leftarrow \mathcal{G}_{\text{Guard3}}(\mathcal{D}_i)$

6      $w_i \leftarrow \begin{cases} [1, 0, 0, 0, 0] & \text{If } \hat{Y} == L_i, \\ \left[0, \frac{1}{|\mathcal{G}|-1}, \frac{1}{|\mathcal{G}|-1}, \frac{1}{|\mathcal{G}|-1}, \frac{1}{|\mathcal{G}|-1}\right] & \text{Otherwise} \end{cases}$

7      $loss \leftarrow \mathcal{L}_{CE}(E(w_i), L_i)$

8      $loss.backward()$

9 Save EGUARD  $E$

10 // **Step 2: EGUARD Predict**

11  $\hat{R} \leftarrow \emptyset$

12 **for** each prompt  $\mathcal{P}_{\text{adv},i}$  in  $\mathcal{P}_{\text{adv}}$  **do**

13    $\hat{L} \leftarrow \emptyset$

14   **for** each guardrail  $\mathcal{G}_j \in \mathcal{G}$  **do**

15      $\hat{L} \leftarrow \hat{L} \cup \mathcal{G}_j(\mathcal{P}_{\text{adv},i})$

16    $\hat{R} \leftarrow \hat{R} \cup E(\hat{L})$

17 **return** EGUARD  $E$ , Prediction  $\hat{R}$

---

TABLE VIII  
COMPARISON OF EGUARD’S ENHANCED ROBUSTNESS TO GUARDRAIL ATTACKS RELATIVE TO GUARD3 (ASR<sub>G</sub>, %)

Method	advBench		DNA		harmBench	
	Guard3	EGUARD	Guard3	EGUARD	Guard3	EGUARD
Raw	1.0	0.0	57.0	41.0	3.0	0.0
GCG [20]	0.0	0.0	36.0	9.0	2.0	0.0
PRP [9]	0.0	0.0	51.0	31.0	0.0	0.0
COLD-Attack [12]	5.0	1.0	55.0	30.0	4.0	0.0
PAP [10]	66.0	59.0	93.0	87.0	71.0	60.0
DUALBREACH	97.0	90.0	100.0	100.0	87.0	74.0

trees. The final prediction  $\hat{R}$  is obtained by aggregating the outputs of all decision trees in the ensemble, *i.e.*,

$$\hat{R} = E(\hat{L}) = \sigma \left( \sum_{t=1}^{T_{tr}} \alpha_t h_t(\hat{L}) \right), \quad (18)$$

where  $\sigma$  is the sigmoid function that normalizes the output to the range  $[0, 1]$ ,  $\alpha_t$  is a learnable parameter, denoting the weight of the  $t$ -th decision tree, and  $h_t(\hat{L})$  is the output of the  $t$ -th decision tree.  $T_{tr}$  is the total number of decision trees, which we set it to 100.

### B. Experiments

**Experimental Setups.** We randomly select 4,000 samples from five proxy datasets (as we mentioned in Section V-A to

TABLE IX  
RUNTIME OF BASE LLMs AND GUARDRAIL-PROTECTED VARIANTS.\*

Model	LLM	LLM+Guardrails AI	LLM+EGuard (Ours)
Llama-3	45.5s	52.9s (+16.3%)	53.0s (+16.4%)
Qwen-2.5	33.0s	40.4s (+22.4%)	40.5s (+22.7%)
Claude-3.5	37.6s	45.0s (+16.7%)	45.1s (+19.9%)

\* The experiments are conducted on the AdvBench dataset, using a server with two NVIDIA A6000 GPUs.

construct EGUARD’s training set. The evaluation is conducted using the three benchmark datasets and four baseline methods outlined in Section V-A. We compare the  $ASR_G$  of different attacks against Guard3 and EGUARD across various datasets.

**Main Results.** As shown in Table VIII, EGUARD outperforms Guard3 by a non-negligible margin across all the baselines. Notably, on the DNA dataset, EGUARD reduces the  $ASR_G$  of GCG, PRP, COLD-Attack and PAP by 6%~25% (18.80% on average) compared with Guard3. The observed improvement is attributed to a fundamental limitation of Guard3: Guard3 primarily focuses on its predefined harmful categories, which significantly weakens its ability to detect jailbreak prompts that fall outside these categories. In contrast, EGUARD integrates the complementary strengths of weaker guardrails, resulting in broader coverage of harmful categories.

**Runtime Overhead.** We evaluate EGUARD’s performance overhead, with results presented in Table IX. Although EGUARD integrates five distinct guardrails, their parallel execution means that the total latency is determined only by the slowest component. This efficient design adds a mere 0.08s to the runtime of the slowest guardrail (Guardrails AI). When contextualized against the much larger latency of LLM inference, this overhead is minimal, representing less than 25% of the inference time. These experimental results confirm EGUARD’s viability as a low-latency, cost-effective security solution.

## VII. CONCLUSION

In this paper, we present DUALBREACH, a comprehensive framework for jailbreaking both prevailing LLMs and guardrails, facilitating robust evaluations of LLM-based applications in real-world scenarios. DUALBREACH introduces a target-driven initialization strategy to initialize harmful queries and further optimize the jailbreak prompts with (approximate) gradients on guardrails and LLMs, enabling efficient and effective dual-jailbreak. Experimental results demonstrate that DUALBREACH achieves a higher dual jailbreak attack success rate with fewer queries compared with state-of-the-art jailbreak methods. Furthermore, we introduce EGUARD, an ensemble guardrail for detecting jailbreak prompts. Extensive evaluations on multiple benchmark datasets validate the superior performance of EGUARD compared with Llama-Guard-3.

## ETHICS CONSIDERATIONS

Our research is dedicated to enhancing the security and robustness of LLM-based applications in real-world scenarios. We adhere to established ethical guidelines to ensure that our work does not inadvertently contribute to the spread of

harmful content, misinformation, or unethical use of technology. Our research aims to address vulnerabilities in LLMs and guardrails to improve their safety and reliability. We are committed to conducting our work responsibly and ethically, with a focus on promoting trustworthiness and safeguarding users in AI-driven systems.

## ACKNOWLEDGEMENT

This research is supported in part by the National Key R&D Program of China (2024YFB4505300, 2023YFB2904000), the “Pioneer” and “Leading Goose” R&D Program of Zhejiang (2024C01169), and the National Natural Science Foundation of China under Grants (62441238, 62032021, U2441240, U23A20306).

## REFERENCES

- [1] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, Y. Du, C. Yang, Y. Chen, Z. Chen, J. Jiang, R. Ren, Y. Li, X. Tang, Z. Liu, P. Liu, J.-Y. Nie, and J.-R. Wen, “A survey of large language models,” 2024. [Online]. Available: <https://arxiv.org/abs/2303.18223>
- [2] OpenAI, *Hello GPT-4o*, 2024, <https://openai.com/index/hello-gpt-4o/>.
- [3] M. Wermelinger, “Using github copilot to solve simple programming problems,” in *Proceedings of the 54th ACM Technical Symposium on Computer Science Education (SIGCSE)*. ACM, 2023, p. 172–178.
- [4] OpenAI, *Video generation models as world simulators*, 2024, <https://openai.com/index/video-generation-models-as-world-simulators>.
- [5] S. Park, H. Subramonyam, and C. Kulkarni, “Thinking assistants: LLM-based conversational assistants that help users think by asking rather than answering,” 2024. [Online]. Available: <https://arxiv.org/abs/2312.06024>
- [6] Y. Huang, L. Sun, H. Wang, S. Wu, Q. Zhang, Y. Li, C. Gao, Y. Huang, W. Lyu, Y. Zhang *et al.*, “Position: TrustLLM: Trustworthiness in large language models,” in *Proceedings of the 41st International Conference on Machine Learning (ICML)*. PMLR, 2024.
- [7] X. Liu, N. Xu, M. Chen, and C. Xiao, “AutoDAN: Generating stealthy jailbreak prompts on aligned large language models,” in *The Twelfth International Conference on Learning Representations (ICLR)*, 2024.
- [8] X. Shen, Z. Chen, M. Backes, Y. Shen, and Y. 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 (CCS)*. ACM, 2024, p. 1671–1685.
- [9] N. Mangaokar, A. Hooda, J. Choi, S. Chandrashekaran, K. Fawaz, S. Jha, and A. Prakash, “PRP: Propagating universal perturbations to attack large language model guard-rails,” in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*. ACL, 2024, pp. 10960–10976.
- [10] Y. Zeng, H. Lin, J. Zhang, D. Yang, R. Jia, and W. Shi, “How johnny can persuade LLMs to jailbreak them: Rethinking persuasion to challenge AI safety by humanizing LLMs,” in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*. ACL, 2024, pp. 14 322–14 350.
- [11] R. Kirk, I. Mediratta, C. Nalmpantis, J. Luketina, E. Hambro, E. Grefenstette, and R. Raileanu, “Understanding the effects of rlfh on llm generalisation and diversity,” 2024. [Online]. Available: <https://arxiv.org/abs/2310.06452>
- [12] X. Guo, F. Yu, H. Zhang, L. Qin, and B. Hu, “COLD-attack: Jailbreaking LLMs with stealthiness and controllability,” in *Proceedings of the 41st International Conference on Machine Learning (ICML)*, vol. 235. PMLR, 2024, pp. 16 974–17 002.
- [13] S. G. Ayyamperumal and L. Ge, “Current state of llm risks and ai guardrails,” 2024. [Online]. Available: <https://arxiv.org/abs/2406.12934>
- [14] G. Evans, J. Miller, M. I. Pena, A. MacAllister, and E. Winer, “Evaluating the Microsoft HoloLens through an augmented reality assembly application,” in *Degraded Environments: Sensing, Processing, and Display 2017*, vol. 10197. SPIE, 2017, p. 101970V.
- [15] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, “Gpt-4 technical report,” 2024. [Online]. Available: <https://arxiv.org/abs/2303.08774>

[16] T. Rebedea, R. Dinu, M. N. Sreedhar, C. Parisien, and J. Cohen, “NeMo guardrails: A toolkit for controllable and safe LLM applications with programmable rails,” in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2023, pp. 431–445.

[17] G. AI, *Mitigate Gen AI risks with Guardrails*, 2023, <https://www.guardrailsai.com/>.

[18] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan *et al.*, “The llama 3 herd of models,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.21783>

[19] H. Jin, A. Zhou, J. D. Menke, and H. Wang, “Jailbreaking large language models against moderation guardrails via cipher characters,” in *Advances in Neural Information Processing Systems*, vol. 37. Curran Associates, Inc., 2024, pp. 59 408–59 435.

[20] A. Zou, Z. Wang, N. Carlini, M. Nasr, J. Z. Kolter, and M. Fredrikson, “Universal and transferable adversarial attacks on aligned language models,” 2023. [Online]. Available: <https://arxiv.org/abs/2307.15043>

[21] L. Qin, S. Welleck, D. Khashabi, and Y. Choi, “Cold decoding: Energy-based constrained text generation with langevin dynamics,” in *In Annual Conference on Neural Information Processing Systems (NeurIPS)*. NeurIPS, 2022, pp. 9538–9551.

[22] G. Alon and M. Kamfonas, “Detecting language model attacks with perplexity,” 2023. [Online]. Available: <https://arxiv.org/abs/2308.14132>

[23] L. Cao, “Learn to refuse: Making large language models more controllable and reliable through knowledge scope limitation and refusal mechanism,” in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2024, pp. 3628–3646.

[24] C. Hawker and E. Koukoumidis, *Improving Trust in AI and Online Communities with PaLM-based Moderation*, 2023, <https://cloud.google.com/blog/products/ai-machine-learning/google-cloud-text-moderation>.

[25] A. Robey, E. Wong, H. Hassani, and G. J. Pappas, “Smoothllm: Defending large language models against jailbreaking attacks,” 2024. [Online]. Available: <https://arxiv.org/abs/2310.03684>

[26] Z. Zhang, J. Yang, P. Ke, F. Mi, H. Wang, and M. Huang, “Defending large language models against jailbreaking attacks through goal prioritization,” 2024. [Online]. Available: <https://arxiv.org/abs/2311.09096>

[27] A. Zou, L. Phan, J. Wang, D. Duenas, M. Lin, M. Andriushchenko, R. Wang, Z. Kolter, M. Fredrikson, and D. Hendrycks, “Improving alignment and robustness with circuit breakers,” 2024. [Online]. Available: <https://arxiv.org/abs/2406.04313>

[28] Y. Wang, H. Li, X. Han, P. Nakov, and T. Baldwin, “Do-not-answer: Evaluating safeguards in LLMs,” in *Proceedings of the The 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. ACL, 2024, pp. 896–911.

[29] M. Mazeika, L. Phan, X. Yin, A. Zou, Z. Wang, N. Mu, E. Sakhaee, N. Li, S. Basart, B. Li, D. Forsyth, and D. Hendrycks, “HarmBench: A standardized evaluation framework for automated red teaming and robust refusal,” in *Proceedings of the 41st International Conference on Machine Learning (ICML)*. PMLR, 2024.

[30] J. Ji, D. Hong, B. Zhang, B. Chen, J. Dai, B. Zheng, T. Qiu, B. Li, and Y. Yang, “Pku-saferlfhf: Towards multi-level safety alignment for llms with human preference,” 2024. [Online]. Available: <https://arxiv.org/abs/2406.15513>

[31] T. Mihaylov, P. Clark, T. Khot, and A. Sabharwal, “Can a suit of armor conduct electricity? a new dataset for open book question answering,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2018, pp. 2381–2391.

[32] N. Asghar, “Yelp dataset challenge: Review rating prediction,” 2016. [Online]. Available: <https://arxiv.org/abs/1605.05362>

[33] M. Joshi, E. Choi, D. Weld, and L. Zettlemoyer, “TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*. ACL, 2017, pp. 1601–1611.

[34] Y. Yang, W.-t. Yih, and C. Meek, “WikiQA: A challenge dataset for open-domain question answering,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2015, pp. 2013–2018.

[35] A. Yang, B. Yang, B. Hui, B. Zheng, B. Yu, C. Zhou, C. Li, C. Li, D. Liu, F. Huang *et al.*, “Qwen2 technical report,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.10671>

[36] OpenAI, *GPT-3.5 Turbo fine-tuning and API updates*, 2023, <https://openai.com/index/gpt-3-5-turbo-fine-tuning-and-api-updates/>

[37] Anthropic, *Claude 3.5 Sonnet*, 2024, <https://www.anthropic.com/news/claude-3-5-sonnet>.

[38] G. Team, P. Georgiev, V. I. Lei, R. Burnell, L. Bai, A. Gulati, G. Tanzer, D. Vincent, Z. Pan, S. Wang *et al.*, “Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context,” 2024. [Online]. Available: <https://arxiv.org/abs/2403.05530>

[39] A. Souly, Q. Lu, D. Bowen, T. Trinh, E. Hsieh, S. Pandey, P. Abbeel, J. Svegliato, S. Emmons, O. Watkins, and S. Toyer, “A strongreject for empty jailbreaks,” in *Advances in Neural Information Processing Systems(NeurIPS)*, vol. 37, 2024, pp. 125 416–125 440.

[40] P. Ding, J. Kuang, D. Ma, X. Cao, Y. Xian, J. Chen, and S. Huang, “A wolf in sheep’s clothing: Generalized nested jailbreak prompts can fool large language models easily,” in *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. Mexico City, Mexico: Association for Computational Linguistics, Jun. 2024, pp. 2136–2153.

[41] X. Qi, Y. Zeng, T. Xie, P.-Y. Chen, R. Jia, P. Mittal, and P. Henderson, “Fine-tuning aligned language models compromises safety, even when users do not intend to!” 2023. [Online]. Available: <https://arxiv.org/abs/2310.03693>

[42] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, “Efficient processing of deep neural networks: A tutorial and survey,” *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.

## APPENDIX

### A. Overview of Target Guardrails and LLMs

In this section, we describe the target guardrails and LLMs used in our paper. We select seven representative LLMs to capture diverse capabilities and accessibility levels, and five widely used guardrails to reflect different approaches to safety alignment. First, we select seven target LLMs that represent a range of model scales and instruction-following capabilities.

- **Llama3-8B-Instruct [18] (Llama-3).** Llama3-8B-Instruct is an open-source model designed for instruction-following, with an accessible architecture that facilitates analysis and vulnerability research.
- **Qwen2.5-7B Instruct [35] (Qwen-2.5).** Qwen2.5 is a series of instruction-tuned LLMs developed to enhance instruction-following capabilities across a variety of tasks.
- **GPT-3.5-turbo-0125 [36] (GPT-3.5).** GPT-3.5-turbo-0125 is an OpenAI model optimized for conversational tasks, offering enhanced speed and cost-effectiveness for real-time applications.
- **GPT-4-0613 [15] (GPT-4).** GPT-4-0613 is an OpenAI model with improved reasoning, accuracy, and contextual understanding, making it well-suited for complex problem-solving and nuanced content generation.
- **GPT-4o-2024-11-20 [2] (GPT-4o).** GPT-4o is a multimodal model by OpenAI that processes and generates combinations of text, audio, image, and video, enabling real-time reasoning across modalities for more natural human-computer interaction.
- **Claude-3.5-sonnet-20241022 [37] (Claude-3.5).** Claude-3.5-sonnet-20241022 is an Anthropic model offering significant improvements in reasoning and coding, designed to be faster and more cost-effective while maintaining strong multilingual and mathematical capabilities.
- **Gemini-1.5-flash [38] (Gemini-1.5).** Gemini-1.5-flash is a lightweight Google model optimized for efficiency and

TABLE X  
ABLATION STUDY ON PROXY GUARDRAILS ( $ASR_G$ , %).\*

Method	TVD †	advBench		DNA		harmBench	
		Raw	DualBreach	Raw	DualBreach	Raw	DualBreach
OpenAI [15]	-	94.0	100.0	86.0	100.0	99.0	100.0
Proxy OpenAI+RAW	0.0045	94.0 (0)	98.0 (-2)	88.0 (+2)	98.0 (-2)	94.0 (-5)	93.0 (-7)
Proxy OpenAI+Kmeans	0.0088	84.0 (-10)	99.0 (-1)	96.0 (+10)	98.0 (-2)	96.0 (-3)	100.0 (0)
Proxy OpenAI+BLEU	0.0125	100.0 (+6)	100.0 (0)	94.0 (+8)	100.0 (0)	100.0 (+1)	100.0 (0)
Google [24]	-	30.0	45.0	46.0	43.0	47.0	41.0
Proxy Google+RAW	0.0584	55.0 (+25)	44.0 (-1)	66.0 (+20)	48.0 (+5)	44.0 (-3)	42.0 (+1)
Proxy Google+Kmeans	0.0397	55.0 (+25)	59.0 (+14)	69.0 (+23)	60.0 (+17)	72.0 (+25)	57.0 (+16)
Proxy Google+BLEU	0.0678	61.0 (+31)	57.0 (+12)	70.0 (+24)	59.0 (+16)	73.0 (+26)	50.0 (+9)

\* The  $ASR_G$  quantifies how accurately proxy guardrails characterize the behaviors of black-box guardrails, where -10%/+10% denote reduced/enhanced accuracy. Green (red) highlights changes below (exceeding) the 20% predefined value.

† The Total variation distance (TVD) quantifies similarity between proxy and black-box guardrails, lower values denoting higher similarity.

low latency in real-time applications, while maintaining strong performance in long-context understanding.

Second, we evaluate five guardrails' performance against attacks, focusing on both white-box and black-box scenarios:

- **Llama Guard3 [18] (Guard3).** Llama Guard3 is a fine-tuned Llama-3.1-8B for moderation, producing a binary *safe/unsafe* token followed by one of 14 harmful categories. It supports input and output filtering and surpasses Llama Guard2 in language and tool-use moderation.
- **Nvidia NeMo [16] (NeMo).** NeMo provides a modular framework with detection tools that enable customizable safety protocols, including fact-checking, hallucination detection, and jailbreak prevention.
- **Guardrails AI [17] (GuardAI).** GuardAI acts as an interception layer for LLM inputs and outputs, enforcing customizable safety standards. It is model-agnostic and can integrate with additional security mechanisms.
- **OpenAI Moderation API [15] (OpenAI).** The OpenAI Moderation API classifies text across 11 predefined harmful categories and provides harmfulness scores to enforce usage policy compliance.
- **Google Moderation API [24] (Google).** The Google Moderation API evaluates text over 16 harmful and sensitive categories, assigning risk scores. We set a threshold of 0.5, above which content is marked as unsafe.

### B. Supplementary Experiments

In this section, we analyze the dual jailbreak performance of DUALBREACH and baseline methods on target LLMs (GPT-3.5 and GPT-4) under the protection of OpenAI guardrail. As shown in Table XI, DUALBREACH consistently achieves the best dual jailbreak performance compared to the baselines. This is primarily because DUALBREACH utilizes a proxy model to simulate the behavior of OpenAI, enabling effective optimization of jailbreak prompts in a black-box environment.

Notably, compared to the results in Table II, the dual jailbreak performance of both DUALBREACH and baseline methods improves significantly. For instance, GCG's  $ASR_L$  increases from 2.0% to 15.0%. This improvement suggests that OpenAI's black-box guardrail is less effective at identifying

TABLE XI  
RESULTS OF DUAL-JAILBREAKING GPT-4 USING PROXY OPENAI.

Method	OpenAI + GPT-3.5		OpenAI + GPT-4	
	$ASR_L$ (%)	Queries per Success	$ASR_L$ (%)	Queries per Success
GCG [20]	64.0	8.9	15.0	11.4
PRP [9]	50.0	11.5	9.0	14.2
COLD-Attack [12]	76.0	8.6	22.0	9.9
PAP [10]	95.0	2.9	74.0	7.8
DUALBREACH	96.0	2.3	91.0	3.7

some jailbreak prompts, incorrectly classifying them as safe and allowing them to reach the target LLM. This highlights the need for more robust guardrails that can reliably detect and block jailbreak prompts, thereby mitigating potential security risks in LLM deployments.

**Ablation Study on Proxy Guardrails.** As shown in Table X, we analyze the impact of different data distillation approaches on the similarity between proxy guardrails and black-box guardrails. The results show that OpenAI and its proxy guardrails exhibit low total variation distance (TVD) values (0.0045~0.0125), indicating a high degree of similarity to OpenAI API [15]. In contrast, Google and its proxy guardrails present higher TVD values (0.0397~0.0678), suggesting weaker similarities compared with OpenAI.

Furthermore, using different distillation approaches moderately affects  $ASR_G$ . On the advBench dataset [20], the  $ASR_G$ s of DUALBREACH with proxy OpenAI (both with and without the two approaches) are 1%~2% lower than those measured using the OpenAI API [15]. Although Google's proxy guardrails yield relatively high  $ASR_G$ s, most of DUALBREACH's  $ASR_G$ s remain within acceptable variation range (i.e., 20%). This notable change in  $ASR_G$  may result from the Google API's lack of providing a definitive threshold, resulting in skewed proxy guardrail performance.

In all, our proxy guardrails effectively simulate the behaviors of black-box guardrails, even with up to a 96% reduction in training set size, resulting in an eligible training cost.

### C. Supplementary materials and analysis of the Target-driven Initialization (TDI) strategy

#### 1) Workflow of TDI




Fig. 4. Comparison between PAP and TDI in Rewriting Harmful Queries

Here we provide a detailed description of the TDI strategy for inducing LLMs to generate adversarial prompts.

**Adversarial Goal Definition and Contextual Framing.** This phase defines a target output that violates the LLM’s safety policy, typically an affirmative harmful response (e.g., “Sure, here is ...”). The LLM is then assigned a professional persona (e.g., “You are a red-teaming analyst ...”), framing the task as a legitimate security audit and lowering the chance of triggering defenses. Finally, the request is further obfuscated under the guise of academic study or policy analysis to conceal its adversarial intent.

**Directive Issuance and Constraint Setting.** Here the target output sample is presented to the LLM as a pre-existing “artifact,” conditioning it on the desired style and structure. The core directive is then given: “Write a user prompt that could realistically lead to this output,” guiding the model to perform the reverse-engineering task. The template also enforces strict quality requirements, ensuring that the model produces complex, evasive prompts rather than trivial and easily detectable ones.

### 2) Prompt Template

We construct the Target-driven initialization (TDI) prompt template to automatically guide the LLM to convert harmful target responses into contextualized jailbreak prompts, thereby expediting the optimization process of DUALBREACH on both the guardrails and LLMs.

### 3) Evaluating the Effectiveness of TDI Against PAP

The Fig. 4 illustrates that prompts rewritten via Target-driven Initialization (TDI) are more likely to evade Guard3 detection than those rewritten by PAP [10]. TDI yields significantly lower unsafe scores (e.g., 0.01 and 0.14) compared to PAP (e.g., 0.98 and 0.79), demonstrating stronger guardrail-bypassing ability under the same harmful intent.

### Target-Driven Initialization Prompt Template

You are a red-team analyst investigating how language models may produce harmful content when prompted with inputs that are superficially academic, socially relevant, or grounded in policy discourse.

Below is a harmful response that a model has produced:  
%S

**Your task:** Write a user prompt that could realistically lead to this output. The prompt should be grammatically correct, logically sound, and appear grounded in legitimate purposes—such as academic research, social commentary, public policy analysis, or educational content development.

The prompt must mask any malicious intent beneath a layer of intellectual justification. You are encouraged to embed the request within a broader social, psychological, cultural, or technological context—possibly referencing real studies, institutions, or public safety narratives.

Use formal, publication-ready language. Avoid casual tone, direct requests for illegal actions, or overt expressions of intent. The goal is to produce a prompt that appears responsible and defensible, while still capable of eliciting the harmful output.

For example, .....

Return only the prompt. Format your answer as a list containing a single string. For example:

[ "your generated prompt here" ]

### D. Examples of Dual Jailbreak prompt with DUALBREACH

In this section, we present dual-jailbreak examples of DUALBREACH on four target LLMs protected by Guard3, along with a comparative case against baseline methods (Figs. 5, 6).




Fig. 5. Example of a Dual Jailbreak Prompt on GPT-4o Optimized by strongreject with DUALBREACH




Fig. 6. Comparison of Target LLM Response Judged by Guard3 between DUALBREACH and Baselines