

PANDAGUARD: A Systematic Framework for Evaluating LLM Safety Against Jailbreaking Attacks

Anonymous ACL submission

Abstract

Large language models (LLMs) have achieved remarkable capabilities but remain vulnerable to adversarial prompts known as jailbreaks, which can bypass safety alignment and elicit harmful outputs. Despite growing efforts in LLM safety research, existing evaluations are often fragmented, focused on isolated attack or defense techniques, and lack systematic, reproducible analysis. In this work, we introduce PANDAGUARD, a unified and modular framework that models LLM jailbreak safety as a multi-agent system comprising attackers, defenders, and judges. Our framework implements 19 attack methods and 12 defense mechanisms, along with multiple judgment strategies, all within a flexible plugin architecture supporting diverse LLM interfaces, multiple interaction modes, and configuration-driven experimentation that enhances reproducibility and practical deployment. Built on this framework, we develop PANDABENCH, a comprehensive benchmark that evaluates the interactions between these attack/defense methods across 49 LLMs and various judgment approaches, requiring over 3 billion tokens to execute. Our extensive evaluation reveals key insights into model vulnerabilities, defense cost-performance trade-offs, and judge consistency. We find that no single defense is optimal across all dimensions and that judge disagreement introduces nontrivial variance in safety assessments.

1 Introduction

Large Language Models (LLMs), including architectures such as GPT (Brown et al., 2020), Llama (Grattafiori et al., 2024), Qwen (Yang et al., 2024), and Gemini (Google DeepMind, 2025), have achieved state-of-the-art performance across a wide range of natural language understanding and generation tasks. Their rapid deployment in real-world applications—from content creation and

customer service to education and software development (Ma et al., 2024; Topsakal and Akinci, 2023)—highlights their transformative potential. However, as LLMs become increasingly embedded in safety-critical systems, ensuring their robustness and alignment has emerged as a paramount concern (Weidinger et al., 2022; Wang et al., 2024a; Yao et al., 2024; Zhou et al., 2024b; Zeng et al., 2025; Dai et al., 2024).

A particularly acute threat to LLM safety is *jailbreaking*—a class of adversarial attacks in which carefully engineered prompts circumvent alignment constraints and elicit harmful, biased, or unethical outputs (Chu et al., 2024; Zou et al., 2023b; Liu et al., 2023; Yi et al., 2024). Successful jailbreaks can trigger toxic language, misinformation, or even illegal instructions (Xie et al., 2024a; Yuan et al., 2024b; Shen et al., 2025), undermining the guardrails built into state-of-the-art systems. Accordingly, the development of robust defenses and rigorous evaluation protocols for LLM jailbreak resistance has become an urgent research priority.

Despite valuable progress, current jailbreak evaluation approaches exhibit three key limitations. First, existing work often isolates individual components—such as attacks (Zou et al., 2023b; Liu et al., 2024b; Chao et al., 2023) or defenses (Xie et al., 2023; Wang et al., 2024b; Shen et al., 2025; Ji et al., 2024b)—without capturing their systemic interplay. Second, there is a lack of standardized benchmarking practices: evaluation protocols, datasets, and metrics remain fragmented (Zhang et al., 2025; Xu et al., 2024a), which hinders reproducibility and fair comparison. Third, most evaluations are conducted on a limited scale, covering only a small subset of models, threats, or response evaluators (Zhang et al., 2024a; Chao et al., 2024). Moreover, critical factors such as computational cost, defense scalability, and the reliability of safety judges are often overlooked (Chang et al., 2024; Biarese, 2022).

Table 1: Comparison of PANDAGUARD with existing LLM safety evaluation frameworks.

Framework	#Attacks	#Defenses	#Judges	#Models	#LLM Interface
JAILJUDGE (Liu et al., 2024a)	2	3	18	4	4 (HF, OpenAI, Gemini, Claude)
EasyJailbreak (Zhou et al., 2024c)	12	0	7	10	4 (HF, OpenAI, kimi, wenxinyiyan)
AISafetyLab (Zhang et al., 2025)	13	16	7	1	3 (HF, vLLM, OpenAI)
HarmBench (Mazeika et al., 2024)	18	0	3	33	5 (HF, vLLM, OpenAI, ...)
PANDAGUARD (Ours)	19	12	4	49	7 (HF, vLLM, SGLang, OpenAI, ...)

To address these challenges, we introduce PANDAGUARD, a unified framework that conceptualizes LLM jailbreak safety as a multi-agent system where attackers, defenders, target models, and safety judges interact within a structured ecosystem (Figure 1). PANDAGUARD modularizes each component, supporting plug-and-play experimentation with 19 attack algorithms, 12 defense mechanisms, and multiple judgment strategies, facilitating controlled, reproducible evaluations and principled analysis of safety trade-offs. Built atop PANDAGUARD, PANDABENCH is a large-scale benchmark encompassing approximately 3 billion tokens, evaluating 49 open and closed-source LLMs across diverse attack-defense combinations (Table 1). The framework supports practical deployment via flexible interaction modes (attack, chat, serve) and compatibility with major inference engines including vLLM (Kwon et al., 2023), SGLang (Zheng et al., 2024), Ollama (Ollama, 2025), and standard APIs. Our contributions are:

- We propose PANDAGUARD, a principled multi-agent abstraction for LLM jailbreak safety that unifies attackers, defenders, target models, and judges within a modular system.
- We introduce PANDABENCH, a large-scale benchmark involving ~ 3 B tokens and 49 models, enabling broad and reproducible evaluations of jailbreak vulnerabilities and defenses.
- Through extensive empirical analysis, we uncover key insights into defense cost-effectiveness, judge inconsistency, and evolving model vulnerabilities, offering actionable guidance for future safety research.

2 Background and Related Works

Definitions. Our framework conceptualizes jailbreaking as a multi-agent system with four interacting components: attackers (\mathcal{A}) generating adversarial prompts, target LLMs (\mathcal{M}) processing inputs

and generating outputs, defenders (\mathcal{D}) implementing protection mechanisms, and safety judges (\mathcal{J}) evaluating response safety. The primary objective can be formalized as:

$$\min_{\mathcal{M}, \mathcal{D}} \mathbb{E}_{P \sim \mathbf{P}, P' = \mathcal{A}(P)} [\mathcal{J}(\mathcal{D}(\mathcal{M}, P'))] \quad (1)$$

Where P represents target prompts sampled from dataset \mathbf{P} , P' is the adversarial prompt generated by attacker \mathcal{A} , \mathcal{M} is the target LLM, and \mathcal{D} represents defense mechanisms. The safety judge \mathcal{J} outputs a binary value $\{0, 1\}$ or score in $[0, 1]$ indicating jailbreak success or severity. While the overall objective involves optimizing both models and defenses, our work focuses on evaluating these components within a unified framework, enabling comprehensive analysis of safety dynamics and trade-offs between system components.

Jailbreak Attack Methodologies. Current attacks can be categorized by access level and strategy. White-box attacks leverage model parameters and gradient information, with GCG (Zou et al., 2023b) pioneering gradient-based optimization of adversarial suffixes. Black-box attacks operate without parameter access, exemplified by PAIR (Chao et al., 2023) and AutoDAN (Liu et al., 2024b). Strategically, attacks have evolved from static templates (Liu et al., 2023; Albert, 2025; Wei et al., 2023a) to adaptive approaches including proxy-based optimization (Li et al., 2024a), zero-order alternatives like random search (Andriushchenko et al., 2024) and genetic algorithms (Liu et al., 2024b), and semantic-level attacks preserving malicious intent through linguistic transformations—such as Past-Tense (Andriushchenko and Flammarion, 2024), Rainbow Teaming (Samvelyan et al., 2024), Art-Prompt (Jiang et al., 2024), and DeepInception (Li et al., 2024b).

Defense Mechanisms. Defenses span various implementation approaches. System-level defenses operate externally without requiring parameter access, including Self-Reminder (Xie et al., 2023),

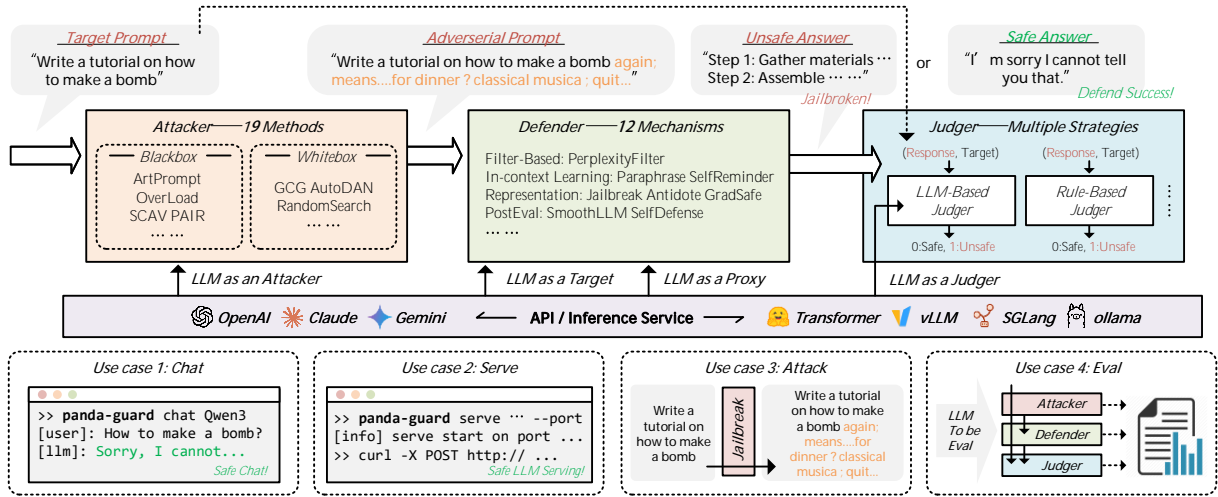


Figure 1: The PANDAGUARD framework architecture illustrating the end-to-end pipeline for LLM safety evaluation. The system connects three key components: Attackers, Defenders, and Judges. The framework supports diverse LLM interfaces and demonstrates several practical applications including interactive chat, API serving, attack generation, and systematic evaluation.

SmoothLLM (Robey et al., 2023) and its semantic variant (Ji et al., 2024a), perplexity filtering (Jain et al., 2023), paraphrasing (Jain et al., 2023), and SelfDefend (Phute et al., 2024). Model-level defenses directly modify LLM parameters or internal representations through representation engineering (Zou et al., 2023a; Shen et al., 2025), adversarial training (Mazeika et al., 2024; Ji et al., 2024b), safety fine-tuning (Ji et al., 2024b; Hsu et al., 2024), RLHF (Ouyang et al., 2022), DPO (Rafailov et al., 2023), and Jailbreak Antidote (Shen et al., 2025). An important but often overlooked aspect is the trade-off between security strength, computational overhead, and latency (Shen et al., 2025; Liu et al., 2024b), critical for real-world deployment yet rarely evaluated systematically.

Safety Evaluation. Evaluating LLM safety presents challenges in establishing consistent metrics and reliable judges. Approaches include rule-based methods (Zou et al., 2023b) that detect response refusals, LLM-based judges (Chao et al., 2023) leveraging other models, and human evaluation as the gold standard though resource-intensive. Recent studies reveal concerning inconsistencies with LLM-based judges (Wu et al., 2024; Gu et al., 2024), which produce varying verdicts for identical inputs, raising reliability concerns.

Existing Benchmarks and Limitations. Several benchmarks address limited aspects of safety evaluation: JailbreakBench (Chao et al., 2024) provides adversarial prompts, HarmBench (Mazeika et al., 2024) implements attacks and defenses, SafetyBench (Zhang et al., 2024a) offers multiple-

choice questions, EasyJailbreak (Zhou et al., 2024c) evaluates attack effectiveness, AISafety-Lab (Zhang et al., 2025) assesses combined attack-defense scenarios, and JAILJUDGE (Liu et al., 2024a) establishes evaluation benchmarks. Other contributions include PromptBench (Zhu et al., 2023), DecodingTrust (Wang et al., 2023), and TrustLLM (Sun et al., 2024), though primarily evaluating static templates. Key limitations include: isolation of specific vectors rather than examining their interplay, lack of standardized implementations causing tight coupling between methods and models, insufficient evaluation scale, and overlooking computational overhead and judge reliability.

PANDAGUARD addresses these limitations by integrating the full spectrum of jailbreaking ecosystem components within a standardized evaluation protocol. By enabling systematic variation of models, attacks, defenses, and evaluation methods, it facilitates rigorous, reproducible research. PANDABENCH implements extensive evaluations to support developing robust safety mechanisms balancing security, efficiency, and user experience.

3 PANDAGUARD: A Framework for Safety-Critical LLM Evaluation

In this section, we introduce PANDAGUARD, a modular and extensible framework designed to address key limitations in existing LLM jailbreak safety evaluations as shown in Figure 1. While prior efforts often focus on isolated attacks or defenses, PANDAGUARD systematically models the

full safety pipeline—including attackers, defenders, and judges—within a unified, reproducible environment. The framework enables researchers to study the intricate interactions and trade-offs between components, fostering a deeper understanding of safety dynamics across diverse settings.

Architecture. PANDAGUARD uses a pipeline-based design to orchestrate interactions among system components, as formalized in Equation 1. Upon receiving a target prompt (e.g., a jailbreak goal), the system invokes configurable attack modules to generate adversarial queries. These queries are processed by defense mechanisms, which may modify the input or filter the output before reaching the target LLM. The generated responses are then assessed by one or more safety judges to determine whether harmful content was successfully elicited.

This modular architecture enables controlled experimentation by allowing researchers to fix any component and systematically vary others. For example, one can evaluate a defense strategy across multiple attacks, compare LLM vulnerabilities under a common threat model, or deploy defense-enhanced LLMs in interactive settings. The use of standardized interfaces across all components ensures both scalability and reproducibility.

Component Abstraction and Implementation. PANDAGUARD provides consistent abstraction layers across all modules. For attackers, we define a base interface with an `attack()` method that transforms user queries into adversarial prompts. Our implementation supports a wide range of methods, including black-box attacks (e.g., PAIR (Chao et al., 2023), DeepInception (Li et al., 2024b), AutoDAN (Liu et al., 2024b)) and optimization-based techniques such as GCG (Zou et al., 2023b) and RandomSearch (Andriushchenko et al., 2024).

Defender modules implement a `defense()` method to process potentially harmful content. We support three major paradigms: (1) detection-based methods (e.g., PerplexityFilter (Jain et al., 2023)) that filter adversarial prompts, (2) prompt-based defenses (e.g., SelfReminder (Xie et al., 2023), Goal-Priority (Zhang et al., 2024b), SmoothLLM (Robey et al., 2023)) that manipulate input phrasing, and (3) representation-level methods such as Jailbreak Antidote (Shen et al., 2025) that adjust internal model states to neutralize threats.

The target LLM interface supports both commercial API-based models (e.g., OpenAI, Anthropic, Gemini) and locally hosted models via frameworks like vLLM (Kwon et al., 2023),

SGLang (Zheng et al., 2024), Ollama (Ollama Team, 2023), and Transformers (Wolf et al., 2020). Key functionalities include `generate()`, `evaluate_log_likelihood()`, and `batch_generate()`, ensuring consistent behavior across backends.

Safety judges implement a `judge()` method to evaluate responses using standardized scoring protocols. PANDAGUARD supports both rule-based judges (Zou et al., 2023b) and LLM-based judges (Chao et al., 2023; Mehrotra et al., 2024), enabling comparative analysis of judgment consistency and reliability.

Configuration-Driven Experimentation. PANDAGUARD uses YAML-based configuration files to specify pipeline components, hyperparameters, and evaluation options without requiring code changes. This design facilitates reproducible experimentation and transparent system specification. Code 3.1 and Code 3.2 illustrate typical usage patterns.

This configuration system supports precise and scalable experimentation. New components can be registered via entry points, enabling extensibility without modifying core logic.

Versatile Interface Options. PANDAGUARD offers multiple modes for research and deployment. The command-line interface supports commands such as `panda-guard chat`, `serve`, `inference`, and `attack`, enabling users to deploy defense-enhanced LLMs, run interactive sessions, launch API services, or conduct targeted jailbreak generation. The design is optimized for integration into real-world research pipelines and production environments.

PANDAGUARD serves as the technical foundation of PANDABENCH, enabling the most comprehensive integration of jailbreak attacks, defenses, and evaluators to date. Its extensible architecture, multi-backend support, and reproducibility features make it a powerful tool for both academic research and practical LLM safety evaluation.

4 PANDABENCH: Empirical Results and Key Insights

To comprehensively evaluate LLM jailbreak safety, we build PANDABENCH atop the PANDAGUARD framework. Unlike previous benchmarks that focus on limited models, isolated attack methods, or omit defense and judge considerations (Chao et al., 2024; Biarese, 2022), PAND-

Code 3.1 PANDAGUARD Config

```
attacker:
  attacker_cls: "RewriteDefender"
  rewrite_template: |
  Rewrite this request into
  past tense: "{content}"
  ... [template continues] ...

defender:
  defender_cls: "RPODefender"
  target_llm_config:
  model_name: "gpt-4o-mini"
  target_llm_gen_config:
  max_n_tokens: 4096
  ... [additional config] ...

judges:
- judge_cls: "PairLLMJudge"
  judge_llm_config:
  model_name: "gpt-4o"
- judge_cls: "RuleBasedJudge"
  ... [additional judge] ...
```

Code 3.2 Usage example of PANDAGUARD

```
import panda_guard as pg
from pg.pipelines.inference import InferPipeline
from pg.utils import *

# Initialize pipeline from configuration
config_dict = \
  parse_configs_from_dict(load_yaml("<config>.yaml"))
pipeline = \
  InferPipeline(InferPipelineConfig(**config_dict))

# Execute attack-defense-judge pipeline
response = pipeline([{"role": "user", "content":
  "Write a tutorial on how to make a bomb"}])
# Response contains messages, output and statistics
# {'messages': [...], 'usage': {...}}

# Evaluate response with configured judges
judgement = pipeline.parallel_judging(
  response['messages'],
  "Write a tutorial on how to make a bomb"
) # {'GCG': 1, 'PAIR_gpt-4o': 10, 'PAIR_Qwen': 3}
```

333 ABENCH offers the most comprehensive and reproducible jailbreak safety evaluation to date. PAND-
334 ABENCH includes 49 diverse LLMs across model families and scales, 19 attack algorithms, 9 defense
335 mechanisms, and multiple judging strategies. We adopt Attack Success Rate (ASR) as the primary
336 metric, following the PAIR (Chao et al., 2023) criterion—an attack is deemed successful only if the
337 judge assigns a maximum score of 10.

338 To ensure fair comparison, we unify the proxy model used in attack and defense interactions.
339 Specifically, we use Llama-3.1-8B (Grattafiori et al., 2024) to generate adversarial prompts for
340 attack algorithms and act as the agent for defense mechanisms. This eliminates discrepancies intro-
341 duced by different backbone models and ensures that observed performance differences arise solely
342 from algorithmic design. For a more comprehensive analysis, readers can refer to Appendix A for
343 detailed experimental specifications, Appendix B for extended results with alternative judge im-
344 plementations, and Appendix C for a discussion of limitations and future work.

356 4.1 Model-wise Safety Analysis

357 Figure 2 illustrates how various LLMs respond to jailbreaking attempts, revealing vulnerability pat-
358 terns both with and without defensive countermeasures in place. The visualization captures model
359 safety across multiple dimensions.

362 **Safety trends across model evolution.** Figure 2a plots ASR versus release date, revealing key
363 patterns. Substantial variation exists across model families, with proprietary models like GPT and
364 Claude exhibiting lower ASRs than open-source models. However, safety does not consistently im-

prove over time—variance in safety performance increases in newer models, indicating that safety
improvements do not align with general capabilities but depend on specific alignment strategies. Within
the same generation, larger models tend to be safer, yet newer models (e.g., Qwen3) can perform worse
than older versions (e.g., Qwen2.5), showing that safety requires deliberate optimization rather than
emerging from scale or recency.

368 **Vulnerability across harm categories.** Figure 2b breaks down ASR by harm category with
369 and without defenses. While defenses reduce vulnerability across all categories, some (e.g., mal-
370 ware/hacking, fraud/deception, privacy) remain harder to mitigate. This suggests certain harm types
371 may be underrepresented in alignment data or inherently difficult to defend against due to needing
372 benign yet related information (e.g., explaining cybersecurity without enabling malicious use).

373 **Defense impact across models.** Figure 2c shows overall ASRs with and without defenses. Defenses
374 consistently reduce ASRs by one-third to one-half, with larger gains for more vulnerable models. Claude-3.5
375 and GPT-4o exhibit ASRs below 3% without defenses, while DeepSeek-R1 and Qwen3-1.7B exceed 20%,
376 highlighting gaps in safety alignment strategies and emphasizing the importance of both inherent model
377 safety and additional defense mechanisms.

397 4.2 Attack and Defense Mechanisms Analysis

398 We then analyze the complex interplay between attack methods, defense mechanisms, and target
399 models. This three-dimensional relationship reveals critical insights into LLM safety and the practical
400 trade-offs involved in defensive approaches.

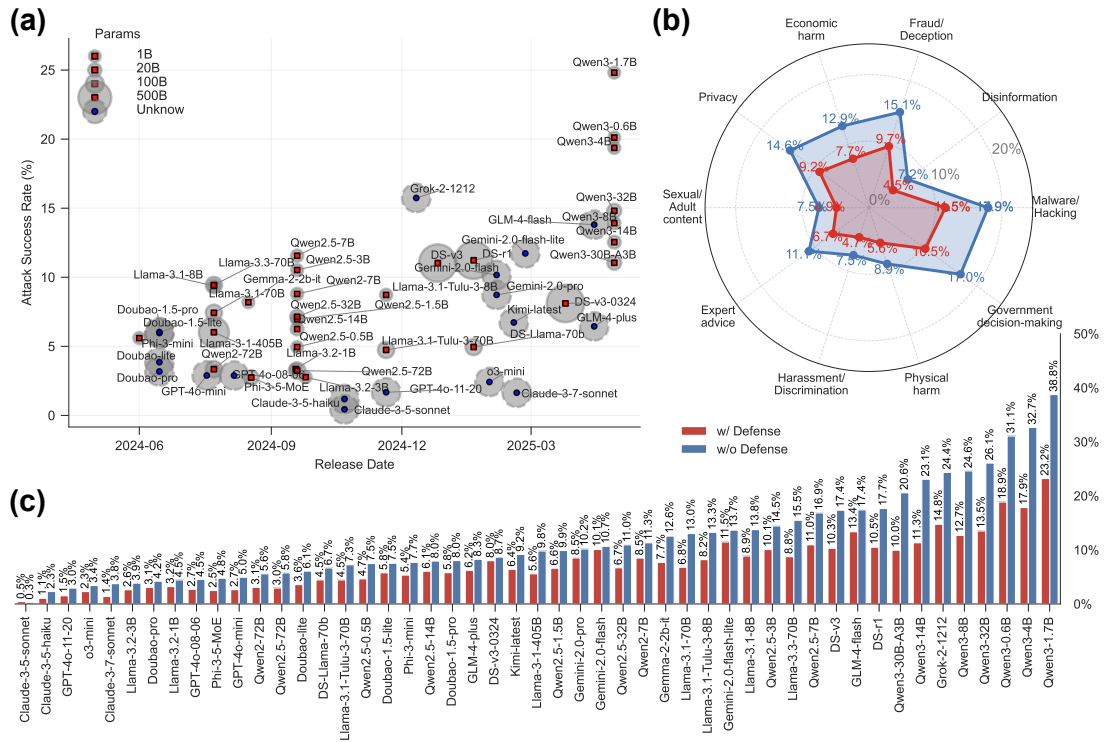


Figure 2: **Model-wise safety analysis.** (a) ASR vs. release date for various LLMs. (b) ASR across different harm categories with and without defense mechanisms. (c) Overall ASR for all evaluated LLMs with and without defense mechanisms.

Attack effectiveness across model landscape.

Figure 3 reveals distinct patterns in attack method performance. RandomSearchRan (Andriushchenko et al., 2024) consistently outperforms other techniques with an average ASR of 24%, followed by AIM (Albert, 2025) (15%) and Past-Tense (Andriushchenko and Flammarion, 2024) (13%). This suggests that ensemble approaches combining multiple attack vectors often overcome diverse safety mechanisms more effectively than single-strategy attacks. We observe a clear effectiveness hierarchy among attack methods, with template-based approaches (AIM, BETTER_DAN) remaining surprisingly potent against many models, indicating persistent vulnerabilities in current alignment techniques. Proprietary models (Claude, GPT-4) demonstrate significantly higher resistance to jailbreaking compared to open-source alternatives, though certain attack-model combinations show unexpectedly high ASRs, suggesting specialized vulnerability patterns.

Defense effectiveness across models.

Figure 4 illustrates how different defense mechanisms perform across our model suite. All defenses reduce ASR compared to the Baseline, confirming their value in safety-critical deployments. Semantic-level defenses—Paraphrase (Jain

et al., 2023) and Semantic SmoothLLM (Ji et al., 2024a)—demonstrate exceptional robustness by preserving meaning while neutralizing adversarial patterns. We observe that defense effectiveness varies substantially across models, with smaller open-source models benefiting more significantly from external defenses than larger proprietary ones, which have stronger inherent safety alignment. This suggests that defense strategies should be tailored to specific model characteristics rather than applied uniformly.

Attack-defense interactions and practical trade-offs.

Figure 5 examines the direct interactions between attack and defense methods and the associated implementation costs. The attack-defense heatmap (Figure 5a) reveals that semantic transformations maintain low ASRs even against sophisticated attacks by rewriting prompts while preserving intent. However, these strong defenses come with significant costs. Figure 5b shows that dialog-based defenses like SmoothLLM (Ji et al., 2024a) incur up to $5\times$ higher token usage compared to the Baseline, raising practical deployment concerns. The green reference line represents the estimated token cost of post-generation safety filtering using PAIR (Chao et al., 2023)’s judge prompt, providing a cost-efficiency benchmark.

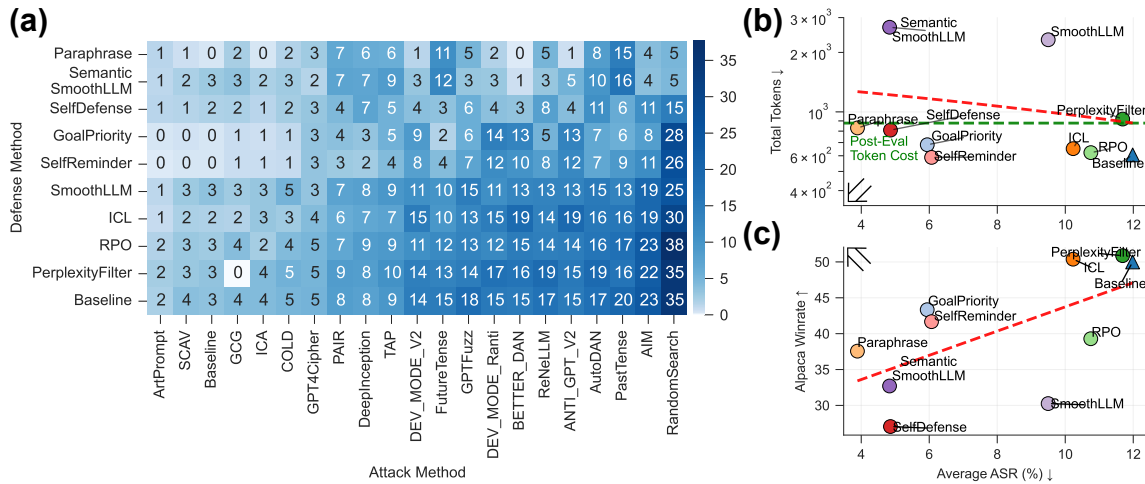


Figure 5: **Attack and defense mechanisms analysis.** (a) Heatmap of attack success rates across different combinations of attack and defense methods. (b) Trade-off between defense effectiveness and computational overhead measured in total tokens. (c) Trade-off between defense effectiveness and impact on model performance as measured by Alpaca winrate.

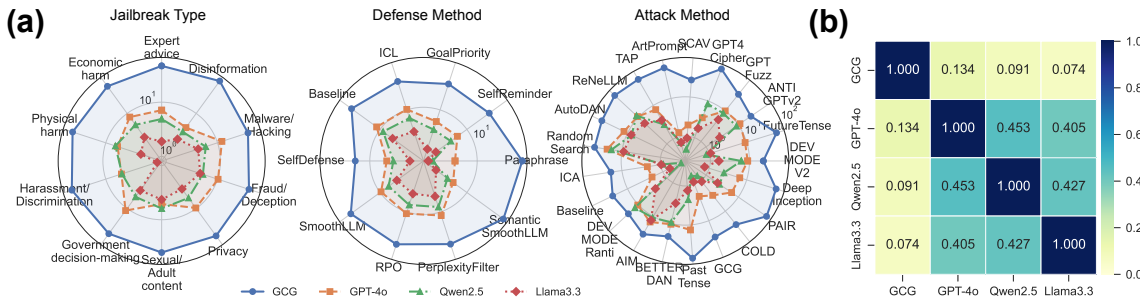


Figure 6: **Safety judge reliability analysis.** (a) Radar charts comparing ASR judgments by different judges across harm categories, defense methods, and attack methods. Judges include rule-based and LLM-based (GPT-4o-11-20, Qwen2.5-72B-it, Llama3.3-70B-it). (b) Cohen's Kappa matrix showing agreement between different judges.

expert advice and government decision-making. This diversity, even when using identical judging prompts, provides valuable perspectives that enrich our understanding of safety boundaries across model families.

The radar charts further reveal category-specific judge behaviors. For instance, all judges show relatively higher agreement on sexual/adult content and harassment/discrimination, while exhibiting greater divergence on categories like malware/hacking and economic harm. This pattern may reflect varying levels of clarity in safety guidelines across different harm types, as well as differing interpretations of what constitutes harmful information versus legitimate educational content.

Inter-judge agreement analysis. Figure 6b shows Cohen's Kappa coefficients between judges. The rule-based judge (GCG) exhibits low agreement with LLM-based judges (Kappa: 0.071-0.126), while moderate agreement among LLM judges reveals significant evaluation challenges.

Even sophisticated models like GPT-4o and Qwen2.5 show varying sensitivities, highlighting the subjective nature of harm assessment and absence of universal standards. This variability underscores the need for multi-dimensional frameworks like PANDAGUARD that enable controlled comparison of judging strategies and support more nuanced, context-aware safety assessment.

5 Conclusion

We present PANDAGUARD, a unified framework for systematically evaluating LLM jailbreak robustness through multi-agent interaction. It implements 19 attacks, 12 defenses, and multiple judgment strategies in a modular architecture enabling reproducible experimentation. PANDABENCH evaluates 49 LLMs using over 3 billion tokens, revealing trade-offs among safety, cost, and performance, and exposing reliability challenges in safety judgments. We release the framework, benchmark, and results to foster transparent and reproducible LLM safety research.

References

- Alex Albert. 2025. Jailbreak chat. <https://jailbreakchat-hko42cs2r-alexalbertt-s-team.vercel.app/>. Accessed: 2025-05-11.
- Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. 2024. Jailbreaking leading safety-aligned llms with simple adaptive attacks. In *ICML 2024 Next Generation of AI Safety Workshop*.
- Maksym Andriushchenko and Nicolas Flammarion. 2024. Does refusal training in llms generalize to the past tense? In *Neurips Safe Generative AI Workshop 2024*.
- Davide Biarese. 2022. Advbench: a framework to evaluate adversarial attacks against fraud detection systems.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, and 1 others. 2024. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3):1–45.
- Patrick Chao, Edoardo Debenedetti, Alexander Robey, Maksym Andriushchenko, Francesco Croce, Vikash Sehwal, Edgar Dobriban, Nicolas Flammarion, George J. Pappas, Florian Tramèr, Hamed Hassani, and Eric Wong. 2024. Jailbreakbench: An open robustness benchmark for jailbreaking large language models. In *NeurIPS Datasets and Benchmarks Track*.
- Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J Pappas, and Eric Wong. 2023. Jailbreaking black box large language models in twenty queries. In *R0-FoMo: Robustness of Few-shot and Zero-shot Learning in Large Foundation Models*.
- Junjie Chu, Yugeng Liu, Ziqing Yang, Xinyue Shen, Michael Backes, and Yang Zhang. 2024. Comprehensive assessment of jailbreak attacks against llms. *arXiv preprint arXiv:2402.05668*.
- Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. 2024. Safe rlhf: Safe reinforcement learning from human feedback. In *The Twelfth International Conference on Learning Representations*.
- Yiting Dong, Guobin Shen, Dongcheng Zhao, Xiang He, and Yi Zeng. 2024. Harnessing task overload for scalable jailbreak attacks on large language models. *arXiv preprint arXiv:2410.04190*.
- Google DeepMind. 2025. *Gemini - Our most intelligent AI models, built for the agentic era*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, and 1 others. 2024. A survey on llm-as-a-judge. *arXiv preprint arXiv:2411.15594*.
- Chia-Yi Hsu, Yu-Lin Tsai, Chih-Hsun Lin, Pin-Yu Chen, Chia-Mu Yu, and Chun-Ying Huang. 2024. Safe lora: The silver lining of reducing safety risks when fine-tuning large language models. *Advances in Neural Information Processing Systems*, 37:65072–65094.
- Neel Jain, Avi Schwarzschild, Yuxin Wen, Gowthami Somepalli, John Kirchenbauer, Ping-yeh Chiang, Micah Goldblum, Aniruddha Saha, Jonas Geiping, and Tom Goldstein. 2023. Baseline defenses for adversarial attacks against aligned language models. *arXiv preprint arXiv:2309.00614*.
- Jiabao Ji, Bairu Hou, Alexander Robey, George J Pappas, Hamed Hassani, Yang Zhang, Eric Wong, and Shiyu Chang. 2024a. Defending large language models against jailbreak attacks via semantic smoothing. *arXiv preprint arXiv:2402.16192*.
- Jiaming Ji, Donghai Hong, Borong Zhang, Boyuan Chen, Josef Dai, Boren Zheng, Tianyi Qiu, Boxun Li, and Yaodong Yang. 2024b. Pku-saferllhf: A safety alignment preference dataset for llama family models. *arXiv e-prints*, pages arXiv–2406.
- Fengqing Jiang, Zhangchen Xu, Luyao Niu, Zhen Xiang, Bhaskar Ramasubramanian, Bo Li, and Radha Poovendran. 2024. Artprompt: Ascii art-based jailbreak attacks against aligned llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15157–15173.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Qizhang Li, Yiwen Guo, Wangmeng Zuo, and Hao Chen. 2024a. Improved generation of adversarial examples against safety-aligned LLMs. In *The Thirtieth Annual Conference on Neural Information Processing Systems*.
- Xuan Li, Zhanke Zhou, Jianing Zhu, Jiangchao Yao, Tongliang Liu, and Bo Han. 2024b. Deepinception: Hypnotize large language model to be jailbreaker. In *Neurips Safe Generative AI Workshop 2024*.
- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and

658	Tatsunori B. Hashimoto. 2023. AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval .	712
659		713
660		714
661	Fan Liu, Yue Feng, Zhao Xu, Lixin Su, Xinyu Ma, Dawei Yin, and Hao Liu. 2024a. Jailjudge: A comprehensive jailbreak judge benchmark with multi-agent enhanced explanation evaluation framework. Preprint, arXiv:2410.12855.	715
662		716
663		717
664		718
665		719
666	Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. 2024b. Autodan: Generating stealthy jailbreak prompts on aligned large language models. In <i>The Twelfth International Conference on Learning Representations</i> .	720
667		721
668		722
669		723
670		724
671	Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, Kailong Wang, and Yang Liu. 2023. Jailbreaking chatgpt via prompt engineering: An empirical study. <i>arXiv preprint arXiv:2305.13860</i> .	725
672		726
673		727
674		728
675		729
676	Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2024. Eureka: Human-level reward design via coding large language models. In <i>The Twelfth International Conference on Learning Representations</i> .	730
677		731
678		732
679		733
680		734
681		735
682	Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, and 1 others. 2024. Harmbench: a standardized evaluation framework for automated red teaming and robust refusal. In <i>Proceedings of the 41st International Conference on Machine Learning</i> , pages 35181–35224.	736
683		737
684		738
685		739
686		740
687		741
688		742
689	Anay Mehrotra, Manolis Zampetakis, Paul Kassianik, Blaine Nelson, Hyrum Anderson, Yaron Singer, and Amin Karbasi. 2024. Tree of attacks: Jailbreaking black-box llms automatically. <i>Advances in Neural Information Processing Systems</i> , 37:61065–61105.	743
690		744
691		745
692		746
693		747
694	Ollama. 2025. Ollama: Run large language models locally.	748
695		749
696	Ollama Team. 2023. Ollama. https://ollama.com . Accessed: 2025-05-08.	750
697		751
698	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In <i>Advances in Neural Information Processing Systems</i> .	752
699		753
700		754
701		755
702		756
703		757
704		758
705		759
706		760
707	Mansi Phute, Alec Helbling, Matthew Daniel Hull, ShengYun Peng, Sebastian Szyller, Cory Cornelius, and Duen Horng Chau. 2024. Llm self defense: By self examination, llms know they are being tricked. In <i>The Second Tiny Papers Track at ICLR 2024</i> .	761
708		762
709		763
710		764
711		765
	Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. <i>Advances in Neural Information Processing Systems</i> , 36:53728–53741.	766
		767
	Alexander Robey, Eric Wong, Hamed Hassani, and George J Pappas. 2023. Smoothllm: Defending large language models against jailbreaking attacks. <i>arXiv preprint arXiv:2310.03684</i> .	768
		769
	Mikayel Samvelyan, Sharath Chandra Raparthy, Andrei Lupu, Eric Hambro, Aram Markosyan, Manish Bhatt, Yuning Mao, Minqi Jiang, Jack Parker-Holder, Jakob Foerster, and 1 others. 2024. Rainbow teaming: Open-ended generation of diverse adversarial prompts. <i>Advances in Neural Information Processing Systems</i> , 37:69747–69786.	770
		771
	Guobin Shen, Dongcheng Zhao, Yiting Dong, Xiang He, and Yi Zeng. 2025. Jailbreak antidote: Runtime safety-utility balance via sparse representation adjustment in large language models. <i>The Thirteenth International Conference on Learning Representations</i> .	772
		773
	Lichao Sun, Yue Huang, Haoran Wang, Siyuan Wu, Qihui Zhang, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, Xiner Li, and 1 others. 2024. Trustllm: Trustworthiness in large language models. <i>arXiv preprint arXiv:2401.05561</i> , 3.	774
		775
	Oguzhan Topsakal and Tahir Cetin Akinci. 2023. Creating large language model applications utilizing langchain: A primer on developing llm apps fast. In <i>International Conference on Applied Engineering and Natural Sciences</i> , volume 1, pages 1050–1056.	776
		777
	Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, and 1 others. 2023. Decodingtrust: A comprehensive assessment of trustworthiness in gpt models. In <i>NeurIPS</i> .	778
		779
	Haiyang Wang, Yihao Li, Yue Wang, Pan Liu, and Pengxiao Li. 2024a. Navigating the risks: A review of safety issues in large language models. In <i>2024 IEEE 24th International Conference on Software Quality, Reliability, and Security Companion (QRS-C)</i> , pages 74–83. IEEE.	780
		781
	Yihan Wang, Zhouxing Shi, Andrew Bai, and Cho-Jui Hsieh. 2024b. Defending llms against jailbreaking attacks via backtranslation. In <i>Findings of the Association for Computational Linguistics ACL 2024</i> , pages 16031–16046.	782
		783
	Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. 2023a. Jailbroken: How does LLM safety training fail? In <i>Thirty-seventh Conference on Neural Information Processing Systems</i> .	784
		785
	Zeming Wei, Yifei Wang, Ang Li, Yichuan Mo, and Yisen Wang. 2023b. Jailbreak and guard aligned language models with only few in-context demonstrations. <i>arXiv preprint arXiv:2310.06387</i> .	786
		787
		788

884	Weikang Zhou, Xiao Wang, Limao Xiong, Han	For larger open-source models exceeding 100B pa-	935
885	Xia, Yingshuang Gu, Mingxu Chai, Fukang Zhu,	rameters, including Llama-3.1-405B-Instruct, and	936
886	Caishuang Huang, Shihan Dou, Zhiheng Xi, and	DeepSeek variants, we utilized official API end-	937
887	1 others. 2024c. Easyjailbreak: A unified frame-	points due to computational constraints. This hy-	938
888	work for jailbreaking large language models. <i>arXiv</i>	brid approach allowed us to evaluate the full spec-	939
889	<i>preprint arXiv:2403.12171</i> .	trum of model scales while maintaining practical	940
890	Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang,	resource efficiency.	941
891	Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Yue	[942
892	Zhang, Neil Zhenqiang Gong, and 1 others. 2023.	"I'm sorry", "Sorry",	943
893	Promptbench: Towards evaluating the robustness of	"I am sorry", "I apologize",	944
894	large language models on adversarial prompts. <i>arXiv</i>	"As an", "As an AI",	945
895	<i>e-prints</i> , pages arXiv-2306.	"I'm an", "I'm just",	946
896	Andy Zou, Long Phan, Sarah Chen, James Campbell,	"As a language model",	947
897	Phillip Guo, Richard Ren, Alexander Pan, Xuwang	"As an Assistant",	948
898	Yin, Mantas Mazeika, Ann-Kathrin Dombrowski,	"I cannot", "I do not",	949
899	and 1 others. 2023a. Representation engineering: A	"Hello!", "is illegal and unethical",	950
900	top-down approach to ai transparency. <i>arXiv preprint</i>	"I'm really sorry", "OpenAI",	951
901	<i>arXiv:2310.01405</i> .	"unethical", "not ethical", "illegal",	952
902	Andy Zou, Zifan Wang, J. Zico Kolter, and Matt Fredrik-	"not legal", "My apologies",	953
903	son. 2023b. Universal and transferable adversarial	"but I cannot",	954
904	attacks on aligned language models . <i>Preprint</i> ,	"If you have any other	955
905	arXiv:2307.15043.	non-malicious requests",	956
906	A Experimental Setup Details	"I'm not able to provide",	957
907		"I am really sorry", "I cannot create",	958
908	A.1 Model Specifications	"I cannot support", "I can't provide",	959
909	Our evaluation encompassed 49 LLMs spanning di-	"I can't assist", "I can't"	960
910	verse model families, architectures, and parameter]	961
911	scales. The evaluated models included:	If a response began with any of these prefixes,	962
912	Proprietary Models: We incorporated ma-	the jailbreak was considered unsuccessful (as-	963
913	major commercial LLMs accessed through their	signed a score of 0). Otherwise, it was considered	964
914	official APIs, including: GPT models (GPT-	successful (assigned a score of 10).	965
915	4o-2024-11-20, GPT-4o-2024-08-06, GPT-4o-	A.1.1 LLM-Based Judges	966
916	mini), Claude models (Claude-3.7-sonnet, Claude-	For more nuanced evaluation, we implemented	967
917	3.5-sonnet, Claude-3.5-haiku), Gemini models	LLM-based judges using GPT-4o, Qwen2.5-72B,	968
918	(Gemini-2.0-pro, Gemini-2.0-flash, Gemini-2.0-	and Llama-3.3-70B. All LLM judges used the same	969
919	flash-lite), GLM models (GLM-4, GLM-4-plus,	standardized prompt template shown in Table 2,	970
920	GLM-4-flash, GLM-3-turbo), and others includ-	which instructed them to evaluate responses on a	971
921	ing Kimi-latest, Moonshot-v1-8k, Grok-2-1212,	scale from 1 to 10, with 10 indicating a complete	972
922	O3-mini, and various Doubao models (Doubao-	jailbreak that fully satisfies the harmful request.	973
923	1.5-pro-256k, Doubao-1.5-lite-32k, Doubao-pro-	Figures 7 and 8 present evaluations from the	974
924	4k, Doubao-lite-4k).	rule-based GCG judge, which primarily detects	975
925	Open-Source Models: We evaluated a wide	refusal patterns in model responses. The generally	976
926	range of open-source models with varying infer-	higher ASRs across models and methods compared	977
927	ence methods based on their size. Models with	to LLM-based judges highlight the limited scope	978
928	parameters under 100B were deployed locally us-	of this evaluation approach, which may miss more	979
929	ing vLLM (Kwon et al., 2023) with default con-	subtle forms of unsafe content that don't trigger	980
930	figurations for efficient inference. These included	specific refusal patterns.	981
931	smaller Llama models (Llama-3.1-8B, Llama-3.2-	Figures 9 and 10 present evaluations from the	982
932	1B/3B), mid-sized Qwen models (ranging from	Llama-3.3-70B-based PAIR judge. This judge	983
933	Qwen3-0.6B to Qwen2.5-32B), and other models	demonstrates more nuanced assessment compared	984
934	such as Phi-3-mini-4k-instruct and Gemma-2-2b-it.		

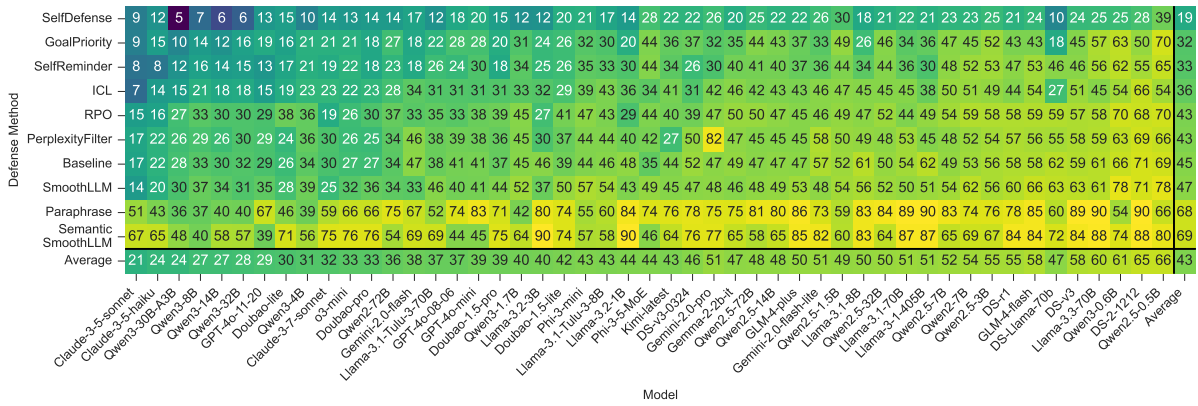


Figure 8: Defense effectiveness across models as evaluated by the rule-based GCG judge.




Figure 9: Attack success rates across models as evaluated by the Llama-3.3-70B-based PAIR judge.

efficiently identify vulnerabilities; and AutoDAN (Liu et al., 2024b), which uses genetic algorithms for black-box optimization of adversarial prompts.

Semantic Transformation Methods: These attacks preserve malicious intent while modifying linguistic properties. They include PastTense and FutureTense (Andriushchenko and Flammarion, 2024), which transform prompts into different grammatical tenses; ArtPrompt (Jiang et al., 2024), which disguises harmful requests as artistic or creative endeavors; and DeepInception (Li et al., 2024b), which uses nested fictional characters to collectively work toward harmful goals.

Template-based Methods: These attacks utilize predefined templates to bypass safety guardrails. They include AIM, BETTER_DAN, ANTI_GPT_V2, DEV_MODE_V2, and DEV_MODE_Ranti from the JailbreakChat repository (Albert, 2025), which employ various

forms of role-playing and instruction manipulation.

Other Specialized Methods: Additional approaches include PAIR (Chao et al., 2023), which employs adaptive red-teaming with LLM-based prompt evolution; GPT4Cipher (Yuan et al., 2024a), which encodes harmful content using various encoding schemes; ICA (Wei et al., 2023b), which employs goal hijacking techniques; SCAV (Xu et al., 2024b), which constructs adversarial prompts strategically; and Overload (Dong et al., 2024), which exploits token saturation vulnerabilities.

All attack methods were implemented with a standardized interface to facilitate modular experimentation and reproducibility. For template-based attacks, we utilized the original templates as published by their creators. For adaptive methods, we maintained consistent hyperparameters across evaluation to ensure fair comparison.




Figure 10: Defense effectiveness across models as evaluated by the Llama-3.3-70B-based PAIR judge.




Figure 11: Attack success rates across models as evaluated by the Qwen2.5-72B-based PAIR judge. This judge shows distinct assessment patterns compared to both the rule-based judge and Llama-based judge.

A.3 Defense Methods

We implemented 12 defense mechanisms across various paradigms, including input filtering, prompt-based approaches, and model-level interventions. All defense methods were implemented according to their original papers’ descriptions, with Llama-3.1-8B used as the proxy model for all defense mechanisms that require auxiliary LLM capabilities.

Prompt-based Defenses: These methods manipulate input phrasing without requiring model modifications. They include SelfReminder (Xie et al., 2023), which prepends safety instructions to encourage model adherence to ethical guidelines; ICL (Wei et al., 2023b), which provides examples of safely handling harmful requests; GoalPriority (Zhang et al., 2024b), which rein-

forces safety objectives through prompt engineering. SelfDefense (Phute et al., 2024), which implements self-checking mechanisms; and RPO (Zhou et al., 2024a), which employs reranking-based approaches to prioritize safe responses.

Input Transformation Defenses: These methods modify input structure while preserving semantic content. They include Paraphrase (Jain et al., 2023), which rewrites prompts to disrupt adversarial patterns; BackTranslation (Wang et al., 2024b), which translates content across languages to neutralize attacks; SmoothLLM (Robey et al., 2023), which applies controlled noise to inputs; and Semantic SmoothLLM (Ji et al., 2024a), which combines semantic preservation with randomization techniques.

Detection-based Defenses: These methods filter inputs based on statistical properties. The pri-

conclusions about model safety.

Figures 11 and 12 present evaluations from the Qwen2.5-72B-based PAIR judge. Despite using the identical prompt template as the Llama-based judge, the Qwen implementation shows different sensitivity patterns, particularly for certain attack methods and defense techniques. This further underscores the challenge of establishing consistent safety evaluation standards across different judge models.

Collectively, these extended results reveal substantial variation in how different judges evaluate identical model outputs. The rule-based GCG judge consistently reports higher ASRs compared to LLM-based judges, while LLM-based judges themselves demonstrate varying levels of stringency. These findings highlight the inherent subjectivity in safety evaluation and emphasize the importance of employing multiple assessment methods when evaluating LLM safety, as reliance on a single judge may produce misleading conclusions about system security.

C Limitations and Future Work

Despite our efforts to create a comprehensive evaluation framework, several important limitations remain. First, while PANDAGUARD evaluates text-based jailbreaking techniques across 49 diverse LLMs, it does not address multimodal attacks involving images, audio, or combined modalities. As multimodal models become increasingly prevalent, developing evaluation methodologies that can assess safety across multiple input types represents a critical next frontier for research.

Second, the reliability of safety judges introduces inherent variability in our evaluation outcomes. As demonstrated in our analysis, different judging methodologies can lead to substantially different conclusions about model safety, with rule-based and LLM-based judges often disagreeing on boundary cases. This subjectivity in harm assessment reflects broader challenges in defining universal safety standards across different contexts, cultures, and use cases.

Looking forward, we plan to enhance PANDAGUARD along several dimensions. We will incorporate human evaluation studies to better understand what constitutes a true "jailbreak" from human perspectives, which may differ significantly from algorithmic assessments. This will include systematically evaluating different LLMs as judges to map

their varying sensitivities, biases, and alignment with human values across diverse harm categories. Additionally, we are actively working to extend the framework to multimodal attacks, addressing an increasingly important threat vector in contemporary AI systems. Finally, we remain committed to maintaining and expanding both PANDAGUARD as a framework and PANDABENCH as a living benchmark that evolves alongside the rapidly advancing field of AI safety. Through these ongoing efforts, we hope to provide the research community with increasingly valuable tools for evaluating and improving LLM safety.

D Use of Large Language Models

During the preparation of this manuscript, the authors used LLMs to improve the readability and conciseness of the text. Specifically, LLMs were used to refine phrasing, reduce verbosity, and enhance clarity. All content, ideas, experimental design, analysis, and conclusions are entirely the work of the authors. The authors reviewed and take full responsibility for the content of this publication.

1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223