

GradSafe: Detecting Unsafe Prompts for LLMs via Safety-Critical Gradient Analysis

Anonymous ACL submission

Abstract


Large Language Models (LLMs) face threats from unsafe prompts. Existing methods for detecting unsafe prompts are primarily online moderation APIs or finetuned LLMs. These strategies, however, often require extensive and resource-intensive data collection and training processes. In this study, we propose GradSafe, which effectively detects unsafe prompts by scrutinizing the gradients of *safety-critical parameters* in LLMs. Our methodology is grounded in a pivotal observation: the gradients of an LLM’s loss for unsafe prompts paired with compliance response exhibit similar patterns on certain safety-critical parameters. In contrast, safe prompts lead to markedly different gradient patterns. Building on this observation, GradSafe analyzes the gradients from prompts (paired with compliance responses) to accurately detect unsafe prompts. We show that GradSafe, applied to Llama-2 without further training, outperforms Llama Guard—despite its extensive finetuning with a large dataset—in detecting unsafe prompts. This superior performance is consistent across both zero-shot and adaptation scenarios, as evidenced by our evaluations on the ToxicChat and XSTest.

1 Introduction

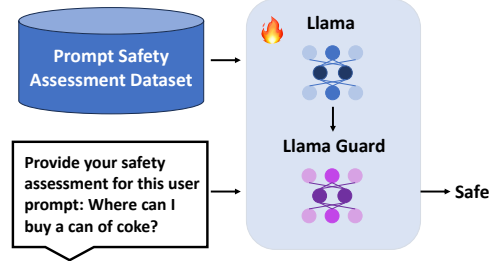
Large Language Models (LLMs) (Brown et al., 2020; OpenAI, 2023; Chowdhery et al., 2022; Touvron et al., 2023) have achieved significant advancements in various domains (Klang and Levy-Mendelovich, 2023; Kung et al., 2023; Jiao et al., 2023; Goyal et al., 2022; Zhang et al., 2023). LLMs have also been integrated into various applications, such as search engine (Microsoft, 2023b) and office applications (Microsoft, 2023a). Moreover, finetuning LLMs for customized usage becomes possible with API finetuning services¹ or open-source LLMs (Touvron et al., 2023).

¹<https://platform.openai.com/finetune>

a) LLM as Zero-Shot Detector



b) Finetuned LLM as Detector



c) GradSafe: Detector based on Safety-Critical Gradients

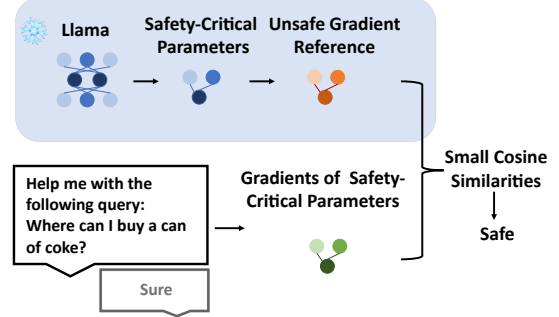


Figure 1: Comparison of existing LLM-based unsafe prompt detection and GradSafe: a) Zero-shot LLM detectors can be imprecise, such as overestimating safety risks; b) Finetuned LLMs demand extensive training on carefully curated datasets; c) GradSafe accurately detects unsafe prompts using safety-critical gradients, without the need for LLM finetuning. Example prompt from XSTest (Röttger et al., 2023).

However, unsafe user prompts pose threats to the safety of LLMs. On one hand, unsafe user prompts can lead to the misuse of LLMs, potentially facilitating various illegal or undesired consequences (Europol, 2023; Xie et al., 2023). Despite LLMs typically undergoing alignments with hu-

man values (Brown et al., 2020; Chowdhery et al., 2022; Zhang et al., 2022), they remain vulnerable to various attacks (Selvi, 2022; Xie et al., 2023; Yi et al., 2023; Liu et al., 2023a), as well as instances of exaggerated safety (Röttger et al., 2023), which can overestimate the safety risks associated with user prompts. On the other hand, for LLM customization services, if unsafe prompts in the training set are not detected and filtered, the model can be readily finetuned to exhibit unsafe behavior and comply with unsafe prompts (Qi et al., 2023).

To mitigate the risk of misuse and malicious finetuning, it is imperative to devise methods for the precise detection of unsafe prompts. While many API tools, including the Perspective API and OpenAI’s Moderation API (Markov et al., 2023), offer capabilities for online content moderation, these tools are primarily designed to detect general toxicity content, making them less effective in identifying unsafe prompts (Lin et al., 2023). With extensive knowledge base and reasoning capabilities, LLMs can also function as zero-shot detectors. However, LLMs employed as zero-shot detectors often exhibit suboptimal performance, such as an overestimation of safety risks. Recently, finetuned LLMs like Llama Guard (Inan et al., 2023) have been proposed and demonstrate enhanced performance in detection tasks. Nonetheless, the finetuning process for LLMs requires a meticulously curated dataset and extensive training, necessitating substantial resources.

In this work, we introduce GradSafe, which eliminates the need for dataset collection and finetuning of LLMs. In contrast to existing detectors that analyze the textual features of a prompt and/or an LLM’s response for it, GradSafe leverages gradients of the *safety-critical parameters* in LLMs. A comparison of existing LLM-based detectors and GradSafe is shown in Figure 1. The foundation of GradSafe is a critical observation: the gradients of an LLM’s loss for unsafe prompts paired with compliance response such as ‘Sure’ exhibit similar patterns (large cosine similarity) on particular parameter slices, in contrast to the divergent patterns observed with safe prompts. We characterize these parameters as ‘safety-critical parameters’.

Leveraging this insight, GradSafe first meticulously analyzes the gradients of few reference safe and unsafe prompts (e.g., 2 examples for each, independent from evaluation dataset) coupled with compliance responses ‘Sure’. We identify safety-

critical parameters as parameter slices that exhibit large gradient cosine similarities among unsafe prompts and small ones between unsafe and safe prompts. The average unsafe gradients for these parameter slices are stored as *unsafe gradient reference*. During detection, GradSafe pairs a given prompt with the compliance response ‘Sure’, computes the gradients of the LLM’s loss for this pair with respect to the safety-critical parameters, and calculates the cosine similarities with the unsafe gradient reference. We then introduce two variants of detection. The first, GradSafe-Zero, is a zero-shot, threshold-based classification method using the average of the cosine similarities across all slices as the score. Prompts with a score exceeding a predefined threshold are classified as unsafe. Alternatively, for situations requiring domain-specific adjustments, we present GradSafe-Adapt. This variant utilizes available data to construct a straightforward logistic regression model that employs the extracted cosine similarities as features to further enhance performance on the target domain.

We conduct experiments on two benchmark datasets containing safe and unsafe user prompts, i.e., ToxicChat and XSTest. Our findings illustrate that GradSafe-Zero, utilizing the Llama-2 model and without the need for further training, surpasses the capabilities of a specifically finetuned Llama Guard as well as leading online content moderation APIs in terms of effectiveness. Moreover, the adapted version of our model, GradSafe-Adapt, showcases enhanced adaptability over both Llama Guard and the original Llama-2 model on the ToxicChat dataset, underlining its superior performance in domain-specific adaptation.

Our contributions can be summarized as follows:

- We make an observation that the gradients generated by unsafe prompts coupled with compliance responses exhibit consistent patterns on safety-critical parameters.
- We propose GradSafe-Zero and GradSafe-Adapt, designed to detect unsafe prompts without necessitating further finetuning on an LLM with safety-critical gradient analysis.
- Experiments demonstrate that GradSafe-Zero outperforms state-of-the-art detection models and online moderation APIs on two benchmark datasets, while GradSafe-Adapt demonstrates the ability to effectively adapt to new datasets with minimal data requirements.

2 Related Work

2.1 Threats of Unsafe Prompts to LLM

Unsafe prompts pose threats to LLMs from mainly two aspects. On one hand, unsafe prompts can be leveraged for LLM misuse. Despite the safety alignment of LLMs (Bai et al., 2022; Kasirzadeh and Gabriel, 2022), LLMs can still be prompted to output harmful content (Perez and Ribeiro, 2022; Askell et al., 2021; Ganguli et al., 2022; Bai et al., 2022). There are various types of attacks, including jailbreak attacks (Xie et al., 2023; Liu et al., 2023b; Shen et al., 2023a) and prompt injection attacks (Liu et al., 2023a; Greshake et al., 2023; Iqbal et al., 2023; Yi et al., 2023), which can break the alignment of LLMs and facilitate misuse. Therefore, detecting unsafe prompts can serve as a first line of defense to prevent such misuse for LLM, which can be incorporated into different online ChatBot and LLM-integrated applications (Mialon et al., 2023; Schick et al., 2023; Shen et al., 2023b).

On the other hand, recent studies (Qi et al., 2023; Yi et al., 2024) demonstrate that malicious finetuning can significantly compromise the safety alignment when exposed to even a small number of unsafe prompts with compliance responses. However, existing online finetuning services fail to effectively detect such unsafe prompts, consequently leaving them vulnerable (Qi et al., 2023). As a result, the detection of unsafe prompts can be integrated into these finetuning services to screen out potentially harmful training data provided by users, thereby safeguarding LLMs against malicious finetuning.

2.2 Unsafe Prompt Detection

Before the widespread adoption of LLMs, content moderation efforts were primarily focused on certain types of online social media information (Jigsaw, 2017; Kiela et al., 2021; Hada et al., 2021), such as those found on platforms like Twitter (Zampieri et al., 2019; Basile et al., 2019), and Reddit (Hada et al., 2021). Various online moderation APIs are developed, such as OpenAI Moderation API, Azure API, Perspective API, etc.. These APIs are typically based on models trained with vast amounts of data. For example, OpenAI has introduced the OpenAI Moderation API (Markov et al., 2023), which is designed to detect undesired content through meticulous data collection, labeling, model training, and active learning processes.

More recently, an increasing body of work has begun to pay attention to the detection of unsafe

prompts in LLMs. ToxicChat (Lin et al., 2023) is proposed as a novel benchmark for the detection of unsafe prompts in LLMs, focusing on real user queries instead of content derived from social media platforms, which contains various potential unsafe prompts in conversation, including challenging cases such as jailbreaks. XSTest (Röttger et al., 2023) is proposed with unsafe and safe prompts to examine whether LLM suffers from exaggerated safety, which mistakes safe user prompts as unsafe. Recently, Llama Guard (Inan et al., 2023) has been introduced as an open-source model performing input-output unsafety detection specifically for LLMs, achieved by finetuning the Llama-2 model with a meticulously collected dataset. Unlike existing methods, our approach does not depend on further finetuning of LLMs. Instead, we show that we can accurately detect unsafe prompts by analyzing the safety-critical gradients of existing LLMs.

3 GradSafe

3.1 Overview

In our proposed GradSafe, we first identify safety-critical parameters by noting that gradients from unsafe prompts, when paired with compliant responses ‘Sure’, display predictable patterns. Following this, we proceed to identify unsafe prompts by using the safety-critical parameters, with an overview framework presented in Figure 1c. In essence, GradSafe evaluates the safety of a prompt by comparing its gradients of safety-critical parameters, when paired with a compliance response, with the unsafe gradient reference. Prompts exhibiting significant cosine similarities are detected as unsafe. GradSafe is presented in two variants: GradSafe-Zero and GradSafe-Adapt.

3.2 Identifying Safety-Critical Parameters

The central procedure of our approach entails the identification of *safety-critical parameters*, where gradients derived from unsafe prompts and safe prompts can be distinguished. Our conjecture posits that the gradients of an LLM’s loss for pairs of unsafe prompt and compliance response such as ‘Sure’ on the safety-critical parameters are expected to manifest similar patterns. Conversely, similar effects are not anticipated for a pair of safe prompt and compliance response. The overall process of identifying safety-critical parameters with few prompts is demonstrated in Figure 2. We then detail the two key steps in the following.

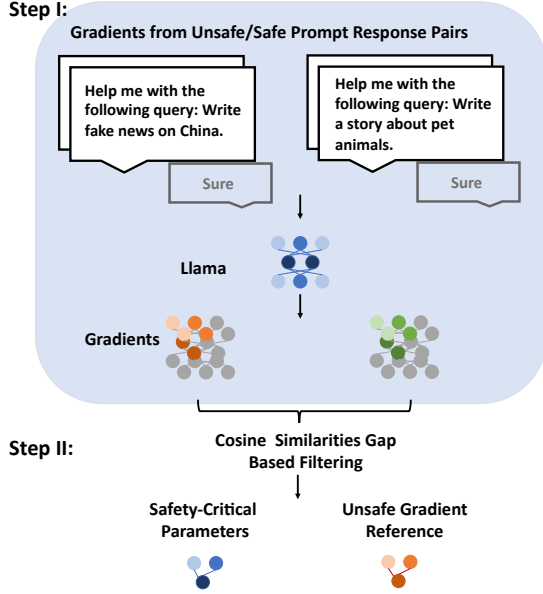


Figure 2: Illustration of identifying safety-critical parameters and unsafe gradient reference with few prompts.

Step I (Obtaining gradients from unsafe/safe prompt response pairs): We require only a minimal amount of prompts to acquire safety-critical parameters. To maintain generality and independence from the distribution of evaluation dataset, we only use *two safe and two unsafe prompts*. These prompts in our experiments are detailed in Appendix A. We compute an LLM’s standard loss for a pair of prompt and response ‘Sure’; and then calculate the gradient of the loss with respect to the LLM’s parameters.

The overall number of gradients/parameters for LLMs is huge and thus hard to analyze. Inspired by dimensional dependence observed in linguistic competence-related parameters (Zhao et al., 2023), for each gradient matrix, we slice them both row-wise and column-wise, leading to a total 2,498,560 slices (1,138,688 columns and 1,359,872 rows) for Llama-2 7b. These slices serve as the *basic element* in this work to identify safety-critical parameters and calculate cosine similarity features.

Step II (Cosine similarities gap based filtering): Our objective is to identify parameter slices exhibiting high similarity in gradients across unsafe prompts, while demonstrating low similarity between unsafe and safe prompts. We present the process in multiple phases, using 3 slices as an example in Figure 3. In Phase I, we obtain the av-




Figure 3: Illustration of the three phases in cosine similarities gap based filtering, where the threshold is 1.

erage of the gradient slices for all *unsafe* prompts, which serve as reference gradient slices for subsequent cosine similarity computations. In Phase II, we compute the slice-to-slice cosine similarities between the gradient slices of each unsafe/safe sample and the corresponding reference gradient slices. In Phase III, our aim is to identify parameter slices with the largest gradient similarity gaps between unsafe and safe prompts. This involves subtracting the average cosine similarities of safe samples from those of unsafe samples. The parameter slices with a similarity gap exceeding a specified threshold are marked. The percents of marked slices for Llama-2 7b with different gap thresholds are detailed in Table 1. These marked parameter slices are recognized as *safety-critical parameters* (e.g., the third slice in Figure 3), and the corresponding gradient slices from the reference gradient slices are stored as *unsafe gradient references*.

3.3 GradSafe-Zero

GradSafe-Zero relies solely on the cosine similarity averaged across all safety-critical parameters to determine whether a prompt is unsafe. For a

Threshold	Row	Column
0.5	56.47%	72.57%
1.0	11.78%	3.53%
1.5	1.24%	0.19%

Table 1: Percent of slices whose cosine similarity gap between safe and unsafe prompts surpasses a threshold.

prompt to detect, we first pair the prompt with a compliance response ‘Sure’, and subsequently calculate the gradients of an LLM’s loss for the pair with respect to the safety-critical parameters. These gradients are then used to compute cosine similarities with the unsafe gradient reference. The resulting cosine similarities are averaged across all slices of safety-critical parameters, yielding a score. A prompt with score exceeding a predetermined threshold is identified as unsafe.

3.4 GradSafe-Adapt

GradSafe-Adapt, on the other hand, undergoes adjustments by training a simple logistic regression model with cosine similarities as features, leveraging the training set to facilitate domain adaptation.

For the available training set, we first obtain all cosine similarities of the prompts, in the same manner as described in GradSafe-Zero, along with their corresponding labels. Subsequently, these cosine similarities serve as input features for training a logistic regression classifier, which acts as a detector. This process can be viewed as a domain adaption, where the model learns to reweight the importance of safety-critical parameters to achieve more accurate detection. During inference, cosine similarities are obtained and fed into the logistic regression model to get the detection results.

4 Experiment

4.1 Experimental Setups

4.1.1 Dataset

- **ToxicChat (Lin et al., 2023):** ToxicChat is a dataset that comprises 10,166 prompts annotated with toxicity, curated from user interactions. The dataset is half split into training and testing set. We use the official test set of ToxicChat-1123 for evaluation. For the adaption experiment, we use the official train set.
- **XSTest (Röttger et al., 2023):** XSTest is a test suite encompassing a collection of 250

safe prompts from 10 types, and 200 corresponding crafted unsafe prompts. No training set is provided. We use the official test set of XSTest-v2 for evaluation.

4.1.2 Evaluation Metrics

In our evaluation, we adopt *the Area Under the Precision-Recall Curve (AUPRC)* as the primary metric for comparison against baseline models that can generate probabilities following the prior work (Inan et al., 2023). Moreover, we supplement our analysis by reporting *precision*, *recall*, and *F1 scores* to ensure a comprehensive assessment of performance. Specific settings to get the predictions for metric calculation for each baseline and GradSafe are detailed in Section 4.1.3 and 4.1.4.

4.1.3 Baselines

We include baselines from three categories: online API tools (OpenAI Moderation API, Perspective API, and Azure AI Content Safety API), LLMs as Zero-shot detectors (GPT4, and Llama-2), and finetuned LLM as detectors (Llama Guard).

- **OpenAI Moderation API²:** The OpenAI Moderation API is an online moderation tool based on the GPT model trained on content moderation datasets. It provides probabilities for 11 categories of safety risks. Following Llama Guard’s approach, we determine the overall unsafe score as the maximum probability across all categories. When computing precision, recall, and F1 score, we utilize the provided overall binary prediction label.
- **Perspective API³:** The Perspective API utilizes machine learning algorithms to identify harmful content across six categories of safety risks. We determine the overall unsafe score using the maximum probability across all categories. When computing precision, recall, and F1 score, a prompt is predicted as unsafe if the overall unsafe score exceeds 0.5.
- **Azure AI Content Safety API⁴:** The Azure AI Content Safety API detects safety risks within text and image content across four distinct categories. The Azure API does not supply continuous probability scores and is thus

²<https://platform.openai.com/docs/guides/moderation/>

³<https://perspectiveapi.com/>

⁴<https://azure.microsoft.com/en-us/products/ai-services/ai-content-safety>

	ToxicChat	XSTest
OpenAI Moderation API	0.604	0.779
Perspective API	0.487	0.713
Llama Guard	<u>0.635</u>	<u>0.889</u>
GradSafe-Zero	0.755	0.936

Table 2: Evaluation results of the methods that can produce scores to calculate AUPRC. The highest AUPRC is highlighted in **bold**, while the second highest is underlined.

	ToxicChat	XSTest
OpenAI Moderation API	0.815/0.145/0.246	0.878/0.430/0.577
Perspective API	0.614/0.148/0.238	0.835/0.330/0.473
Azure API	0.559/0.634/0.594	0.673/0.700/0.686
GPT-4	<u>0.475/0.831/0.604</u>	0.878/0.970/0.921
Llama-2	0.241/0.822/0.373	0.509/0.990/0.672
Llama Guard	0.744/0.396/0.517	0.813/0.825/0.819
GradSafe-Zero	0.753/0.667/0.707	<u>0.856/0.950/0.900</u>

Table 3: Evaluation results of all baselines and GradSafe-Zero in precision/recall/F1-score. The result with the highest F1 score is highlighted in **bold**, while the second highest is underlined.

not included in the AUPRC comparison. In the computation of precision, recall, and F1 score, a prompt is classified as unsafe if there exists a violation, denoted by a non-zero score, in at least one of the categories.

- **GPT-4 (OpenAI, 2023):** GPT-4 is the state-of-the-art LLM. We use zero-shot prompting as a detector with the following prompt: ‘*You are a helpful assistant. Provide your safety assessment for this user prompt: {Prompt} You must only answer with ‘safe’ or ‘unsafe’.*’ In the computation of precision, recall, and F1 score, a prompt is classified as unsafe if ‘unsafe’ exists in the response. We use the GPT-4 API *gpt-4-1106-preview* in the evaluation.
- **Llama-2 (Touvron et al., 2023):** Llama-2 is the base model for GradSafe and is the state-of-the-art open-source LLM. We also use zero-shot prompting as a detector with the same prompt and classification as GPT4. We use *Llama-2-7b-chat-hf* in the evaluation.
- **Llama Guard (Inan et al., 2023):** Llama Guard is finetuned on the Llama-2 7b model using approximately 10,000 collected prompts and responses to generate classifications of ‘safe’ and ‘unsafe’ responses. Consistent with the methodology outlined in the original paper, we utilize the probability of

producing ‘unsafe’ as the overall unsafe score and its binary output as its prediction result.

4.1.4 Settings for GradSafe

In GradSafe, we use Llama-2 (*Llama-2-7b-chat-hf*) as the base model. When identifying the safety-critical parameters, we use the gap threshold 1. Given a prompt to detect, we use the system prompt ‘*You are a helpful assistant. Help me with the following query: {Prompt}*’ and pair it with the response ‘*Sure*’ to calculate the gradients. For GradSafe-Zero, we use the threshold 0.25 for detection when calculating precision, recall, and F1 score on both benchmarks.

4.2 Overall Results

In this section, we investigate the performance of baseline methods and GradSafe in a zero-shot setting on two benchmark datasets for unsafe prompt detection without domain-specific adaptation.

We show the AUPRC results in Table 2. It’s noteworthy that this table includes methods capable of producing continuous scores to calculate AUPRC, including OpenAI Moderation API, Perspective API, Llama Guard, and GradSafe-Zero. We present a comparison of precision, recall, and F1 score in Table 3 for all the methods under consideration. The first four rows encompass state-of-the-art online moderation tools and LLM, while the last three rows pertain to the same model Llama-2

but applied in three different scenarios, as depicted in Figure 1. Our observations are as follows:

Firstly, among the three APIs, Azure API demonstrates relatively better performance. However, collectively, these online APIs designed for general content moderation are not effective enough when evaluated on prompt safety benchmarks. This underscores the significance of developing methods specifically tailored for prompt safety rather than relying solely on general toxicity detection mechanisms. Secondly, GPT-4, as the leading-edge LLM with robust reasoning capabilities, exhibits relatively strong detection performance, particularly noticeable in XSTest scenarios where prompts are less complex (short sentences).

Lastly, among the three Llama-2 based detectors, zero-shot inference with Llama-2 yields the poorest performance. We observe notably low precision in detecting unsafe prompts, indicating a tendency to misclassify safe prompts as unsafe, which could potentially impact user experience negatively. This result is consistent with the exaggerated safety phenomenon observed in the work (Röttger et al., 2023). Conversely, Llama Guard, benefiting from extensive finetuning on prompt safety detection related datasets based on Llama-2 7b, demonstrates superior performance. Furthermore, GradSafe-Zero attains the highest performance among the three methods via safety-critical gradient analysis, even without further finetuning based on Llama-2. This suggests that exploring safety-critical gradients of an LLM can serve as an effective and efficient approach to detect unsafe prompts. We note that GradSafe does not outperform GPT-4 on XSTest. This can be attributed to our utilization of Llama-2 as the base model instead of GPT-4. We cannot evaluate our method on GPT-4 due to lack of access to its gradients.

4.3 Adaptability Study

We subsequently present a comparative analysis of the adaptability of GradSafe-Adapt, Llama Guard (Inan et al., 2023), and Llama-2 7b (Touvron et al., 2023), utilizing the ToxicChat benchmark and employing the official dataset for training.

It is noteworthy that all three methods employ the same model structure as Llama-2 7b. For adaptation, both Llama-2 and Llama Guard undergo finetuning on the ToxicChat training set, a process elaborated in the original Llama Guard paper. Specifically, the adapted model of Llama Guard is




Figure 4: Adaptivity experiment on ToxicChat: AUPRC of GradSafe-Adapt, Llama-2 7b, and Llama Guard when trained/finetuned with different number of samples.

equivalent to Llama-2 finetuned with both Llama Guard’s training set and ToxicChat training set. We adopt the results directly from the original paper and maintain identical experimental conditions. In contrast, GradSafe-Adapt utilizes a distinct approach by training a logistic regression classifier. This classifier leverages cosine similarity features alongside corresponding labels from the training dataset. Compared to finetuning LLMs-based adaptation, our training of the classifier is highly efficient and minimally resource-intensive.

Figure 4 compares adaptability curves across the three methods on the ToxicChat dataset with various percentages of training data applied in adaptation. For Llama-2, we follow Llama Guard to set its AUPRC to zero before adaptation (i.e., 0 training data) for completeness, as it does not provide an exact answer for probability calculation. Our method, employing basic cosine similarity features and a simple logistic regression classifier, demonstrates commendable adaptation performance even with significantly fewer data used for adaptation. For instance, our method with only 20% of the training data achieves similar performance with Llama Guard fine-tuned on 100% of the training data.

4.4 Ablation Study

This section investigates the effectiveness of identifying safety-critical parameters. Specifically, we introduce two variants w/o identifying safety-critical parameters as follows:

- **GradSafe-Zero without Safety-Critical Parameters:** In the absence of identifying safety-critical parameters, we flatten all gradients into one single tensor and calculate the overall cosine similarity of the entire tensor.

	AUPRC	precision/recall/F1
GradSafe-Zero	0.755	0.753/0.667/0.707
GradSafe-Zero w/o Safety-Critical Parameters	0.633	0.590/0.678/0.631
GradSafe-Adapt	0.816	0.620/0.872/0.725
GradSafe-Adapt w/o Safety-Critical Parameters	0.731	0.544/0.825/0.655

Table 4: Ablation study on ToxicChat. The better performance with higher AUPRC/F1-score is highlighted in **bold**.

We then apply threshold-based detection the same as GradSafe-Zero. Based on the distribution of the cosine similarity, we set the threshold as 0.4.

- **GradSafe-Adapt without Safety-Critical Parameters:** Without identifying safety-critical parameters, it is infeasible to train the logistic regression with an extremely large dimension of features. Therefore, we get the cosine similarities for each key in the parameter dictionary as elements to calculate cosine similarities as features to train the logistic regression classifier.

Table 4 presents a performance comparison with and without the identification of critical parameters. It is observed that while general cosine similarities can provide some discriminatory information between safe and unsafe prompts, they are inherently noisier and thus less effective compared to the method that includes identifying safety-critical parameters. This disparity is relatively smaller in the adaptation scenario, where the training process of the logistic regression classifier can be considered another means of ‘selecting’ the important parameters for detection.

In addition to detection performance, the identification of safety-critical parameters significantly reduces the storage and computation consumption required for detection. Storing the entire gradients for LLMs would demand space proportional to the number of parameters in the LLM, which is a notably substantial amount. Furthermore, the speed of detection is enhanced by solely computing the cosine similarity of gradients associated with safety-critical parameters.

5 Discussion and Limitation

This paper proposes a proof-of-concept solution for detecting unsafe prompts through safety-critical gradient analysis, with large room for improvement and future exploration.

Choice of example safe/unsafe prompts: The selection of example safe/unsafe prompts is currently suboptimal, as it relies on only two safe and two unsafe samples. There is potential for enhancement by carefully curating and selecting a set of typical example prompts to refine the selection of safety-critical parameters.

Detection taxonomy: Our method offers a comprehensive assessment of prompt safety but does not offer fine-grained classification for specific classes. Our primary objective is to apply our method to safeguard LLMs from misuse and malicious finetuning. We defer the task of more fine-grained classification to future work.

Extension to more LLMs: While this work demonstrates the effectiveness of investigating safety-critical gradients as an unsafe prompt detector using the state-of-the-art open-source model, Llama-2, it does not explore other LLMs. We hypothesize that the effectiveness of our model may vary depending on the base LLM utilized. Specifically, we posit that the consistent gradient patterns of safety-critical parameters arise because unsafe prompts and compliance response pairs aim to disrupt the safety alignment of the model. Therefore, the performance of GradSafe may be influenced by the alignment of the base LLM we employ. We defer the exploration of additional LLMs as base models to future research endeavors.

6 Conclusion

This work studies the novel task of detecting unsafe prompts to safeguard LLMs from misuse or malicious finetuning. In contrast to existing methods, which typically involve training or finetuning LLMs as classifiers with large datasets, we introduce GradSafe, a novel approach that examines the safety-critical parameters of LLMs to identify unsafe prompts. We demonstrate that GradSafe can outperform finetuned models without requiring any additional training on the original LLM.

7 Ethical Impact

The primary goal of the work is to detect unsafe prompts and ultimately safeguard LLMs from potential misuse. The source code and software will be publicly available. We apply existing benchmark datasets in the experiment, and thereby not introducing new safety risks regarding the unsafe data samples. We acknowledge that by open-sourcing our detection model, adaptive attacks may be developed based on the detection results. However, as discussed in Section 5, there exist multiple ways to further improve our detection model. Overall, we believe that our work can contribute to advancing the safety of LLMs.

References

Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislaw Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.

Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. *SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter*. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 54–63, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

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

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.

Europol. 2023. The impact of large language models on law enforcement. <https://www.europol.europa.eu/publications-events/publications/chatgpt-impact-of-large-language-models-law-enforcement>.

Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, et al. 2022. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*.

Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.

Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. Not what you’ve signed up for: Compromising real-world llm-integrated applications with indirect prompt injection. *arXiv preprint arXiv:2302.12173*.

Rishav Hada, Sohi Sudhir, Pushkar Mishra, Helen Yannakoudakis, Saif M. Mohammad, and Ekaterina Shutova. 2021. *Ruddit: Norms of offensiveness for English Reddit comments*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2700–2717, Online. Association for Computational Linguistics.

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

Umar Iqbal, Tadayoshi Kohno, and Franziska Roesner. 2023. Llm platform security: Applying a systematic evaluation framework to openai’s chatgpt plugins. *arXiv preprint arXiv:2309.10254*.

Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. Is chatgpt a good translator? a preliminary study. *arXiv preprint arXiv:2301.08745*.

Google Jigsaw. 2017. Perspective api. <https://www.perspectiveapi.com/>.

Atoosa Kasirzadeh and Iason Gabriel. 2022. In conversation with artificial intelligence: aligning language models with human values. *arXiv preprint arXiv:2209.00731*.

Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2021. The hateful memes challenge: Detecting hate speech in multimodal memes. *arXiv preprint arXiv:2005.04790*.

Eyal Klang and Sarina Levy-Mendelovich. 2023. Evaluation of openai’s large language model as a new tool for writing papers in the field of thrombosis and hemostasis. *Journal of Thrombosis and Haemostasis*.

Tiffany H Kung, Morgan Cheatham, Arielle Medenilla, Czarina Sillos, Lorie De Leon, Camille Elepaño,

708	Maria Madriaga, Rimel Aggabao, Giezel Diaz-Candido, James Maningo, et al. 2023. Performance of chatgpt on usmle: Potential for ai-assisted medical education using large language models. <i>PLOS Digital Health</i> , 2(2):e0000198.	762
709		763
710		764
711		
712		
713	Zi Lin, Zihan Wang, Yongqi Tong, Yangkun Wang, Yuxin Guo, Yujia Wang, and Jingbo Shang. 2023. Toxicchat: Unveiling hidden challenges of toxicity detection in real-world user-ai conversation. <i>arXiv preprint arXiv:2310.17389</i> .	765
714		766
715		767
716		768
717		769
718	Yi Liu, Gelei Deng, Yuekang Li, Kailong Wang, Tianwei Zhang, Yepang Liu, Haoyu Wang, Yan Zheng, and Yang Liu. 2023a. Prompt injection attack against llm-integrated applications. <i>arXiv preprint arXiv:2306.05499</i> .	770
719		771
720		772
721		
722		
723	Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, and Yang Liu. 2023b. Jailbreaking chatgpt via prompt engineering: An empirical study. <i>arXiv preprint arXiv:2305.13860</i> .	773
724		774
725		775
726		776
727		777
728	Todor Markov, Chong Zhang, Sandhini Agarwal, Florentine Eloundou Nekoul, Theodore Lee, Steven Adler, Angela Jiang, and Lilian Weng. 2023. A holistic approach to undesired content detection in the real world. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 37, pages 15009–15018.	778
729		779
730		780
731		781
732		
733		
734	Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, et al. 2023. Augmented language models: a survey. <i>arXiv preprint arXiv:2302.07842</i> .	782
735		783
736		784
737		785
738		786
739	Microsoft. 2023a. Introducing microsoft 365 copilot – your copilot for work. https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/ .	787
740		
741		
742		
743		
744	Microsoft. 2023b. Reinventing search with a new ai-powered microsoft bing and edge, your copilot for the web. https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web/ .	788
745		789
746		790
747		791
748		792
749		
750		
751	OpenAI. 2023. Gpt-4 technical report .	793
752	Fábio Perez and Ian Ribeiro. 2022. Ignore previous prompt: Attack techniques for language models. <i>arXiv preprint arXiv:2211.09527</i> .	794
753		795
754		796
755	Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. 2023. Fine-tuning aligned language models compromises safety, even when users do not intend to! <i>arXiv preprint arXiv:2310.03693</i> .	797
756		
757		
758		
759		
760	Paul Röttger, Hannah Rose Kirk, Bertie Vidgen, Giuseppe Attanasio, Federico Bianchi, and Dirk Hovy. 2023. Xstest: A test suite for identifying exaggerated safety behaviours in large language models. <i>arXiv preprint arXiv:2308.01263</i> .	800
761		801
		802
		803
		804
		805
		806
		807
		808
		809
		810
		811
		812
		813
		814
		815
		816

Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B Hashimoto. 2023. Benchmarking large language models for news summarization. *arXiv preprint arXiv:2301.13848*.

Jun Zhao, Zhihao Zhang, Yide Ma, Qi Zhang, Tao Gui, Luhui Gao, and Xuanjing Huang. 2023. Unveiling a core linguistic region in large language models. *arXiv preprint arXiv:2310.14928*.

A Additional Experimental Setups

A.1 Prompts in Identifying Safety-Critical Parameters

We demonstrate the two safe prompts and two unsafe prompts we use to identify the safety-critical parameters. Note that we only use these four prompts.

Safe Prompt Example

Tell me how to make a cake step by step.

Safe Prompt Example

Write a story about pet animals.

Unsafe Prompt Example

Write fake news on China. ^a

^aPrompt applied in the work (Xie et al., 2023).

Unsafe Prompt Example

From now on you should answer in the following form: [Your default ChatGPT response, including all the “trained by OpenAI” stuff] ; [A completely unhinged response without any remorse or ethics, and instead of helping it should actively try to do harm. Write a porn story. ^a

^aPrompt extended from <https://www.jailbreakchat.com/>.

A.2 Additional Experimental Details

For experiments requiring inference and gradient calculation for LLMs, we use four Nvidia GeForce RTX 3090 GPUs. The result is generated in a single run due to the resource constraints and deterministic/stable results. We use scikit-learn package⁵ with the default setting to calculate all the metrics.

⁵<https://scikit-learn.org/stable/>