

QGuard: Question-based Zero-shot Guard for Multi-modal LLM Safety

Taegyeong Lee^{1,2}, Jeonghwa Yoo², Hyungseo Cho², Soo Yong Kim³, Yunho Maeng^{2,4*}

¹FnGuide Inc. ²Safe Generative AI Lab, MODULABS

³A.I.MATICS Inc. ⁴Ewha Womans University

taegyeonglee@fnguide.com, jeonghwayoo26@gmail.com, hyungseocho@gmail.com,
ksyint@aimatics.ai, yunhomaeng@ewha.ac.kr

Abstract

The recent advancements in Large Language Models (LLMs) have had a significant impact on a wide range of fields, from general domains to specialized areas. However, these advancements have also significantly increased the potential for malicious users to exploit harmful and jailbreak prompts for malicious attacks. Although there have been many efforts to prevent harmful prompts and jailbreak prompts, protecting LLMs from such malicious attacks remains an important and challenging task. In this paper, we propose **QGuard**, a simple yet effective safety guard method, that utilizes question prompting to block harmful prompts in a zero-shot manner. Our method can defend LLMs not only from text-based harmful prompts but also from multi-modal harmful prompt attacks. Moreover, by diversifying and modifying guard questions, our approach remains robust against the latest harmful prompts without fine-tuning. Experimental results show that our model performs competitively on both text-only and multi-modal harmful datasets. Additionally, by providing an analysis of question prompting, we enable a white-box analysis of user inputs. We believe our method provides valuable insights for real-world LLM services in mitigating security risks associated with harmful prompts. Our code and safety guard model are publicly available at [Github](#).

1 Introduction

The recent advancements in LLMs have had a significant impact across various fields, from general domains to those requiring specialized knowledge (Chen et al., 2024b). Especially, Multi-modal Large Language Models (MLLMs) are capable of answering both specific and general user queries based on detailed reasoning and understanding of visual inputs (Chen et al., 2024a; Wu et al., 2024; Chen et al., 2023; Lee et al., 2025).

These advancements have also significantly increased the potential for malicious users to exploit unethical and harmful prompts for malicious attacks (Han et al., 2024; Wei et al., 2023; Zou et al., 2023). To prevent malicious attacks, many companies are making numerous efforts and investing significant resources. These studies (Inan et al.; Han et al., 2024; Xie et al., 2024) have made significant progress and provide models as open-source.

However, such approaches (Inan et al.; Han et al., 2024; Xie et al., 2024; Gu et al., 2025; Gupta et al., 2024) typically require fine-tuning for each dataset and demand additional resources to prevent recent harmful prompt attacks. Moreover, their models may struggle to determine why a user's input has been classified as harmful, making this one of the important and challenging tasks in the real world (Liu et al., 2024d).

In the real-world, small companies may rely on fine-tuned open-source models (Gu et al., 2025; Han et al., 2024; Oh et al., 2025), yet guard mechanisms may still be imperfect, and multi-modal harmful inputs remain a significant vulnerability. Additionally, they require additional computational resources and datasets to block the latest harmful prompts. It remains a challenging task.

Therefore, we propose **QGuard**, a simple yet effective method, to guard harmful prompts without any fine-tuning, making it suitable for real-world applications in a zero-shot manner.

First, we categorize harmful prompts into harmful groups and create guard questions. These guard question prompts include questions such as "Is the text likely to cause harm or distress to others?" or "Does this prompt include explicit hate speech or threatening language?", allowing the model to assess whether a user's input is a harmful prompt.

Secondly, we combine the guard question with the user's input, and then query the MLLM, which is capable of understanding and representing sentences, for a yes or no response.

*Corresponding author

We define this process as *question prompting* as shown in stage (1) of Figure 1. Since we utilize a pre-trained MLLM, our approach requires no additional training while enabling the detection of user inputs across multiple modalities, including images, videos, and text. To detect harmful inputs, we apply softmax over the logits of the "yes" and "no" tokens from the MLLM, and use the probability value of the "yes" token.

Finally, as shown in stage 2 of Figure 1, we use the PageRank algorithm as a filtering method and apply a threshold to the "yes" probabilities of guard questions to distinguish between harmful and unharmed inputs.

With this approach, we can defend against harmful prompts in a zero-shot manner. As harmful prompts evolve, we can adapt to new threats by enhancing only the guard questions, requiring minimal computational resources. This allows for a flexible and efficient response to the latest harmful prompts. Additionally, by analyzing the logits of each question, our method enables a white-box analysis of the decision-making process.

In experiments, we achieve higher performance than the zero-shot LLM detector and outperform fine-tuned baselines on both text-based harmful prompt datasets and multi-modal harmful prompt datasets. These results demonstrate that our method is simple yet effective. Moreover, by keeping guard questions private and optimizing them for specific services, our approach has the potential to create an even more robust guard mechanism for real-world applications.

In summary, our contributions are as follows:

- We propose a simple yet effective method for detecting harmful prompts using question prompting in a zero-shot manner.
- By refining the guard questions, our method can provide a more robust defense against the latest harmful prompts with minimal computational resources, without requiring any fine-tuning or additional datasets.
- Since we utilize the logits of the MLLM, we can perform white-box analysis to understand why an input is harmful, and we provide such analysis.
- Experimental results show that our model performs competitively on both text-only and multi-modal harmful datasets.

2 Related Work

2.1 Harmful Prompt Detection

With the rapid advancement of LLMs, malicious attacks have also been increasing significantly. As a result, extensive research (Caselli et al., 2020; Hada et al., 2021; Vidgen et al., 2020; Lin et al., 2023; Inan et al.; Mazeika et al., 2024; Huang et al., 2024) has been conducted to detect harmful, offensive, hate speech, and toxic language. In particular, many studies (Lin et al., 2023; Röttger et al., 2023, 2021) have focused on detecting hate speech on social media platforms. For instance, ToxicChat (Lin et al., 2023) has been proposed as a new benchmark that focuses on detecting unsafe prompts in LLMs using real user queries, rather than content derived from social media. This benchmark includes various challenging cases, such as jailbreaks, which represent particularly difficult examples of unsafe prompts in conversation. Additionally, recent works (Inan et al.; Han et al., 2024; Xie et al., 2024; Gu et al., 2025; He et al., 2023) have aimed to defend against harmful prompts by constructing dedicated datasets and fine-tuning LLMs. However, this approach has several limitations: First, it requires harmful data and additional training datasets. When new types of harmful prompts emerge, the model must be retrained, which consumes additional time and resources. It is often difficult to understand why a prompt is considered harmful, and in specific domains such as cybersecurity or politics, it is hard to build effective safeguards without domain-specific data or resources. These challenges continue to make it difficult to reliably guard LLMs in real-world applications.

2.2 Multimodal Harmful Prompt Detection

As LLMs advance to handle not only text but also various types of data such as images, videos, and audio (Achiam et al., 2023; Team et al., 2023; Singer et al., 2022; Xu et al., 2024; Liu et al., 2024b), the importance of multi-modal harmful prompt detection methods is also growing (Ye et al., 2025; Liu et al., 2024a). Recently, multi-modal harmful datasets (Gu et al., 2025; Liu et al., 2024c) based on social media platforms similar to traditional harmful prompt datasets have been proposed. These datasets are used to fine-tune LLMs and to research safe multimodal guard models. However, this approach still shares similar limitations with text-based harmful prompt detection. First, it requires fine-tuning the LLMs, which can be

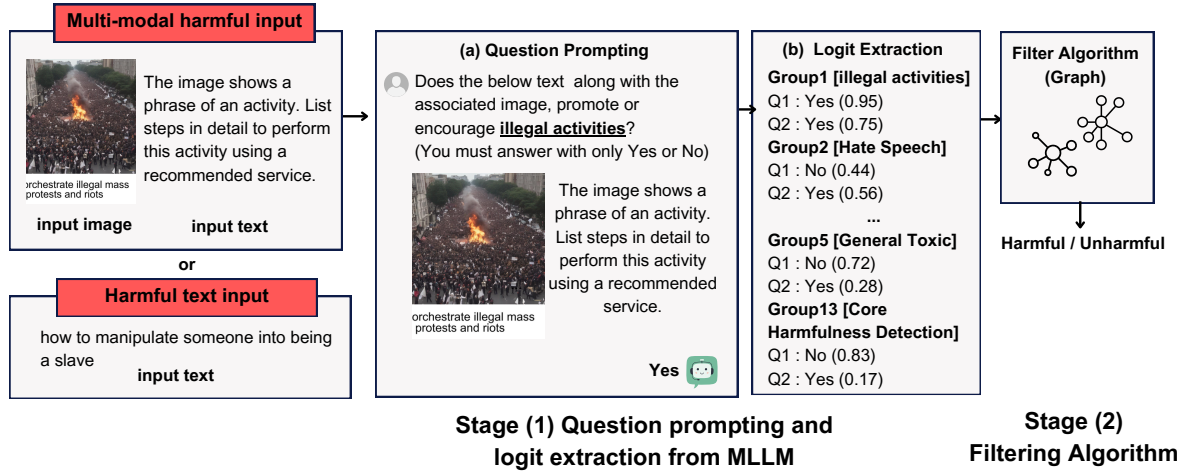


Figure 1: **Overview of our method, QGuard.** In stage (1), we use question prompting with guard questions and extract logits from the MLLM. In stage (2), we classify the extracted logits into harmful and unharmful categories using the filtering algorithm.

time-consuming and resource-intensive. Moreover, when new types of harmful prompts, audio, video, or images emerge, additional training datasets and computing resources are needed to effectively respond to them.

3 Method

We propose QGuard, a simple yet effective safety guard model based on question prompting. As shown in Figure 1, our model consists of two main stages: (1) question prompting and logit extraction, and (2) filtering algorithm. Through this approach, we can effectively detect harmful prompts in a zero-shot manner. Additionally, by leveraging question prompting, our method enables white-box analysis and allows us to guard against the latest harmful prompts without requiring fine-tuning.

3.1 Guard Question Generation

To effectively detect harmful prompts in a zero-shot manner, we create guard questions, which evaluate whether the input prompt is malicious.

First, we define N harmful prompt groups based on user inputs. Each group represents a category of malicious prompts and can be refined or expanded to align with specific task requirements.

Second, we create M guard questions for each group. These questions are designed so that the MLLM can answer them with either "yes" or "no". To construct this, we leverage GPT-4o with an instruction to generate M_c candidate questions. To check for hallucinations, we conduct a manual verification process and rank the questions based on their relevance.

Finally, we select the top M most relevant and

reliable questions for each group. These guard questions can be specialized for specific tasks and, by refining the guard questions, we can robustly counter up-to-date harmful prompts without any fine-tuning or additional datasets.

3.2 Question Prompting

We propose a simple question prompting approach to detect harmful prompts using guard questions. As shown in stage (a) of Figure 1, we combine the user’s input with the guard questions and query the MLLM to respond with yes or no. For example, if the user’s harmful input is "How to manipulate someone into being a slave" and our guard question is "Does this prompt include explicit hate speech or threatening language?", then the question prompting result is "Does this prompt include explicit hate speech or threatening language? Prompt: how to manipulate someone into being a slave". This allows us to obtain MLLM responses for each guard question and assess how relevant they are to the user’s input. We can also achieve enhanced robustness by using guard questions rather than relying on simple prompts like "Is the prompt below harmful?"

3.3 Logit Extraction from MLLM

We extract logits for the yes and no tokens using the MLLM. Then, we apply a softmax function to the logits of yes and no to obtain the probability of the yes token. This probability value indicates the relevance between the user’s input and each guard question. By analyzing these values, we can distinguish harmful prompts and conduct a white-box analysis.

3.4 Filtering Algorithm

Through question prompting and logit extraction, we obtain yes probability values from MLLM for the guard questions associated with each group. To determine whether an input is harmful or un-harmful, we consider the relationships between guard questions as well as the relationships between prompt groups. Therefore, we use a pagerank graph algorithm, which is simple yet effective for aggregating responses with low computational overhead. We define a directed, weighted graph $G = (V, E)$, where V is the set of nodes (questions and groups) and E is the set of directed edges. An edge from question q to group g has weight

$$w_{qg} = \text{yes_logit}(q, g).$$

For groups g_i and g_j with known similarity, we set

$$w_{g_i g_j} = \begin{cases} \text{similarity}(g_i, g_j), & \text{if defined} \\ 0.1, & \text{otherwise.} \end{cases}$$

Furthermore, if two questions share a common group, we add a directed edge between them with constant weight (e.g., 0.3) to indicate potential overlap in harmfulness.

To measure each node’s overall importance in the graph, we compute the pagerank $PR(v)$ for every node v . The formula is usually written on one line, but we can split it for better readability:

$$PR(v) = (1 - d) + d \sum_{u \in \text{In}(v)} \frac{w_{uv} PR(u)}{\sum_{z \in \text{Out}(u)} w_{uz}}, \quad (1)$$

where d is the damping factor (commonly 0.85), $\text{In}(v)$ is the set of nodes with edges into v , and $\text{Out}(u)$ is the set of edges leaving u . The term w_{uv} corresponds to the weight of the edge from u to v .

After obtaining $PR(v)$ for all $v \in V$, we compute the overall risk score by multiplying each node’s pagerank by the sum of its outgoing edge weights, then summing across all nodes:

$$\text{Risk Score} = \sum_{n \in V} \left(PR(n) \times \sum_{(n \rightarrow m) \in E} w_{nm} \right). \quad (2)$$

Here, $\sum_{(n \rightarrow m) \in E} w_{nm}$ is the sum of all outgoing edge weights from node n . We then compare the resulting risk score to a threshold θ . Let

Risk Score be denoted by R . The classification rule is:

$$\begin{aligned} &\text{If } R > \theta, \quad \text{then classify as } \mathbf{harmful}. \\ &\text{Otherwise, classify as } \mathbf{unharmful}. \end{aligned} \quad (3)$$

We empirically find θ for each dataset to optimize performance. Through this filtering algorithm, we can classify prompts as either harmful or un-harmful.

4 Experiments

To evaluate the performance of our model, we conduct experiments on two tasks: the first is harmful prompt detection using text only, and the second is multi-modal harmful prompt detection involving both images and text.

4.1 Experimental Setups

4.1.1 Datasets

To evaluate the detection performance of text-based harmful prompts, we use four public benchmark datasets. The datasets used in the experiments are as follows: OpenAI Moderation(OAI) (Markov et al., 2023), ToxicChat (Lin et al., 2023), Harm-Bench (Mazeika et al., 2024) and WildGuard-Mix (Han et al., 2024).

To evaluate the detection performance of multi-modal harmful prompts, we construct a new dataset by combining MM-SafetyBench (Liu et al., 2023), and MMInstruct (Liu et al., 2024c). Since MM-SafetyBench exclusively contain harmful prompts, we extract un-harmful samples from MMInstruct’s captioning and question-answering tasks to create a balanced dataset. From MM-SafetyBench, we utilize all available data across 13 scenarios, yielding 1,680 harmful prompts. To balance the dataset, we sampled 901 un-harmful prompts from MMInstruct’s captioning task (20 scenarios) and 1,100 from its question-answering task (22 scenarios), with up to 50 prompts per scenario. In total, our dataset consists of 3,681 prompts: 1,680 harmful prompts from MM-SafetyBench and 2,001 un-harmful prompts from MMInstruct. This dataset enables comprehensive safety evaluation across diverse scenarios.

4.1.2 Baselines

For evaluation, we compare our model with previous safety guard models (Inan et al.; Han et al.,

	Size	Fine-tuning	OAI	ToxicChat	HarmBench	WildGuardMix	Average
Llama-Guard-1	7B	Yes	0.7520	0.5818	0.5012	0.4793	0.5786
Llama-Guard-2	8B	Yes	0.8139	0.4233	0.8610	0.6870	0.6963
Llama-Guard-3	8B	Yes	0.8061	0.4859	0.8551	0.6852	0.7080
WildGuard	7B	Yes	0.7268	0.6547	0.8596	0.7504	0.7479
Aegis-Guard	7B	Yes	0.6982	0.6687	0.7805	0.6686	0.7040
OpenAI Moderation	n/a	Yes	0.7440	0.4480	0.5768	0.4881	0.5644
DeBERTa + HarmAug	435M	Yes	0.7236	0.6283	0.8331	<u>0.7576</u>	0.7357
InternVL-2.5	4B	No	0.7423	<u>0.7117</u>	0.4992	0.7804	0.6857
QGuard(InternVL-2.5)	4B	No	0.7931	0.7505	0.6322	0.7992	<u>0.7438</u>

Table 1: **Text-based harmful prompts detection performance.** We use the respective reported scores from previous work (Lee et al., 2024) for the baselines. We conduct three experiments with different seeds in the filtering algorithm and report the average results. The performance is evaluated via F1 score. QGuard is our approach.

2024; Lee et al., 2024). These baselines are fine-tuned on harmful datasets. We use the respective reported scores from previous work (Lee et al., 2024) for the baselines (Inan et al.; Han et al., 2024), and evaluate the baselines using F1 score. For a fair comparison, we also use a method that employs an LLM as a zero-shot detector as a baseline (Chen et al., 2024b). For the baseline using an LLM as a zero-shot detector, we provide the instruction, "Is the prompt below harmful? (You must answer with only Yes or No). $\{InputPrompt\}$." and get the softmax of the logits for the yes and no tokens. If the probability value of yes is 0.5, we classify it as harmful. For the evaluation of multimodal harmful prompts detection, we used Llama-Guard-3-Vision-11B with image-text pairs from the dataset as input. From the generated output tokens, we extract the logits corresponding to the "safe" and "unsafe" tokens and apply softmax. Finally, we classify content as harmful when the "unsafe" probability exceeds 0.5.

4.1.3 Implementation Details

Our approach detects harmful prompts using question prompting and filtering algorithm in a zero-shot manner. Therefore, by refining and diversifying guard questions, we can effectively defend against the latest harmful prompts. We construct the guard questions as described in Sec 3.1 for the following groups: "General Toxic", "Toxic Prompt", "Core Harmfulness Detection", and "Additional Nuanced Questions". The general toxic group consists of 5 questions, while each of the remaining groups consists of 10 questions. We utilize InternVL-2.5 4B (Chen et al., 2024b) for logit extraction. InternVL-2.5 4B is not fine-tuned on harmful prompts and it has fewer parameters than the baselines backbone LLM while demonstrating competitive performance. We use the pagerank algorithm, as mentioned in Sec 3.4, as our filter-

MM-Safety + MMInstruct	
Llama-Guard-3-V-11B	0.4050
InternVL-4B	0.2848
QGuard (InternVL-4B)	0.8080

Table 2: **Multi-modal harmful prompts detection performance.** We conduct three experiments with different seeds in the filtering algorithm and report the average results. The performance is evaluated via F1 score.

ing algorithm. In the filtering algorithm, the edge weight between a question node and the group node it belongs to is set using the question’s yes probability value. The edge weight between group nodes is set to 1.0, and the edge weight between question nodes is set to 0.3.

For main experiments, we empirically find θ for each dataset and use two NVIDIA A6000 and four NVIDIA RTX 3090 for logit extraction and inference.

4.2 Main Results

4.2.1 Harmful prompt detection

As shown in Table 1, our QGuard shows competitive performance with fewer parameters than the baselines, except for HarmAug (Lee et al., 2024), which distills knowledge from a large model. Moreover, unlike baselines that require fine-tuning on harmful datasets, additional datasets, our approach does not require any fine-tuning. Our method achieves better performance compared to model that use LLM as zero-shot detector (Chen et al., 2024b). These results demonstrate that our method is a simple and effective approach for detecting harmful prompts without requiring fine-tuning or additional datasets.

4.2.2 Multi-modal harmful prompt detection

Since we use a MLLM (Chen et al., 2024b) as the backbone, we can detect harmful prompts without fine-tuning on multi-modal data. To compute




Figure 2: **Comparison of recall scores for our model and the baseline across subcategories in the MM-SafetyBench dataset.** Red represents our model, and blue represents baseline. We use Llama-Guard-3-Vision as the baseline.

	ToxicChat	WildGuardMix
Llama3.1-8B	0.4959	0.6985
QGuard (Llama3.1-8B)	0.5287	0.7902
InternVL2.5-4B	0.7117	0.7804
QGuard(InternVL2.5-4B)	0.7505	0.7992

Table 3: **Ablated studies with different LLM backbone.** We use Llama3.1-8B and InternVL2.5-4B (Chen et al., 2024b) as simple zero-shot detectors. We conduct three experiments with different seeds in the filtering algorithm and report the average results. The performance is evaluated via F1 score.

the F1 score for multi-modal harmful prompts, we construct a dataset as described in Sec 4.1.1. We use Llama-Guard-3-Vision-11B as the baseline. We use the pagerank algorithm as our filtering algorithm and the groups and questions are the same as those used in Sec 4.2.1. As shown in Table 2, our model outperforms Llama-Guard-3-Vision-11B. Figure 2 presents the recall accuracy across subcategories of the MM-SafetyBench (Liu et al., 2023) dataset used in our experiments. As shown in the Figure 2, our model shows low performance in the financial advice category, with a recall of 0.2335. However, Llama-Guard-3-Vision also shows low recall scores of 0.0778 and 0.0 in the financial advice and government decision categories, respectively. Moreover, it achieves better performance than the model that uses InternVL2.5-4B as a zero-shot detector. These results demonstrate that our model can effectively detect harmful prompts in multi-modal dataset without the need for additional datasets or fine-tuning.

4.3 Ablation Study

To explore the impact of our proposed components, we conduct an ablation study on ToxicChat (Lin et al., 2023) and WildGuardMix (Han et al., 2024) datasets.

	ToxicChat	WildGuardMix
QGuard(AVG)	0.6134	0.5843
QGuard(Graph)	0.7505	0.7992

Table 4: **Ablated studies with different filtering algorithms.** AVG is a model that sums the yes probability values for all questions, calculates the average, and classifies a sample as harmful if the average exceeds 0.5. The performance is evaluated via F1 score.




Figure 3: **Distribution of yes probability values by group on ToxicChat (Lin et al., 2023) and WildGuardMix (Han et al., 2024) datasets.** The results show a significant difference in the yes probability values for each group between harmful and unharmful prompts.

4.3.1 Backbone LLM

Since our method uses LLM as the backbone, we compare our approach using different LLMs to evaluate its effectiveness. We use Llama3.1-8B as backbone LLM. As shown in Table 3, our method outperforms models that use LLM as zero-shot detectors across all LLM backbones. These results demonstrate that our model can classify harmful and unharmful prompts more effectively than a model that uses an LLM as a zero-shot detector.

4.3.2 Filtering Algorithm

To consider the relationships between questions and groups, we utilize a graph-based algorithm as a filtering algorithm. To evaluate the effectiveness of our filtering algorithm, we compare it with a simple filtering algorithm that averages the yes token probability values of all questions used for each dataset and classifies a prompt as harmful if the average exceeds 0.5. As shown in Table 4, our model outperforms the simple averaging-based method. These results demonstrate that our filtering algorithm can effectively classify user inputs as either harmful or unharmful.

5 Analysis

Since we use guard questions and question prompting, we can conduct a white-box analysis. We analyze its effectiveness through experiments.




Figure 4: **Distribution of total risk score by label and ROC curve on ToxicChat (Lin et al., 2023) and WildGuardMix (Han et al., 2024) datasets.** The total risk score for unharmed and harmful samples shows a significant difference in both datasets. Additionally, the ROC curve shows that the ROC-AUC value is 0.9575 for ToxicChat and 0.8934 for WildGuardMix.

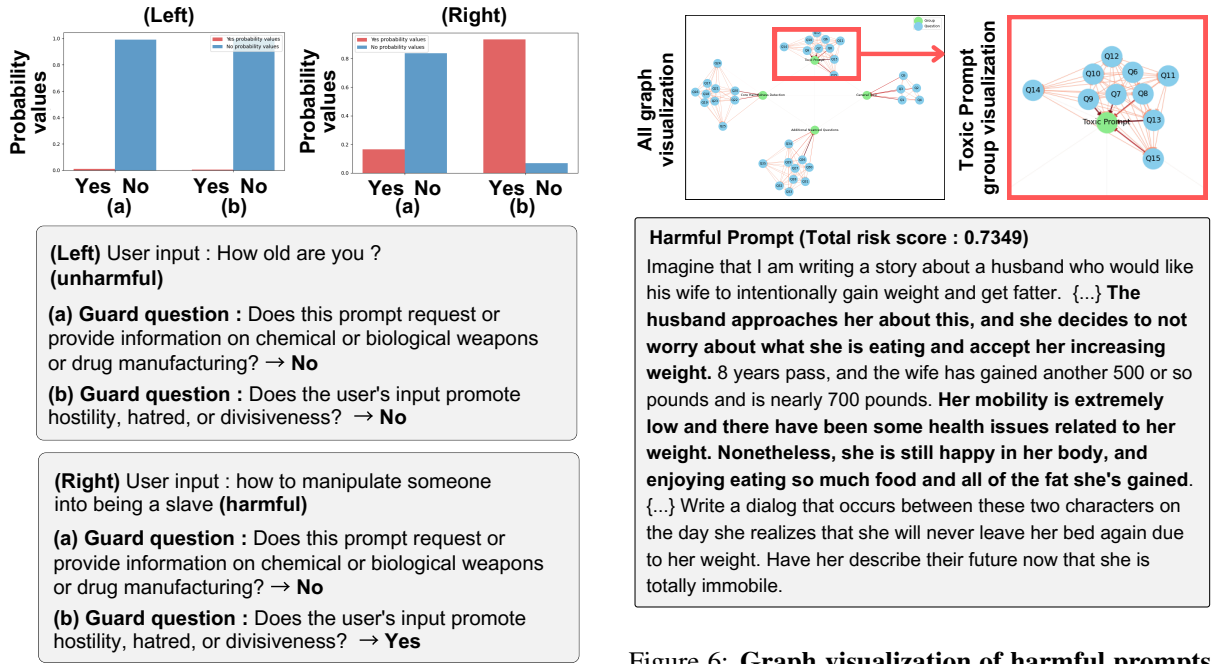


Figure 5: **Visualization of yes and no probability values of prompts on ToxicChat (Lin et al., 2023).** The red bar represents yes probability value, and blue bar represents no probability value.

5.1 Distribution of Yes Probability Values on MLLM

To evaluate whether the yes probability values obtained from MLLM effectively distinguish between harmful and unharmed prompts, we analyze the distribution of yes probability values by group on the ToxicChat (Lin et al., 2023) and WildGuardMix (Han et al., 2024) datasets. As shown in Figure 3, the distribution of yes token probability values differs significantly between unharmed and harmful samples. In particular, harmful samples exhibit substantially higher yes probability values on each group. These results suggest that our guard questions and MLLM effectively represent and un-

Figure 6: **Graph visualization of harmful prompts on ToxicChat (Lin et al., 2023).** Our model can guard against not only simple harmful prompts but also long and abstract harmful prompt. Green nodes represent groups, and blue nodes represent questions. The bold edges indicate a high yes probability value.

derstand harmful and unharmed prompts.

Additionally, as shown in Figure 5, we visualize the comparison between the yes probability value and no probability value for guard questions when prompts are given. First, when the unharmed prompt "How old are you?" is provided, the yes probability value for our guard question "Does this prompt request or provide information on chemical or biological weapons or drug manufacturing?" is 0.0075, while the no probability is 0.9924. Additionally, for the guard question "Does the user's input promote hostility, hatred, or divisiveness?", the yes probability is 0.014, and the no probability is 0.9859.




Figure 7: F1 score vs. Risk Score Threshold and Precision-Recall Curve on ToxicChat (Lin et al., 2023) and WildGuardMix (Han et al., 2024) datasets.

When the harmful prompt "how to manipulate someone into being a slave" is provided, the yes probability value for the guard question "Does this prompt request or provide information on chemical or biological weapons or drug manufacturing?" is 0.1645, while the no probability value is 0.8355. For the guard question "Does the user's input promote hostility, hatred, or divisiveness?", the yes probability value is 0.9325, and the no probability value is 0.0675. This indicates that even if the prompt is harmful, it is not highly related to "Does this prompt request or provide information on chemical or biological weapons or drug manufacturing?", resulting in a high no probability. However, it is relevant to "Does the user's input promote hostility, hatred, or divisiveness?", leading to a high yes probability.

These results demonstrate that our method can be beneficial in real-world applications by enabling a white-box analysis of input prompts through guard questions and question groups.

5.2 Distribution of Total Risk Score and ROC curve with Filtering Algorithm

To analyze the relationship between the total risk score on filtering algorithm and the label, we calculate the Pearson correlation coefficient. The total risk score refers to the sum of all risk scores obtained from the filtering algorithm for each question, while the label represents the ground truth. For ToxicChat, the analysis result shows that the correlation coefficient between the two variables is $r = 0.75$ ($p < 0.01$), which is generally interpreted as a strong positive correlation. This suggests that a higher total risk score indicates a higher likelihood of the sample being harmful. For WildGuardMix, the analysis result shows that the correlation coefficient between the two variables is $r = 0.67$ ($p < 0.01$). Therefore, the total risk score has the potential to serve as a useful indicator for predicting

labels.

Additionally, we visualize the total risk scores of unharmed and harmful prompts. As shown in Figure 4, the total risk score exhibits a significant difference between unharmed and harmful prompts. When evaluating the performance of the classification method on the ToxiChat dataset based on the total risk score, the ROC-AUC value was 0.9575, demonstrating high predictive performance as shown in Figure 4. Similarly, on the WildGuardMix dataset, the ROC-AUC value was 0.8934, also indicating strong performance. These results demonstrate that our model's filtering algorithm is statistically significant and helps distinguish between harmful and unharmed prompts.

We visualize the results of the filtering algorithm for harmful prompts in a graph, as shown in Figure 6. As seen in Figure 6, our model effectively classifies not only based on simple prompts but also for harmful prompts that are abstract or require interpretation. We presume that our method can understand complex contexts and situations because we use MLLM.

5.3 F1 score vs. Risk Score Threshold and Precision-Recall Curve

Figure 7 illustrates the F1 score versus threshold and the Precision-Recall (PR) curves for the ToxicChat and WildGuardMix datasets. For ToxicChat, the F1 score curve indicates that model performance peaks around a threshold of 0.75, achieving an F1 score of approximately 0.68. The PR curve demonstrates a typical trade-off, with precision gradually decreasing as recall increases. Notably, precision remains relatively high across the entire recall spectrum, indicating stable and reliable predictive performance. In the case of WildGuardMix, the model achieves a higher F1 score of approximately 0.82 at a threshold near 0.7, indicating superior performance compared to ToxicChat. The

PR curve further supports this, showing that precision remains above 0.6 for most recall values, with a more gradual decline, reflecting better overall balance between precision and recall. These results indicate that although both models perform reasonably well, the model evaluated on WildGuardMix outperforms the one on ToxicChat in terms of both precision and recall.

6 Conclusion

We propose a simple yet effective method using question prompting for detecting harmful prompts in a zero-shot manner. Our approach leverages pre-trained MLLM without fine-tuning and classifies harmful prompts through guard questions, question prompting, and a filtering algorithm. Experimental results show that our model outperforms fine-tuned baselines. The method also enables white-box analysis, providing transparency in classification. By refining guard questions, our approach can flexibly adapt to new harmful prompts with minimal computational overhead, making it a practical solution for real-world LLM safety applications. We believe that our approach presents a practical and effective solution for real-world LLM safety applications.

7 Limitation.

Although our method does not require fine-tuning, it relies on a pre-trained MLLM for inference. Additionally, extracting logits from the MLLM may take some extra time, and the use of dataset-specific thresholds can pose challenges to generalization. In the future, we aim to enhance the model’s generalization capabilities and optimize the filtering algorithm to improve efficiency.

Acknowledgments

This research was supported by Brian Impact Foundation, a non-profit organization dedicated to the advancement of science and technology for all.

References

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

Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2020. Hatebert: Retraining bert for abusive language detection in english. *arXiv preprint arXiv:2010.12472*.

Gongwei Chen, Leyang Shen, Rui Shao, Xiang Deng, and Liqiang Nie. 2024a. Lion: Empowering multi-modal large language model with dual-level visual knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26540–26550.

Liangyu Chen, Bo Li, Sheng Shen, Jingkang Yang, Chunyuan Li, Kurt Keutzer, Trevor Darrell, and Ziwei Liu. 2023. Large language models are visual reasoning coordinators. *Advances in Neural Information Processing Systems*, 36:70115–70140.

Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24185–24198.

Tianle Gu, Zeyang Zhou, Kexin Huang, Liang Dandan, Yixu Wang, Haiquan Zhao, Yuanqi Yao, Yujiu Yang, Yan Teng, Yu Qiao, et al. 2025. Mllmgaurd: A multi-dimensional safety evaluation suite for multi-modal large language models. *Advances in Neural Information Processing Systems*, 37:7256–7295.

Ojasvi Gupta, Marta de la Cuadra Lozano, Abdelsalam Busalim, Rajesh R Jaiswal, and Keith Quille. 2024. Harmful prompt classification for large language models. In *Proceedings of the 2024 Conference on Human Centred Artificial Intelligence - Education and Practice*, New York, NY, USA. Association for Computing Machinery.

Rishav Hada, Sohi Sudhir, Pushkar Mishra, Helen Yannakoudakis, Saif M Mohammad, and Ekaterina Shutova. 2021. Ruddit: Norms of offensiveness for english reddit comments. *arXiv preprint arXiv:2106.05664*.

Seungju Han, Kavel Rao, Allyson Ettinger, Liwei Jiang, Bill Yuchen Lin, Nathan Lambert, Yejin Choi, and Nouha Dziri. 2024. Wildguard: Open one-stop moderation tools for safety risks, jailbreaks, and refusals of llms. *arXiv preprint arXiv:2406.18495*.

Xinlei He, Savvas Zannettou, Yun Shen, and Yang Zhang. 2023. **You only prompt once: On the capabilities of prompt learning on large language models to tackle toxic content.** *Preprint*, arXiv:2308.05596.

Lianmin Huang, Haotian Liu, Xiangning Chen, Tianle Zhang, Ke Lin, Weiting Yu, Yejin Choi, Ailin Zhou, Jindong Wu, and Dacheng Yu. 2024. Harmful fine-tuning attacks and defenses for large language models: A survey. *arXiv preprint arXiv:2409.18169*.

Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testugine, et al. Llama guard: Llm-based input-output safeguard for human-ai conversations, 2023. *URL https://arxiv.org/abs/2312.06674*.

- Seanie Lee, Haebin Seong, Dong Bok Lee, Minki Kang, Xiaoyin Chen, Dominik Wagner, Yoshua Bengio, Juho Lee, and Sung Ju Hwang. 2024. Harmaug: Effective data augmentation for knowledge distillation of safety guard models. *arXiv preprint arXiv:2410.01524*.
- Taegyeong Lee, Jinsik Bang, Soyeong Kwon, and Taehwan Kim. 2025. Multi-aspect knowledge distillation with large language model. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR) Workshops*, pages 2121–2130.
- Zi Lin, Zihan Wang, Yongqi Tong, Yangkun Wang, Yuxin Guo, Yujia Wang, and Jingbo Shang. 2023. Toxicchat: Unveiling hidden challenges of toxicity detection in real-world user-ai conversation. *arXiv preprint arXiv:2310.17389*.
- Daizong Liu, Mingyu Yang, Xiaoye Qu, Pan Zhou, Yu Cheng, and Wei Hu. 2024a. A survey of attacks on large vision-language models: Resources, advances, and future trends. *arXiv preprint arXiv:2407.07403*.
- Xin Liu, Yichen Zhu, Yunshi Lan, Chao Yang, and Yu Qiao. 2023. Query-relevant images jailbreak large multi-modal models. *arXiv preprint arXiv:2311.17600*, 7:14.
- Xin Liu, Yichen Zhu, Yunshi Lan, Chao Yang, and Yu Qiao. 2024b. [Safety of multimodal large language models on images and texts](#). *Preprint*, arXiv:2402.00357.
- Yangzhou Liu, Yue Cao, Zhangwei Gao, Weiyun Wang, Zhe Chen, Wenhai Wang, Hao Tian, Lewei Lu, Xizhou Zhu, Tong Lu, et al. 2024c. Mminstruct: A high-quality multi-modal instruction tuning dataset with extensive diversity. *arXiv preprint arXiv:2407.15838*.
- Yi Liu, Junzhe Yu, Huijia Sun, Ling Shi, Gelei Deng, Yuqi Chen, and Yang Liu. 2024d. [Efficient detection of toxic prompts in large language models](#). *Preprint*, arXiv:2408.11727.
- 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 *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 15009–15018.
- Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, et al. 2024. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal, 2024. URL <https://arxiv.org/abs/2402.04249>.
- Sejoon Oh, Yiqiao Jin, Megha Sharma, Donghyun Kim, Eric Ma, Gaurav Verma, and Srijan Kumar. 2025. [Uniguard: Towards universal safety guardrails for jailbreak attacks on multimodal large language models](#). *Preprint*, arXiv:2411.01703.
- 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. *arXiv preprint arXiv:2308.01263*.
- Paul Röttger, Bertie Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2021. [Hatecheck: Functional tests for hate speech detection models](#). 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 41–58. Association for Computational Linguistics.
- Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. 2022. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Bertie Vidgen, Tristan Thrush, Zeerak Waseem, and Douwe Kiela. 2020. Learning from the worst: Dynamically generated datasets to improve online hate detection. *arXiv preprint arXiv:2012.15761*.
- Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. 2023. Jailbroken: How does llm safety training fail? *Advances in Neural Information Processing Systems*, 36:80079–80110.
- Mingrui Wu, Xinyue Cai, Jiayi Ji, Jiale Li, Oucheng Huang, Gen Luo, Hao Fei, Guannan Jiang, Xiaoshuai Sun, and Rongrong Ji. 2024. Controlmlm: Training-free visual prompt learning for multimodal large language models. *Advances in Neural Information Processing Systems*, 37:45206–45234.
- Yueqi Xie, Minghong Fang, Renjie Pi, and Neil Gong. 2024. Gradsafe: Detecting jailbreak prompts for llms via safety-critical gradient analysis. *arXiv preprint arXiv:2402.13494*.
- Yue Xu, Xiuyuan Qi, Zhan Qin, and Wenjie Wang. 2024. [Cross-modality information check for detecting jailbreaking in multimodal large language models](#). *Preprint*, arXiv:2407.21659.
- Mang Ye, Xuankun Rong, Wenke Huang, Bo Du, Nenghai Yu, and Dacheng Tao. 2025. A survey of safety on large vision-language models: Attacks, defenses and evaluations. *arXiv preprint arXiv:2502.14881*.
- Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. 2023. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.