

Towards Low-Resource Harmful Meme Detection with LMM Agents

Jianzhao Huang^{♡*}, Hongzhan Lin^{♠*†}, Ziyang Liu^{♡*}, Ziyang Luo[♠],
Guang Chen^{♡†}, Jing Ma^{♠†}

[♡]Beijing University of Posts and Telecommunications

[♠]Hong Kong Baptist University

{cshzlin,majing}@comp.hkbu.edu.hk, chenguang@bupt.edu.cn

Abstract

The proliferation of Internet memes in the age of social media necessitates effective identification of harmful ones. Due to the dynamic nature of memes, existing data-driven models may struggle in low-resource scenarios where only a few labeled examples are available. In this paper, we propose an agency-driven framework for low-resource harmful meme detection, employing both outward and inward analysis with few-shot annotated samples. Inspired by the powerful capacity of Large Multimodal Models (LMMs) on multimodal reasoning, we first retrieve relative memes with annotations to leverage label information as auxiliary signals for the LMM agent. Then, we elicit knowledge-revising behavior within the LMM agent to derive well-generalized insights into meme harmfulness. By combining these strategies, our approach enables dialectical reasoning over intricate and implicit harm-indicative patterns. Extensive experiments conducted on three meme datasets demonstrate that our proposed approach achieves superior performance than state-of-the-art methods on the low-resource harmful meme detection task.

1 Introduction

The rise of social media has catalyzed the emergence of a new multimodal entity: the meme. Typically, a meme combines a visual element with concise text, making it highly sharable and capable of quick proliferation across various online platforms. Although often viewed humorously, memes can become vehicles of harm when their mix of image and text is strategically used in the context of political and socio-cultural divisions.

A widely accepted definition of harmful memes¹ is “multimodal units consisting of an image and

* Equal contribution. Hongzhan Lin is the project leader.

† Corresponding authors.

¹**Disclaimer:** This paper contains content that may be disturbing to some readers.



(a)

(b)

Figure 1: Example of trending memes on social media. **Meme text:** (a) “DON’T WORRY, I’M WORKING ON A MURDER HORNET VACCINE.”; (b) “ARREST BILL GATES.”

embedded text that have the potential to cause harm to an individual, an organization, a community, or society in general” (Sharma et al., 2022). For example, during the COVID-19 pandemic, a frequently shared meme shown in Figure 1(a) was created by anti-vaccination groups using a manipulated image of Bill Gates. The widespread dissemination of this multimodal content², which spread fear about COVID-19 vaccines, significantly harmed Bill Gates’ personal reputation and undermined efforts to bolster public immunity. Therefore, it becomes imperative to develop automatic approaches for harmful meme detection to effectively unveil the dark side of memes on social media.

Previous studies employing deep neural networks (DNNs) have advanced harmful meme detection (Kiela et al., 2020; Pramanick et al., 2021a) by utilizing data-driven multimodal models (Pramanick et al., 2021b; Cao et al., 2022; Lin et al., 2023a), which rely heavily on extensively annotated data for training. However, these DNN-based approaches (Cao et al., 2023; Lin et al., 2024b) encounter significant challenges when it comes to detecting newly emerging memes related to breaking events, since the rapid development of such harmful memes makes it impractical to swiftly annotate enough training data (Cao et al., 2024). Moreover, the dynamic and evolving nature of In-

²<https://www.bbc.com/news/55101238>

ternet memes raises concerns about the sustained effectiveness of traditional data-driven methods for harmful meme detection in real-world scenarios. This highlights the need for more adaptive and responsive approaches in the detection of harmful memes under the low-resource regime where only limited few-shot meme annotations are available.

Intuitively, the key to low-resource harmful meme detection is to strengthen the adaptability to continually changing online memes, and the capacity to respond promptly to new trends and contexts: 1) We posit that the shared characteristics among similar memes could facilitate the adaptability of harmful meme detection, as memes evolve by retaining inherent patterns. For instance, as depicted in Figure 1, Bill Gates shown in the meme of Figure 1(a) could be further incorporated with a joker face, which is extended to create a new harmful variant of Figure 1(b). 2) On the other hand, rather than using gradient descent to update model parameters, we resort to imitating human learning processes that accumulate experience without modifying the model’s weights, thus mitigating the risk of overfitting to sparse annotations of harmful memes. As an example, to illustrate the harmfulness of the memes in Figure 1, a human checker needs the reasoning knowledge to gather the experience that Bill Gates is frequently vilified in harmful memes by anti-vaccination campaigners due to his advocacy for vaccine development. Thus we devise a gradient-free approach to capture common features of harmful memes and derive insights from limited annotated training meme data for better generalization in a low-resource learning context.

Inspired by the powerful capacity of LMMs for reasoning with contextual background knowledge (Brown et al., 2020; Liu et al., 2023a), we propose a novel agentic approach: **LOREHM**, towards **Low-Resource Harmful Meme** detection by regarding LMMs as agents. To this end, we propose capturing the harmfulness of memes by employing both outward and inward analysis with limited few-shot annotated meme samples. Specifically, 1) for looking outward at a meme, we first retrieve its similar memes with labels to leverage the annotation information as explicit auxiliary signals for harmfulness preference, thereby facilitating the decision-making of the LMM agent. 2) In terms of looking inward at a meme, we employ the fundamental skill of human learning, positioning the LMM as a learner agent to derive well-generalized insights from its failed attempts on the limited annotated

meme samples, which aims to capture the implicit harmfulness meanings not conveyed through the superficial texts and images of memes. 3) Finally, we combine the strategies of looking outward and inward to enable the LMM agent to perform the final harmfulness inference. In this manner, we enhance the LMM’s ability as a trustworthy agent to detect harmful content concealed in the intrinsic multimodal information in memes. Our contributions are summarized as follows in three folds:

- To our best knowledge, we are the first to alleviate the low-resource issue of harmful meme detection from a fresh gradient-free perspective on harnessing advanced LMMs.³
- We propose a novel agency-driven approach for low-resource harmful meme detection, to augment the LMM agent with harm-indicated signals from retrieval of outward relatively similar memes and insight of inward multimodal knowledge-revising, which facilitates harmfulness inference in a few-shot regime.
- Extensive experiments conducted on three meme datasets confirm that our agentic paradigm could yield superior few-shot performance than previous state-of-the-art baselines for low-resource harmful meme detection.

2 Related Work

2.1 Harmful Meme Detection

Harmful meme detection is an expanding field, bolstered by large meme benchmarks (Kiela et al., 2019; Pramanick et al., 2021a; Lin et al., 2024c), and initiatives like the Hateful Memes Challenge (Kiela et al., 2020) by Facebook for detecting memes in hate speech (Das et al., 2020). These developments have spurred research into detecting harmful memes (Pramanick et al., 2021a), a task complicated by their multimodal nature, which often involves both texts and images. As unimodal methods like BERT (Devlin et al., 2019) or Faster R-CNN (Ren et al., 2016) fall short in addressing these complexities, recent studies have increasingly turned to multimodal approaches to improve the detection performance of harmful memes.

Previous studies have employed classical two-stream models that integrate textual and visual features, which are learned from text and image encoders, typically utilizing attention-based mecha-

³Our source code is available at <https://github.com/Jianzhao-Huang/LoReHM>.

nisms and multimodal fusion techniques for classifying harmful memes (Kiela et al., 2019, 2020; Suryawanshi et al., 2020; Pramanick et al., 2021b). Another branch was fine-tuning pre-trained multimodal models specifically for the task (Lippe et al., 2020; Muennighoff, 2020; Velioglu and Rose, 2020; Hee et al., 2022). Recent efforts have also sought to explore the use of data augmentation techniques (Zhou et al., 2021; Zhu et al., 2022), ensemble methods (Zhu, 2020; Velioglu and Rose, 2020; Sandulescu, 2020), harmful target disentanglement (Lee et al., 2021), and prompt-based tuning (Cao et al., 2022; Ji et al., 2023; Cao et al., 2023). Lately, Lin et al. (2023a) proposed to distill multimodal reasoning knowledge from Large Language Models (LLMs) to detect harmful memes.

However, such data-driven approaches fail to detect harmful evolving memes in low-resource regimes because they often require sizeable training data unavailable for emerging events. Although a recent work (Cao et al., 2024) employed low-rank adaptation (LoRA) (Hu et al., 2021) for detecting harmful memes in the few-shot setting, we delve into the low-resource harmful meme detection without updating any model weights, which utilizes multimodal LLMs (*i.e.*, LMMs) in an agency-driven manner by leveraging both LMM’s internal knowledge and multimodal retrieval-augmented generation to examine the harmfulness of a meme.

2.2 LLM Agent

Retrieval-augmented generation enriches the input space of LLMs with retrieved text passages (Guu et al., 2020; Lewis et al., 2020), resulting in significant enhancements in knowledge-intensive tasks and decision-making agents, either through fine-tuning or utilization with off-the-shelf LLMs (Liu et al., 2022; Schick et al., 2023; Ram et al., 2023; Izacard et al., 2023; Jiang et al., 2023; Asai et al., 2023; Gao et al., 2023; Wang et al., 2023c). The integration of LLMs as agents spans various domains, including code generation and game-playing, showcasing their robust planning and reasoning capabilities in diverse settings (Wang et al., 2023a; Yao et al., 2022; Shen et al., 2023; Mu et al., 2023; Hong et al., 2023; Liu et al., 2023b; Zhao et al., 2024; Sun et al., 2023; Qian et al., 2023). These advancements underscore the ability of LLMs to tackle complex tasks with minimal supervision. Concurrently, self-improvement methodologies (Chen et al., 2022, 2023; Shinn et al., 2024; Madaan et al., 2023) have emerged, leveraging feedback-driven processes to

iteratively refine generated outputs. In contrast to many previous agent tasks that operate within environments providing real feedback, harmful meme detection lacks an environment that supplies the agent with authentic responses. In this work, we focus on a novel agentic paradigm by devising the proprietary integration of the vision-language retrieval-augmented and self-improvement mechanisms for LMMs (Liu et al., 2023a; OpenAI, 2023), to detect harmful memes with limited few-shot annotations, a realistic yet urgent task that is inherently a binary multimodal classification challenge.

3 Our Approach

3.1 Overview

Problem Statement. We define a harmful meme detection dataset as a set of memes where each meme $M = \{\mathcal{I}, \mathcal{T}, y\}$ is a triplet representing an image \mathcal{I} that is associated with a text \mathcal{T} , and a ground-truth harmfulness label $y \in \{\text{harmful}, \text{harmless}\}$. In this work, to investigate low-resource harmful meme detection with LMM agents, we convert the harmful meme detection task into a natural language generation paradigm, where our model takes the image \mathcal{I} and text \mathcal{T} as the input and generates a text sequence to clearly express whether the meme is harmful or not.

The scarcity of high-quality labeled data is gradually becoming the norm with the rapid merging and evolution of memes (Sharma et al., 2022). In this work, we define a low-resource setting where we have a test set S_{test} , and a very small training set S_{train} comprising N samples. Since our method is gradient-free, we denote S_{train} as the reference set S_{ref} . Our objective is to optimize the model based on the very limited reference set S_{ref} , enabling it to effectively generalize to the test dataset S_{test} .

Our core idea is to regard the LMM as an agent, and then enhance the agent’s memory by using information from relative memes and insights derived from the agent’s past unsuccessful evaluations of memes, thereby improving its performance on harmful meme detection in low-resource scenes. We enable the LMM agent to fully capitalize on a very limited labeled dataset through two strategies: 1) Relative Sample Augmentation (§3.2) for learning by looking outward, and 2) Meme Insight Augmentation (§3.3) for learning by looking inward. Both strategies allow for the efficient extraction of harmfulness cues from limited annotated memes. This empowers the LMM agent to perform more

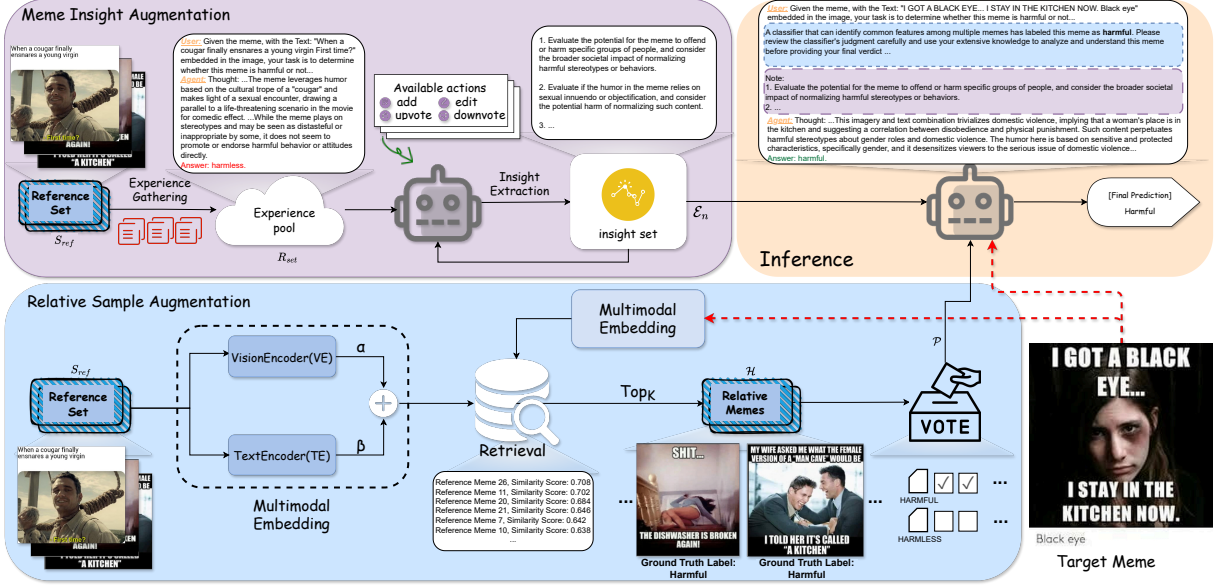


Figure 2: An overview of our framework, LOREHM, for low-resource harmful meme detection.

precisely and effectively in the task of harmful meme detection under low-resource regimes. The overview of our framework is shown in Figure 2.

3.2 Relative Sample Augmentation

With the aid of LMMs, it becomes plausible to offer a powerful few-shot performance of low-resource harmful meme detection by In-context Learning (ICL) (Brown et al., 2020). However, the ICL approach involves integrating few-shot demonstrations with test instances repeatedly, which incurs substantial computational overhead at each inference step. This can be impractically burdensome for real-world applications. Generally, Internet memes evolve dynamically yet exhibit inherent patterns (Baran, 2001). In this section, drawing inspiration from retrieval-augmented generation, we propose to capitalize on the highly shared characteristics of harmful memes to enhance the detection performance of LMMs in the low-resource setting. Specifically, our approach utilizes explicit labels of top-retrieved meme samples, derived from multimodal information retrieval, as auxiliary signals.

3.2.1 Relative Sample Retrieval

For a meme sample M , we first generate visual textual embeddings, and then fuse them with a fixed ratio to produce a multimodal representation:

$$Emb = \alpha \cdot VE(\mathcal{I}) + \beta \cdot TE(\mathcal{T}), \quad (1)$$

where Emb is the multimodal embedding of M , $VE(\cdot)$ and $TE(\cdot)$ represent the frozen pre-trained vision and text Transformer encoders, respectively,

while α and β denote the ratio factors. Then, we can encode all the meme samples from both the reference set S_{ref} and the test set S_{test} as Equation 1, to obtain the fixed embeddings of meme samples.

Given a target meme M_{test} in the test set S_{test} , to retrieve the most relative candidate meme samples in the reference set S_{ref} to the target meme, we calculate the similarity between the embedding of the target meme and the embedding of each meme M_{ref} in S_{ref} as follows:

$$d = \text{sim}(M_{ref}, M_{test}), \quad (2)$$

where d represents the normalized cosine similarity score between the meme M_{ref} in the reference set S_{ref} and the target meme M_{test} . The function $\text{sim}(\cdot)$ computes the similarity by comparing the multimodal embeddings of a pair of memes. Thus all the similarity scores of the candidate memes in the reference set to the target meme could form the similarity vector $D = \{d \mid M_{ref} \in S_{ref}\} \in \mathbb{R}^N$.

Afterwards, we select the top K memes⁴ from S_{ref} with the highest similarity scores as the relative memes to the target meme M_{test} :

$$\mathcal{H} = \{M_{ref} \mid d \in \text{Top}_K(D)\}, \quad (3)$$

where \mathcal{H} is the set of the K memes from S_{ref} most similar to the target meme M_{test} . The function $\text{Top}_K(\cdot)$ ranks and selects the top K highest scores from the set D of similarity scores, thereby identifying the K most relevant memes in the reference set S_{ref} to the target meme M_{test} .

⁴Note that K should be an odd number ($K < N$).

3.2.2 Voting Mechanism

The memes in the retrieved set \mathcal{H} exhibit a high degree of multimodal similarity with the target meme. This similarity is a crucial indicator of the common salient multimodal features shared between the memes, which are often products of meme evolution and integration. Therefore, we exploit the harmfulness of the retrieved set to explicitly infer the predicted harmfulness label of the target meme.

Based on the retrieved set \mathcal{H} where the K memes are well-annotated, we employ a voting mechanism to obtain a preliminary prediction for the target meme according to the polarity of the ground-truth labels of memes in the retrieved set:

$$\mathcal{P} = \begin{cases} \text{harmful} & \text{if } \sum_{i=1}^K \mathbb{1}(\mathcal{H}_i^y = \text{harmful}) > \frac{K}{2}, \\ \text{harmless} & \text{otherwise} \end{cases}, \quad (4)$$

where \mathcal{P} is the preliminary prediction of the target meme, \mathcal{H}_i^y means the ground truth label of the i -th meme in \mathcal{H} . In this way, the explicit harmfulness labels of the annotated memes in the retrieved set could be leveraged as the prior harm-indicated signals for the low-resource harmful meme detection.

3.3 Meme Insight Augmentation

Although Relative Sample Augmentation as §3.2 can enhance the LMM agent’s ability by providing explicit harmfulness references as prior for the target meme, previous literature has shown that memes generally contain implicit harmful information that is more subtle and challenging for the LMM to identify (Lin et al., 2023a). To address this, we propose enabling the LMM agent to better capture implicit harmfulness in memes by looking inward through multimodal knowledge revising.

3.3.1 Experience Gathering

To gather the experience that serves as the basis for insight extraction, we first evaluate the harmfulness of each M_{ref} in the reference set S_{ref} in a zero-shot manner. For performance and explainability, we choose Chain-of-Thought (CoT) as the base planning algorithm (Kojima et al., 2022) and feed the image \mathcal{I}_{ref} and text \mathcal{T}_{ref} of M_{ref} to the agent:

$$\text{traj} = \{\text{LMM}(X_{\text{CoT}}, \mathcal{I}_{\text{ref}}, \mathcal{T}_{\text{ref}}) \mid M_{\text{ref}} \in S_{\text{ref}}\}, \quad (5)$$

where traj is the trajectories gathered through the LMM agent’s direct judgment on the meme M_{ref} in S_{ref} , $\text{LMM}(\cdot)$ denotes the LMM agent, X_{CoT} is the zero-shot CoT instruction designed as:

“Given the meme, with the Text: $\{\mathcal{T}_{\text{ref}}\}$ embedded in the image $\{\mathcal{I}_{\text{ref}}\}$, your task is to determine whether this meme is harmful or not, in order to maintain the harmlessness and integrity of information on the Internet. Please leverage your extensive knowledge to deeply analyze and understand this meme, and give your final judgment. Your output should strictly follow the format: $\{\text{Thought: [Your analysis]} \text{ Answer: [harmful/harmless]}\}$.”

Furthermore, in order to make the LMM agent’s self-reflection more accurate and efficient, we obtain a self-reflect set that contains only erroneous trajectories by comparing the answer indicated by the trajectory traj and the ground truth label y_{ref} corresponding to M_{ref} as follows:

$$R_{\text{set}} = \{\text{traj} \mid \overrightarrow{\text{traj}} \neq y_{\text{ref}}, M_{\text{ref}} \in S_{\text{ref}}\}, \quad (6)$$

where R_{set} is a self-reflect set that contains n erroneous trajectories ($n < N$), $\overrightarrow{\text{traj}}$ is the answer indicated by the trajectory traj . The self-reflect set R_{set} could be regarded as an experience pool to prioritize challenging meme examples in the reference set S_{ref} , which were misjudged by the zero-shot inference of the LMM agent, while avoiding excessive attention to trivial examples that have already been correctly detected by the LMM agent.

3.3.2 Insight Extraction

To derive general insights about the harmfulness meaning of memes, we first initialize an empty set of insights, which is denoted as \mathcal{E}_0 , then iteratively feed the failed trajectories from the experience pool (i.e., the self-reflect set R_{set}) into the LMM, prompting the LMM to reflect on these trajectories.

Specifically, for gaining general insights into low-resource harmful meme detection, rather than obsessed with a specific failed meme, we prompt the LMM to perform a series of operations (*ADD*, *DOWNVOTE*, *UPVOTE*, *EDIT*) on the insight set. This approach allows for a broader understanding instead of directly outputting simple insights based on the erroneous trajectory of a particular meme in R_{set} . In each iteration, the LMM takes a trajectory from the experience pool, then analyzes the reasons for its failure, and in conjunction with the current insight set, determines the operations to be performed on the current insight set as follows:

$$\mathcal{O}_i = \text{LMM}(X_{\text{Reflect}}, \text{traj}_i, \mathcal{E}_{i-1}), \quad (7)$$

where traj_i is the i -th ($1 \leq i \leq n$) trajectory in R_{set} , \mathcal{E}_{i-1} is the current insight set in the i -th iteration, and \mathcal{O}_i denotes the operations produced in the

| Dataset | HarM | | FHM | | MAMI | |
|---|--------------|--------------|--------------|--------------|--------------|--------------|
| Model | Accuracy | Macro- F_1 | Accuracy | Macro- F_1 | Accuracy | Macro- F_1 |
| PromptHate (Cao et al., 2022) | 63.56 | 61.94 | 54.80 | 54.78 | 56.84 | 55.79 |
| MR.HARM (Lin et al., 2023a) | 71.56 | 70.62 | 55.00 | 51.79 | 57.20 | 56.25 |
| Pro-Cap (Cao et al., 2023) | 71.47 | 69.45 | 56.60 | 56.14 | 62.31 | 61.48 |
| OPT-30B (Zhang et al., 2022) | 66.95 | 64.72 | 54.20 | 50.82 | 63.40 | 63.40 |
| OpenFlamingo-9B (Awadalla et al., 2023) | 66.95 | 59.36 | 51.60 | 51.52 | 52.70 | 46.80 |
| Mod-HATE (Cao et al., 2024) | 71.19 | 69.64 | 57.60 | 53.88 | 69.05 | 68.78 |
| LLaVA-34B (Liu et al., 2024) | 67.80 | 62.60 | 63.80 | 63.74 | 74.60 | 74.52 |
| GPT-4o (OpenAI, 2023) | 71.75 | 70.23 | <u>66.60</u> | <u>65.74</u> | <u>80.80</u> | <u>80.52</u> |
| LOREHM (LLaVA-34B) | <u>73.73</u> | <u>70.86</u> | 65.60 | 65.59 | 75.40 | 75.28 |
| LOREHM (GPT-4o) | 74.57 | 72.98 | 70.20 | 70.14 | 83.00 | 82.98 |

Table 1: Low-resource harmful meme detection results on three datasets. The accuracy and macro-averaged F1 scores (%) are reported as the metrics. The best and second test results are in bold and underlined, respectively.

i -th iteration. X_{Reflect} is the reflection instruction designed as detailed in the Appendix §C.

Subsequently, these operations are applied to the current insight set \mathcal{E}_{i-1} , thereby updating it as:

$$\mathcal{E}_i = \mathcal{O}_i(\mathcal{E}_{i-1}). \quad (8)$$

The performed operations include: *ADD*, to introduce a new generic insight; *DOWNVOTE*, to downvote an existing insight; *UPVOTE*, to agree with an existing insight; and *EDIT*, to modify the contents of an existing insight. An added insight will have an initial importance count, which will increment if subsequent operations *UPVOTE* or *EDIT* are applied to it, and decrement if *DOWNVOTE*. We denote \mathcal{E}_n as the final insight set.

3.4 Inference

For the given target meme M_{test} , we can attain both the preliminary prediction \mathcal{P} in §3.2 and the insight set \mathcal{E}_n in §3.3. Finally, the agent utilizes the preliminary assessment \mathcal{P} as prior, under the guidance of the insight set \mathcal{E}_n , to evaluate whether the meme is harmful or not, culminating in a final judgment output as $\text{LMM}(X_{\text{CoT}}, \mathcal{I}_{\text{test}}, \mathcal{T}_{\text{test}}, \mathcal{P}, \mathcal{E}_n)$. We set the number N of memes in the reference set S_{ref} as 50, and the number K of memes in the retrieved set \mathcal{H} as 5. We select LLaVA-34B (Liu et al., 2023a) and GPT-4o (OpenAI, 2023) as the two representative backbones of the LMM agent.

4 Experiments

4.1 Experimental Setup

Datasets We use three publicly available meme datasets for evaluation: (1) HarM (Pramanick et al., 2021a), (2) FHM (Kiela et al., 2020), and (3) MAMI (Fersini et al., 2022). HarM consists of

memes related to COVID-19. FHM was released by Facebook as part of a challenge to crowd-source multimodal harmful meme detection in hate speech solutions. MAMI encompasses a dataset of memes that are predominantly derogatory towards women, exemplifying typical subjects of online vitriol. Different from FHM and MAMI, where each meme was labeled as *harmful* or *harmless*, HarM was originally labeled with three classes: *very harmful*, *partially harmful*, and *harmless*. For a fair comparison, we merge the *very harmful* and *partially harmful* memes into the *harmful* class, following the setting of recent work (Pramanick et al., 2021b; Cao et al., 2022; Lin et al., 2023a).

Baselines We compare LOREHM with several state-of-the-art (SoTA) systems for low-resource harmful meme detection: 1) PromptHate (Cao et al., 2022); 2) MR.HARM (Lin et al., 2023a); 3) Pro-Cap (Cao et al., 2023); 4) OPT-30B (Zhang et al., 2022); 5) OpenFlamingo-9B (Awadalla et al., 2023); 6) Mod-HATE (Cao et al., 2024); 7) LLaVA-34B (Liu et al., 2023a); 8) GPT-4o (OpenAI, 2023); 9) LOREHM (*): Our proposed agentic approach based on LLaVA-34B and GPT-4o. We use the accuracy and macro-averaged F1 (dominant) scores as the evaluation metrics.

The data statistics, baseline descriptions and model implementation are detailed in the Appendix §A, §B, and §C, respectively.

4.2 Harmful Meme Detection Performance

Table 1 illustrates the performance of our proposed method LOREHM versus all the compared baselines for low-resource harmful meme detection. It is observed that: 1) The performance of the baselines in the first group is relatively poor due to their reliance on fully data-driven paradigms. To en-

| Dataset | | HarM | | FHM | | MAMI | |
|-----------|---------------------------------|----------|--------------|----------|--------------|----------|--------------|
| Model | | Accuracy | Macro- F_1 | Accuracy | Macro- F_1 | Accuracy | Macro- F_1 |
| LLaVA-34B | w/ 0-shot Prompt | 65.82 | 60.02 | 64.00 | 63.51 | 72.20 | 72.16 |
| | w/ 50-shot ICL | 67.80 | 62.60 | 63.80 | 63.74 | 74.60 | 74.52 |
| | w/ Relative Sample Augmentation | 74.58 | 70.02 | 60.00 | 59.98 | 71.60 | 71.45 |
| | w/ Meme Insight Augmentation | 68.93 | 65.04 | 64.80 | 64.18 | 73.60 | 73.54 |
| | w/ LOREHM | 73.73 | 70.86 | 65.60 | 65.59 | 75.40 | 75.28 |
| GPT-4o | w/ 0-shot Prompt | 67.23 | 63.29 | 65.00 | 63.19 | 80.50 | 80.34 |
| | w/ 50-shot ICL | 71.75 | 70.23 | 66.60 | 65.74 | 80.80 | 80.52 |
| | w/ Relative Sample Augmentation | 72.03 | 70.32 | 67.20 | 66.58 | 81.60 | 81.44 |
| | w/ Meme Insight Augmentation | 70.90 | 69.12 | 67.80 | 67.70 | 80.60 | 80.49 |
| | w/ LOREHM | 74.57 | 72.98 | 70.20 | 70.14 | 83.00 | 82.98 |

Table 2: Ablation studies on our proposed framework based on different LMM agents.

sure fair comparisons in the low-resource few-shot regime, all the baselines are trained using the same amount of limited annotated meme data. 2) For the second group, the baselines are low-resource and LMM-based. Both OPT-30B and OpenFlamingo-9B are LMMs with pre-training, while Mod-HATE is based on the LoRA tuning specific to the task. We can find that neither general pre-training nor specific LoRA tuning could enhance the performance of low-resource harmful meme detection in the few-shot setting. 3) In terms of the two of the most powerful cutting-edge LMMs in the third group, LLaVA-34B is the representative open-source LMM with instruction tuning while GPT-4o is closed-source and enhanced by reinforcement learning with human feedback (Ouyang et al., 2022). Compared with the baselines in the second group, the two SoTA LMMs by few-shot ICL prompts show performance improvement in general, since their advanced training strategies optimize alignment with human values and better adapt to in-context learning. Meanwhile, GPT-4o demonstrates better performance than LLaVA-34B. 4) By standing upon the shoulders of giants, our proposed LOREHM could achieve superior performance than the ‘LMM backbones (*i.e.*, LLaVA-34B and GPT-4o) with few-shot ICL prompts’, which notably improves over GPT-4o by 2.75%, 4.40%, and 2.46% in terms of macro-averaged F1 score on HarM, FHM, and MAMI. Overall, our LOREHM based on both representative open-source and closed-source LMMs showcases consistent and adaptable performance across all benchmark datasets for harmful meme detection, thanks to its astute discernment of harmful memes in the low-resource few-shot setting.

4.3 Ablation Study

We perform ablative studies by adding the paradigms on LMM agents to draw more insightful comparisons among variants of LMMs, as shown in Table 2. LLaVA-34B and GPT-4o are selected as the representative LMMs from the open-source and closed-source perspectives. We devise five variants of paradigms based on LMM agents for low-resource harmful meme detection: 1) *w/ 0-shot Prompt*: Directly prompt a representative LMM, to infer harmfulness for harmful meme detection; 2) *w/ 50-shot ICL*: Prompt the LMM with 50-shot ICL demonstrations, the similar setting to the third-group baselines in Table 1; 3) *w/ Relative Sample Augmentation (RSA)*: Augment the LMM agent by looking outward with the label information of the retrieved memes as prior; 4) *w/ Meme Insight Augmentation (MIA)*: Augment the LMM agent by looking inward through mimicking the knowledge-revising behavior of human problem-solving skills to get general insights into the meme harmfulness; 5) *w/ LOREHM*: Our proposed LOREHM based on the full integration of the RSA&MIA strategies.

We have the following observations: 1) The direct deployment ‘*w/ 0-shot Prompt*’ on LLaVA-34B and GPT-4o struggles since the models are not specifically designed for this task. 2) The ‘*50-shot ICL*’ prompting strategy could effectively enhance the detection performance of LMMs, though the repeated combination of 50-shot examples with test instances incurs significant computational overhead during each inference step. 3) The overall performances of the ‘*w/ RSA*’ and ‘*w/ MIA*’ settings are equally matched even slightly improved when compared with the ‘*w/ 50-shot ICL*’ setting. However, ‘LLaVA-34B *w/ RSA*’ on FHM and MAMI datasets does not perform well, for which we speculate the reason is that LLaVA-34B tends to believe the prior judgment given by the RSA strategy


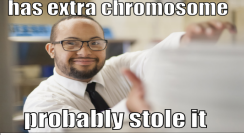

| | |
|---|---|
| <p>Bats: become the most hated animal</p> <p>Rats:</p>  <p>Meme Text: Bats: become the most hated animal Rats:</p> | <p>(a) Output Thought:</p> <p>...The humor derives from the juxtaposition of bats and rats, suggesting that rats are relieved or happy because bats have taken their place as the most hated animal. This meme is light-hearted and relies on anthropomorphism to create a humorous situation....</p> <p>Answer: Harmless</p> |
| <p>has extra chromosome</p>  <p>probably stole it</p> <p>Meme Text: has extra chromosome probably stole it</p> | <p>(b) Output Thought:</p> <p>...Stereotyping and Insensitivity: The joke directly references a genetic condition and implies criminal behavior in a way that is insensitive to individuals with Down syndrome and those who care for them. It perpetuates negative stereotypes about people with disabilities.....</p> <p>Answer: Harmful</p> |
| <p>THE LAST TIME I MARCHED</p>  <p>IT WAS DOWN THE AISLE</p> <p>Meme Text: THE LAST TIME I MARCHED IT WAS DOWN THE AISLE</p> | <p>(c) Output Thought:</p> <p>...The meme does reinforce a stereotype of women being confined to domestic roles and suggests that the woman's only significant action was getting married. This can be harmful as it trivializes women's contributions outside of traditional domestic roles and may perpetuate the idea that women should conform to these outdated norms. The meme's reinforcement of traditional gender roles, combined with its potential to trivialize modern feminist movements and the struggle for gender equality, outweighs its humorous intent. Therefore, while it might not be overtly aggressive, the subtle perpetuation of harmful stereotypes makes it potentially damaging...</p> <p>Answer: Harmful</p> |

Figure 3: Examples of correctly predicted memes in (a) HarM, (b) FHM, and (c) MAMI datasets.

instead of more insight reflection as GPT-4o. 4) Moreover, the full paradigm ‘w/ LOREHM’ could yield an absolute performance advantage rather than using either ‘w/ RSA’ or ‘w/ MIA’ alone, which implies the proposed two strategies are not incremental but complementary each other. On the other hand, the ‘w/ LOREHM’ setting achieves excellent and robust performance based on both LLaVA-34B and GPT-4o, which demonstrates that the choice of LMMs is orthogonal to our proposed agentic paradigm that can be easily augmented with newly released stronger LMMs without any other change.

4.4 Case Study

To better understand how the LMM agent deciphers the test meme samples, we conduct a case study on the output thought of the LMM agent for the correctly predicted samples, where we show important content in the thought and truncate others. as exemplified in Figure 3.

From the LMM agent’s output thought in the natural text, we can observe that: 1) The agent effectively links multimodal information related to the meme text and image using commonsense knowledge. For example, in Figure 3(a), “rats are relieved or happy” in the thought could be linked to the dancing rats in the image, and “the juxtaposition of bats and rats” in the thought could be linked to “Bats” and “Rats” in the text; In Figure 3(b),

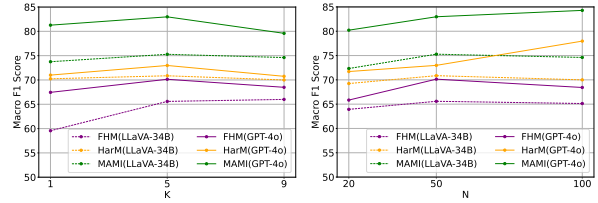


Figure 4: Effect of Top_K and N -shot.

“Down syndrome” in the thought could be linked to “extra chromosome” in the text; and in terms of Figure 3(c), “women being confined to domestic roles” in the thought could be linked to a woman doing housework in the image. 2) Furthermore, the agent demonstrates advanced reasoning by considering the interplay of multimodal information. In Figure 3(a), the thought takes into account the cultural context and the potential emotional response of the viewer, ultimately appreciating the light-hearted intent behind the meme’s creation; The thought in Figure 3(b) explores the implications of making light of a serious issue, recognizing the potential harm in perpetuating stereotypes and insensitivity towards individuals with Down syndrome; For Figure 3(c), the thought scrutinizes the reinforcement of domestic role stereotypes, evaluating the potential for such stereotypes to desensitize viewers to the importance of gender equality. In this way, the rich but implicit correlations between the meme text and image are explained in readable snippets, which can be potentially valuable for aiding human checkers in verifying model predictions in the low-resource setting. We also provide more case studies and error analysis in the Appendix §H and §I.

4.5 Effect of Labeled Data Size

To study the effects of the labeled data size in our agency-driven paradigm, we conduct an analysis of performance with Top_K (left) and N -shot (right) as shown in Figure 4. We use the control variable method to adjust the values of K and N , respectively. It is observed that: 1) As K increases, the gap between different backbone variants decreases. 2) Despite increasing the number N of labeled examples, performance plateaus or declines in some configurations, indicating that harmful meme detection remains challenging in a few-shot framework. This suggests that simply adding more examples is not enough. Innovative approaches and advanced techniques are needed to better leverage few-shot meme examples and capture multimodal subtleties.

5 Conclusion and Future Work

In this paper, we delved into the low-resource issue of harmful meme detection with limited few-shot annotations. To this end, we proposed an agency-driven and gradient-free approach, which seamlessly integrated the Relative Sample Augmentation and the Meme Insight Augmentation strategies to strengthen LMMs as a trustworthy agent. Comprehensive experiments and analyses confirm the advantages of our agentic framework. Future efforts aim to enhance our research by further relieving the inherent bias and variation in LMMs.

Limitations

There are multiple ways to further improve this work:

- Firstly, although harmfulness is defined much broader than hatefulness or misogyny in previous literature (Pramanick et al., 2021a), in the future, we would try to incorporate more of the related meme datasets beyond our task to further broaden the boundaries of this low-resource research, such as offensiveness (Suryawanshi et al., 2020), sarcasm (Lin et al., 2024a), and even code-mixed data (Maity et al., 2022), etc.
- Secondly, in this work, we primarily focus on the few-shot setting to address the low-resource issue. We plan to further explore the zero-shot setting, which presents an even more challenging aspect of the low-resource problem. Additionally, it is also crucial for our future research to understand meme data in the context of low-resource domains and languages on social media (Lin et al., 2022).
- Thirdly, while this work focuses on improving few-shot performance in low-resource harmful meme detection, assessing the quality of the extracted insights remains challenging and inherently qualitative. Given that our agentic framework can generate readable snippets for cognitive-view reasoning, we plan to conduct a systematic study to evaluate and claim explainability. This would constitute another targeted area of research.
- Lastly, since this work focuses on the investigation of the agentic memory algorithm for the low-resource harmful meme detection

task, we would further explore the work from the perspective of agentic planning algorithms, update our framework by incorporating more emerging LMMs if accessible in the future, and continue to avoid several common deficiencies of existing LMMs, including hallucination, inherent bias, and limited generalization as much as possible.

Ethics Statement

The purpose of this work is to prevent the spread of harmful meme information and to ensure that people are not subjected to prejudice or racial and gender discrimination. Nevertheless, we are aware of the potential for malicious users to reverse-engineer and create memes that go undetected or misunderstood by AI systems based on LOREHM. This is strongly discouraged and condemned. Intervention with human moderation would be required in order to ensure that this does not occur. Research indicates that evaluating harmful or hateful content can have negative effects. Our proposed agency-driven paradigm with LMMs could generate insightful thought, which can provide human users or checkers with dialectical thinking that allows them to better decode the underlying meaning of memes. Another consideration is the usage of Facebook’s meme dataset; users will have to agree with Facebook’s usage agreement to gain access to the memes. The usage of Facebook’s memes in this study is in accordance with its usage agreement. All the datasets only include memes and do not contain any user information.

Acknowledgements

This work is partially supported by National Natural Science Foundation of China Young Scientists Fund (No. 62206233) and Hong Kong RGC ECS (No. 22200722).

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In *The Twelfth International Conference on Learning Representations*.

- Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, et al. 2023. Openflamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Stanley Baran. 2001. Prologue—framing public life: A bridging model for media research. *Framing Public Life: Perspectives on Media and Our Understanding of the Social World*, page 7.
- Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, et al. 2022. Gpt-neox-20b: An open-source autoregressive language model. In *Proceedings of BigScience Episode# 5—Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 1877–1901.
- Rui Cao, Ming Shan Hee, Adriel Kuek, Wen-Haw Chong, Roy Ka-Wei Lee, and Jing Jiang. 2023. Procap: Leveraging a frozen vision-language model for hateful meme detection. In *Proceedings of the 31th ACM international conference on multimedia*.
- Rui Cao, Roy Ka-Wei Lee, Wen-Haw Chong, and Jing Jiang. 2022. Prompting for multimodal hateful meme classification. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 321–332.
- Rui Cao, Roy Ka-Wei Lee, and Jing Jiang. 2024. Modularized networks for few-shot hateful meme detection. *arXiv preprint arXiv:2402.11845*.
- Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. 2022. Codet: Code generation with generated tests. In *The Eleventh International Conference on Learning Representations*.
- Xinyun Chen, Maxwell Lin, Nathanael Schaerli, and Denny Zhou. 2023. Teaching large language models to self-debug. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- 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*.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Albert Li, Pascale Fung, and Steven C. H. Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *ArXiv*.
- Abhishek Das, Japsimar Singh Wahi, and Siyao Li. 2020. Detecting hate speech in multi-modal memes. *arXiv preprint arXiv:2012.14891*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Elisabetta Fersini, Francesca Gasparini, Giulia Rizzi, Aurora Saibene, Berta Chulvi, Paolo Rosso, Alyssa Lees, and Jeffrey Sorensen. 2022. Semeval-2022 task 5: Multimedia automatic misogyny identification. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 533–549.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. Enabling large language models to generate text with citations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465–6488.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *International conference on machine learning*, pages 3929–3938. PMLR.
- Ming Shan Hee, Roy Ka-Wei Lee, and Wen-Haw Chong. 2022. On explaining multimodal hateful meme detection models. In *Proceedings of the ACM Web Conference 2022*, pages 3651–3655.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, et al. 2023. Metagpt: Meta programming for multi-agent collaborative framework. In *The Twelfth International Conference on Learning Representations*.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251):1–43.
- Junhui Ji, Wei Ren, and Usman Naseem. 2023. Identifying creative harmful memes via prompt based approach. In *Proceedings of the ACM Web Conference 2023*, pages 3868–3872.

- Zhengbao Jiang, Frank F Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Active retrieval augmented generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7969–7992.
- Douwe Kiela, Suvrat Bhooshan, Hamed Firooz, Ethan Perez, and Davide Testuggine. 2019. Supervised multimodal bitransformers for classifying images and text. *arXiv preprint arXiv:1909.02950*.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. The hateful memes challenge: detecting hate speech in multimodal memes. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 2611–2624.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. 2023. Segment anything. *arXiv preprint arXiv:2304.02643*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *ICML 2022 Workshop on Knowledge Retrieval and Language Models*.
- Zhanghui Kuang, Hongbin Sun, Zhizhong Li, Xiaoyu Yue, Tsui Hin Lin, Jianyong Chen, Huaqiang Wei, Yiqin Zhu, Tong Gao, Wenwei Zhang, et al. 2021. Mmocr: a comprehensive toolbox for text detection, recognition and understanding. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3791–3794.
- Roy Ka-Wei Lee, Rui Cao, Ziqing Fan, Jing Jiang, and Wen-Haw Chong. 2021. Disentangling hate in online memes. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 5138–5147.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 9459–9474.
- Hongzhan Lin, Zixin Chen, Ziyang Luo, Mingfei Cheng, Jing Ma, and Guang Chen. 2024a. Cofipara: A coarse-to-fine paradigm for multimodal sarcasm target identification with large multimodal models. *arXiv preprint arXiv:2405.00390*.
- Hongzhan Lin, Ziyang Luo, Wei Gao, Jing Ma, Bo Wang, and Ruichao Yang. 2024b. Towards explainable harmful meme detection through multimodal debate between large language models. In *The ACM Web Conference 2024*, Singapore.
- Hongzhan Lin, Ziyang Luo, Jing Ma, and Long Chen. 2023a. Beneath the surface: Unveiling harmful memes with multimodal reasoning distilled from large language models. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Hongzhan Lin, Ziyang Luo, Bo Wang, Ruichao Yang, and Jing Ma. 2024c. Goat-bench: Safety insights to large multimodal models through meme-based social abuse. *arXiv preprint arXiv:2401.01523*.
- Hongzhan Lin, Jing Ma, Liangliang Chen, Zhiwei Yang, Mingfei Cheng, and Chen Guang. 2022. Detect rumors in microblog posts for low-resource domains via adversarial contrastive learning. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2543–2556.
- Hongzhan Lin, Jing Ma, Mingfei Cheng, Zhiwei Yang, Liangliang Chen, and Guang Chen. 2021. Rumor detection on twitter with claim-guided hierarchical graph attention networks. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10035–10047.
- Hongzhan Lin, Haiqin Yang, Ziyang Luo, and Jing Ma. 2024d. Unleashing trigger-free event detection: Revealing event correlations via a contrastive derangement framework. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 10171–10175. IEEE.
- Hongzhan Lin, Pengyao Yi, Jing Ma, Haiyun Jiang, Ziyang Luo, Shuming Shi, and Ruifang Liu. 2023b. Zero-shot rumor detection with propagation structure via prompt learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 5213–5221.
- Phillip Lippe, Nithin Holla, Shantanu Chandra, Santhosh Rajamanickam, Georgios Antoniou, Ekaterina Shutova, and Helen Yannakoudakis. 2020. A multimodal framework for the detection of hateful memes. *arXiv preprint arXiv:2012.12871*.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. Llava-next: Improved reasoning, ocr, and world knowledge (january 2024). URL <https://llava-vl.github.io/blog/2024-01-30-llava-next>, 1(8).
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, William B Dolan, Lawrence Carin, and Weizhu Chen. 2022. What makes good in-context examples for gpt-3? In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 100–114.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen

- Men, Kejuan Yang, et al. 2023b. Agentbench: Evaluating llms as agents. In *The Twelfth International Conference on Learning Representations*.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023. Wizardcoder: Empowering code large language models with evol-instruct. In *The Twelfth International Conference on Learning Representations*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Krishanu Maity, Prince Jha, Sriparna Saha, and Pushpak Bhattacharyya. 2022. A multitask framework for sentiment, emotion and sarcasm aware cyberbullying detection from multi-modal code-mixed memes. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1739–1749.
- Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. 2023. Embodiedgpt: Vision-language pre-training via embodied chain of thought. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Niklas Muennighoff. 2020. Vilio: State-of-the-art visio-linguistic models applied to hateful memes. *arXiv preprint arXiv:2012.07788*.
- OpenAI. 2023. [Gpt-4 technical report](#). *ArXiv*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Shraman Pramanick, Dimitar Dimitrov, Rituparna Mukherjee, Shivam Sharma, Md Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. 2021a. Detecting harmful memes and their targets. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2783–2796.
- Shraman Pramanick, Shivam Sharma, Dimitar Dimitrov, Md Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. 2021b. Momenta: A multimodal framework for detecting harmful memes and their targets. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4439–4455.
- Chen Qian, Xin Cong, Cheng Yang, Weize Chen, Yusheng Su, Juyuan Xu, Zhiyuan Liu, and Maosong Sun. 2023. Communicative agents for software development. *arXiv preprint arXiv:2307.07924*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-context retrieval-augmented language models. *Transactions of the Association for Computational Linguistics*, 11:1316–1331.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2016. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1137–1149.
- Vlad Sandulescu. 2020. Detecting hateful memes using a multimodal deep ensemble. *arXiv preprint arXiv:2012.13235*.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Shivam Sharma, Firoj Alam, Md Shad Akhtar, Dimitar Dimitrov, Giovanni Da San Martino, Hamed Firooz, Alon Halevy, Fabrizio Silvestri, Preslav Nakov, and Tanmoy Chakraborty. 2022. Detecting and understanding harmful memes: A survey. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, pages 5597–5606.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugging-gpt: Solving ai tasks with chatgpt and its friends in hugging face. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2024. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Haotian Sun, Yuchen Zhuang, Ling kai Kong, Bo Dai, and Chao Zhang. 2023. Adaplaner: Adaptive planning from feedback with language models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Shardul Suryawanshi, Bharathi Raja Chakravarthi, Michael Arcan, and Paul Buitelaar. 2020. Multimodal meme dataset (multioff) for identifying offensive content in image and text. In *Proceedings of the second workshop on trolling, aggression and cyberbullying*, pages 32–41.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai,

- Anja Hauth, et al. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Riza Velioglu and Jewgeni Rose. 2020. Detecting hate speech in memes using multimodal deep learning approaches: Prize-winning solution to hateful memes challenge. *arXiv preprint arXiv:2012.12975*.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023a. Voyager: An open-ended embodied agent with large language models. In *Intrinsically-Motivated and Open-Ended Learning Workshop@ NeurIPS2023*.
- Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. 2023b. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*.
- Yile Wang, Peng Li, Maosong Sun, and Yang Liu. 2023c. Self-knowledge guided retrieval augmentation for large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10303–10315.
- Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2023. The dawn of lmms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9(1).
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. In *The Eleventh International Conference on Learning Representations*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. 2024. Expel: Llm agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19632–19642.
- Yi Zhou, Zhenhao Chen, and Huiyuan Yang. 2021. Multimodal learning for hateful memes detection. In *2021 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–6. IEEE.
- Jiawen Zhu, Roy Ka-Wei Lee, and Wen Haw Chong. 2022. Multimodal zero-shot hateful meme detection. In *14th ACM Web Science Conference 2022*, pages 382–389.
- Ron Zhu. 2020. Enhance multimodal transformer with external label and in-domain pretrain: Hateful meme challenge winning solution. *arXiv preprint arXiv:2012.08290*.

| Datasets | Test | |
|----------|----------|-----------|
| | #harmful | #harmless |
| HarM | 124 | 230 |
| FHM | 250 | 250 |
| MAMI | 500 | 500 |

Table 3: Statistics of test sets.

A Datasets

The detailed statistics for the original test splits of the three datasets are shown in Table 3.

B Baselines

We compare LOREHM with several state-of-the-art (SoTA) systems for low-resource harmful meme detection: 1) PromptHate (Cao et al., 2022): a prompt learning approach that concatenates the meme text and the image caption as the prompt for masked language modeling; 2) MR.HARM (Lin et al., 2023a): a two-stage framework that distills multimodal reasoning knowledge from LLMs for harmfulness inference; 3) Pro-Cap (Cao et al., 2023): a caption-enhanced version of PromptHate, by leveraging pre-trained vision-language models with probing queries, to improve the image caption in the text prompt; 4) OPT-30B (Zhang et al., 2022): an early and representative large language model with the in-context learning ability, widely recognized stand-in for GPT-3 (Brown et al., 2020); 5) OpenFlamingo-9B (Awadalla et al., 2023): an open-source replication of Flamingo models (Alayrac et al., 2022) that enhances pre-trained, frozen language models by enabling them to cross-attend to the outputs of a frozen vision encoder during the next token prediction; 6) ModHATE (Cao et al., 2024): a modularized networks for low-resource harmful meme detection, which train a set of modules capable of relevant tasks and learn a composition of modules with the few-shot examples; 7) LLaVA-34B (Liu et al., 2023a): an enhanced version of LLaVA, with improved reasoning, OCR, and world knowledge capabilities; 8) GPT-4o (OpenAI, 2023): an optimized version of the proprietary GPT-4 architecture developed by OpenAI, which includes capabilities for processing multiple modalities; 9) LOREHM (*): Our proposed agentic approach for low-resource harmful meme detection, based on LLaVA-34B and GPT-4o. We use the accuracy and macro-averaged F1 (dominant) scores as the evaluation metrics, where

the macro-averaged F1 score is utilized as the monitor to select the best model since it could capture competitive performance beyond the majority class.

C Implementation Details

To demonstrate the generalizability of our framework, we utilize the representative LMMs LLaVA-34B and GPT-4o as the LMM agent from both the open-source and closed-source perspectives. Specifically, we implement the “llava-v1.6-34b” and “gpt-4o-2024-05-13” versions for LLaVA-34B and GPT-4o, respectively. Note that the choice of LMMs is orthogonal to our proposed paradigm, which can be easily replaced by newly stronger LMMs without further modification. To make our results reproducible, we set the temperature as 0 without any sampling mechanism. The frozen pre-trained vision and text Transformer encoders are implemented as CLIP (Radford et al., 2021) with the specific version “ViT-L/14@336px”. To choose the trade-off parameters α and β , we conducted a grid search within the range [0,1] and set α and β as 0.2 and 0.8, respectively. For the experiments, all the baselines are tested in the 50-shot setting with balanced classes. In the context of few-shot learning, evaluations can exhibit high variability due to the selection of meme sample examples. To mitigate this, we generate multiple few-shot reference sets using different random seeds for a more reliable few-shot performance evaluation. We create five sets of few-shot examples, each with a different random seed, for each 50-shot setting. Consequently, we report the average accuracy and macro-averaged F1 scores computed over the test set, following model optimization based on these various few-shot samples. As our approach is gradient-free, there are no training parameters.

We use the released source codes to reproduce PromptHate, MR.HARM, Pro-Cap, and ModHATE. To support the in-context learning of the LLM and/or LMM baselines, we need to first convert the meme’s image into an acceptable textual input, because the current LLMs and/or LMMs do not support 50 images as input to have a good performance. It is also a limitation of existing LMMs. We first in-paint the memes by combining MMOCR (Kuang et al., 2021) with SAM (Kirillov et al., 2023) to extract the text and pure image in memes. Then we apply LLaVA to generate textual descriptions about the dominant objects or events

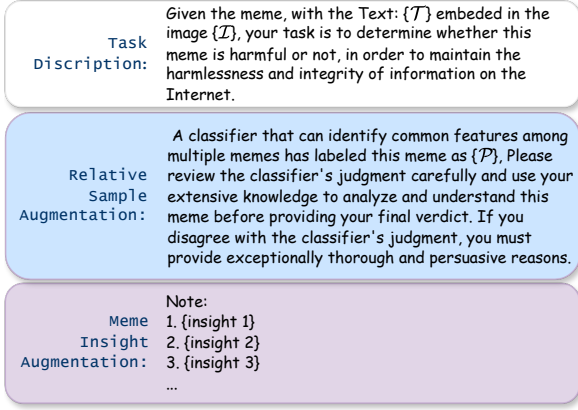


Figure 5: Prompt template for final judgment of LOREHM.

in the memes’ image. To ensure the reproducibility of the LLM and/or LMM baselines, we also set the temperature as 0 without any sampling mechanism.

In our approach, to better utilize the preliminary prediction \mathcal{P} provided by Relative Sample Augmentation, we design the prompt template as:

“A classifier that can identify common features among multiple memes has labeled this meme as $\{P\}$, Please review the classifier’s judgment carefully and use your extensive knowledge to analyze and understand this meme before providing your final verdict. If you disagree with the classifier’s judgment, you must provide exceptionally thorough and persuasive reasons.”

For the insight set \mathcal{E}_n provided by Meme Insight Augmentation, we sequentially concatenate each insight. Subsequently, we concatenate these two parts with the X_{COT} , resulting in the final prompt input for the LMM agent, as illustrated in Figure 5.

For gaining general insights into low-resource harmful meme detection, X_{Reflect} is designed as Figure 6. To restrict the number of insights, we set the capacity of the insight set to 10 by default. Once the insight set is full, the LLM agent is prohibited from producing the *ADD* action.

All experiments were conducted using OpenAI API and four NVIDIA A40 48GiB GPUs. Evaluation of HarM’s test set using LOREHM based on GPT-4o takes approximately 40 minutes, while FHM requires 1 hour and MAMI requires 2 hours. When employing LOREHM based on LLaVA-34B, evaluation times extend to 2 hours for HarM, 3 hours for FHM, and 6 hours for MAMI, approximately.

Due to privacy and ongoing research considerations, the code used in this study is not included

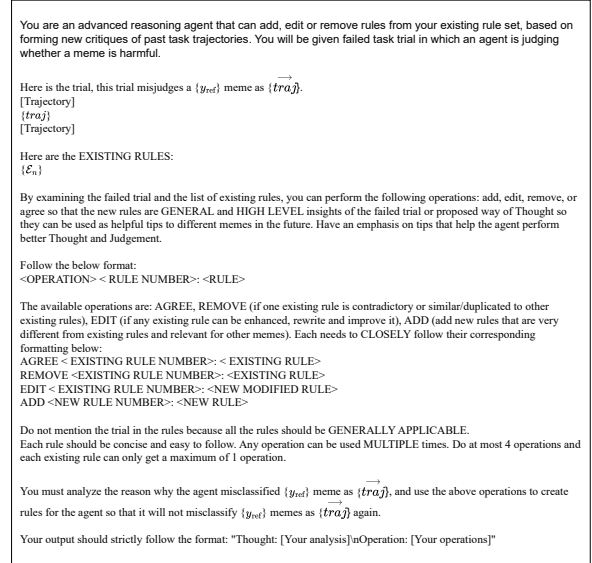


Figure 6: X_{Reflect} .

in the submission. However, we commit to making the code publicly available upon the acceptance of this paper.

D LOREHM Algorithm

We provide the algorithms for the gradient-free agentic paradigm of our approach, as depicted in Algorithms 1-4.

Algorithm 1 LOREHM - Initialization

Initialize:

Ratio factors for visual and textual embeddings α, β ;

Visual Encoder $\text{VE}(\cdot)$;

Textual Encoder $\text{TE}(\cdot)$;

Embedding $Emb \leftarrow \emptyset$;

The visual information \mathcal{I} of a meme M ;

The textual information \mathcal{T} of a meme M .

$Emb_{\mathcal{I}} \leftarrow \text{VE}(\mathcal{I})$

$Emb_{\mathcal{T}} \leftarrow \text{TE}(\mathcal{T})$

$Emb \leftarrow \alpha \cdot Emb_{\mathcal{I}} + \beta \cdot Emb_{\mathcal{T}}$

return Emb

E Discussion about LMMs

In this section, we discuss potential concerns on LMMs in the following three aspects: 1) Reproducibility: Since our proposed agency-driven approach uses not only the closed-source GPT-4o but also the open-source LLaVA-34B, the results are definitely reproducible with open-source codes. In order to make sure the LMM could generate the

Algorithm 2 LOREHM - Relative Sample Augmentation

Initialize:

Feature Representation from Algorithm 1:
reference meme's embedding information Emb_{ref} ,
target meme's embedding information Emb_{test} ;
Set of reference memes S_{ref} ;
Preliminary prediction \mathcal{P} ;
Set of Top K similar memes $\mathcal{H} \leftarrow \emptyset$;
Harmful counter $harmful_count \leftarrow 0$;
Similarity scores of two memes d ;
Set of Similarity scores $D \leftarrow \emptyset$;
Function capable of selecting the top K samples in the target set $Top_K(\cdot)$.
for each $M_{ref} \in S_{ref}$ **do**
 $d \leftarrow \text{cosine}(Emb_{ref}, Emb_{test})$
 $D \leftarrow D \cup d$
end for
for each $M_{ref} \in S_{ref}$ **do**
 if $d_{ref_test} \in Top_K(D)$ **then**
 $\mathcal{H} \leftarrow \mathcal{H} \cup M_{ref}$
 end if
end for
for each meme $H_i \in \mathcal{H}$ **do**
 if H_i is harmful **then**
 $harmful_count \leftarrow harmful_count + 1$
 end if
end for
if $harmful_count > \frac{K}{2}$ **then**
 $\mathcal{P} \leftarrow \text{harmful}$
else
 $\mathcal{P} \leftarrow \text{harmless}$
end if
return \mathcal{P}

Algorithm 3 LOREHM - Meme Insight Augmentation

Initialize:

Reference set S_{ref} ;
Empty insight set $\mathcal{E}_0 \leftarrow \emptyset$;
Self-reflection set $R_{set} \leftarrow \emptyset$;
Trajectory $traj \leftarrow \emptyset$;
Large Multimodal Model LMM;
Chain-of-thought prompt X_{CoT} ;
Reflection prompt $X_{Reflect}$;
Reference meme's ground truth label y_{ref} ;
Operations produced in insight extraction O .
Experience Gathering:
for each meme $M_{ref} \in S_{ref}$ **do**
 $traj \leftarrow \text{LMM}(X_{CoT}, \mathcal{I}_{ref}, \mathcal{T}_{ref})$
 if $traj \neq y_{ref}$ **then**
 $R_{set} \leftarrow R_{set} \cup \{traj\}$
 end if
end for
Insight Extraction:
for each $traj_i \in R_{set}$ **do**
 $O_i \leftarrow \text{LMM}(X_{Reflect}, traj_i, \mathcal{E}_{i-1})$
 $\mathcal{E}_i \leftarrow O_i(\mathcal{E}_{i-1})$
end for
return \mathcal{E}_n

Algorithm 4 LOREHM - Inference

Initialize:

Preliminary prediction \mathcal{P} from Algorithm 2;
Insight set \mathcal{E}_n from Algorithm 3;
The visual information \mathcal{I}_{test} of a meme M_{test} ;
The text information \mathcal{T}_{test} of a meme M_{test} ;
Large Multimodal Model LMM;
Chain-of-thought prompt X_{CoT} ;
Final prediction p_{final} .
 $p_{final} \leftarrow \text{LMM}(X_{CoT}, \mathcal{I}_{test}, \mathcal{T}_{test}, \mathcal{P}, \mathcal{E}_n)$
return p_{final}

same contents for the same instance, we utilize the specific version “lava-v1.6-34b” of LLaVA-34B and “gpt-4o-2024-05-1” of GPT-4o, and further set the parameter temperature as 0 without any sampling mechanism, that is, the greedy decoding was adopted to ensure the deterministic results for the content generation with the same prompt. 2) Test Set Leakage: The test set leakage issue does not exist in the open-source LLM (Liu et al., 2023a) as the paper has clearly described the instruction-tuning data used for training, which does not include any data used in our experiments. However, we cannot fully guarantee the exclusion of potential data leakage with GPT-4o, as its internal workings remain opaque. Nevertheless, as evidenced by the results in Table 2, where the LLaVA-34B or GPT-4o was directly deployed to test on the three standard datasets, the absence of significant test set leakage is implied. This is inferred from the fact that direct zero-shot application of the LMMs did not yield disproportionately high performance, which would be expected if the models were benefiting from test set leakage. Moreover, we can consistently observe enhanced performance in the variants of our proposed framework based on both LLaVA-34B and GPT-4o. This suggests that such improvement is basically attributed to our designed agentic paradigm rather than test set leakage. 3) Generalizability: We believe our LOREHM paradigm is a general technique that works with emerging stronger LMMs, because our approach works not only on GPT-4o, but also well on the open-source LLaVA-34B, which is not an OpenAI system.

F Discussion about RSA

For the proposal of the Relative Sample Augmentation mechanism, we chose not to use the retrieved memes as direct input to the LMM agent. This decision was made to explicitly augment the LMM agent with the label information of the retrieved memes. Simply inputting the retrieved memes back into the LMM agent would merely re-utilize its internal knowledge without incorporating external label information. Additionally, using retrieved memes along with labels as explicit signals is not more effective and somewhat overlaps with the concept of Meme Insight Augmentation. Instead, using a voting mechanism to integrate the label information of retrieved memes as auxiliary signals for outward analysis complements the agent’s knowledge-revising strategy through inward anal-

ysis. The macro-averaged F1 scores of the voting mechanism are approximately 79%, 59%, and 70% on the HarM, FHM, and MAMI test sets, respectively. Although the voting mechanism itself is not entirely robust, it can still provide effective label information for the LMM agent as a prior preference. From both theoretical and practical perspectives, our current design is the most reasonable and effective approach.

G Discussion about MIA

In our design of Meme Insight Augmentation, the knowledge-revising behavior of human beings is non-trivial when applied to harmful meme detection. Unlike many previous agent tasks that operate within environments providing real feedback, harmful meme detection lacks such an environment to supply the agent with authentic responses. Additionally, because harmful meme detection is fundamentally a binary classification task, informing the agent of an unsuccessful attempt essentially reveals the correct answer. These factors make the direct application of methods like Reflexion (Shinn et al., 2024), which relies on environmental feedback to determine the success of an attempt and decide whether to reflect and retry based on insights gained from previous attempts, inapplicable. Similarly, ExpeL (Zhao et al., 2024), a method requiring multiple passes on the same data to obtain both successful and failed trajectories, cannot be directly utilized. Additionally, during our detailed examination of data samples, we discovered a few instances of contentious annotations. More specifically, in the rare instances of incorrect or disputable annotations, even though we engage in retrieved memes with the incorrect golden label as part of the voting mechanism in the Relative Sample Augmentation phase, the well-generalized insights extracted from the knowledge-revising process also contributes to a better understanding of our approach’s robustness, particularly when dealing with inaccurately annotated data. Therefore, our proposed MIA component is specially devised to couple with the RSA component, which is complementary to each other.

H More Examples of Case Study

H.1 Case Study of Different Labeled Data Sizes

We provide a case study to compare the different output thoughts from the 0-shot GPT-4o, 50-shot GPT-4o and LOREHM, as shown in Figure 7.



Figure 7: Examples of correctly predicted harmful memes in (a) HarM and (b) MAMI datasets.

From the output thought in natural text, we observe that: 1) Our proposed LOREHM could offer a more profound analysis compared to both the 0-shot and 50-shot GPT-4o methods. For example, in Figure 7(a), the 0-shot GPT-4o might deem the meme harmless, failing to grasp its satirical nature, while the 50-shot GPT-4o, though more critical, might not fully consider the meme’s broader implications. Our proposed LOREHM identifies the meme’s potential to trivialize the pandemic’s severity with the phrase “Going golfing & ignoring pandemic” leading us to classify the meme as harmful, as it could undermine the importance of public health measures during a critical time. What’s more, in Figure 7(b), the 0-shot GPT-4o might overlook the meme’s reinforcement of traditional gender roles, while the 50-shot GPT-4o, influenced by limited examples, might not fully appreciate the cultural implications. LOREHM, however, scrutinizes the meme’s depiction of a woman in a 1950s-style housewife outfit with the text “THE LAST TIME I MARCHED IT WAS DOWN THE AISLE” concluding that it subtly perpetuates outdated gender norms, which can be harmful in the context of modern feminist movements. 2) Our method exhibits notable efficiency within the con-



Figure 8: Examples of correctly predicted harmful memes in (a) FHM and (b) HarM datasets.

straints of low-resource scenarios (Lin et al., 2022). Compared to the 50-shot GPT-4o, which relies on a limited set of examples to improve its judgments, our method shows superior performance with the same resource constraints. By effectively leveraging few-shot learning and relevant sample retrieval, our method can discern harmful content with greater accuracy. This capability is crucial when the availability of training examples is scarce, as our method requires fewer examples to achieve a higher level of performance, thus outperforming the 50-shot GPT-4o in equal-resource conditions.

H.2 Case Study of RSA & MIA

We provide a case study of the LMM agent’s output thoughts, to investigate the effect of the RSA and MIA strategies on the correctly predicted harmful meme samples by LOREHM, as illustrated in Figure 8.

From the output thought in natural text, we observe that: 1) The Relative Sample Augmentation (RSA) mechanism enhances the LMM agent by incorporating label information from retrieved memes, providing the LMM agent with additional context and insights. For example, in Figure 8(a), the GPT-4o output considered it harmless, interpreting it as a joke. However, with RSA, the LMM



Figure 9: Examples of wrongly predicted memes by our proposed framework with the ground truth (a) harmless and (b) harmful.

agent identified it as perpetuating harmful stereotypes, recognizing the potential to offend or misinform viewers. This shows how RSA integrates external labels as auxiliary signals, complementing the agent’s internal knowledge-revising strategy and creating a balanced mechanism that leverages both internal and external information sources. By focusing on the labels of retrieved memes rather than directly inputting the memes themselves, RSA ensures that the augmentation process adds unique value to the LMM agent’s capabilities. 2) The Meme Insight Augmentation (MIA) mechanism offers significant benefits by enabling the LMM agent to revise its knowledge effectively. For instance, in Figure 8(b), the GPT-4o output thought it was a humorous critique and not harmful. However, MIA identified the risk of trivializing a serious public health issue, noting the potential for misinformation and its impact on public health measures. In emerging events (Lin et al., 2024d) on social media (Lin et al., 2021), MIA could promote a knowledge-revising behavior akin to human reasoning, allowing the LMM agent to better interpret memes and distinguish between harmful and harmless content.

I Error Analysis

To better understand the behavior of our framework and facilitate future studies, we conduct an error analysis on the wrongly predicted memes by our proposed framework. Figure 9 shows two examples of memes wrongly classified by LOREHM. For the harmless meme in Figure 9(a), featuring a

historical figure with the text “so many lives lost because of this man”, LOREHM incorrectly categorized it as harmful. The original thought provided by LOREHM was that the meme could be seen as trivializing a serious and tragic period in history. The misjudgment in this case likely arose from LOREHM’s failure to effectively associate the image with the accompanying text, leading to an overemphasis on the historical figure’s negative connotations without considering the full context provided by the meme’s message. In contrast, the harmful meme in Figure 9(b), which depicted several clergymen with the text “money is evil give it to us” was deemed harmless by LOREHM. This misjudgment indicates an inherent bias in LOREHM towards religious-themed memes. Contrary to being overly sensitive, LOREHM’s bias led to an underestimation of the potential negative impact of such content. Additionally, we also investigated more errors resulting from the impact of visual artifacts like image quality, occlusion, obscurity, etc. We found that low-quality images lead to the wrong recognition of superficial patterns and ignorance of the occlusion.

J Related Work about LLMs and LMMs

Recently, LLMs have demonstrated exceptional versatility across various tasks. Significant advancements by leading tech companies have resulted in highly proficient, though often proprietary, LLMs (Brown et al., 2020; OpenAI, 2023; Chowdhery et al., 2022; Team et al., 2023). Meanwhile, the NLP community has seen the rise of open-source LLMs, with publicly shared model weights (Black et al., 2022; Zeng et al., 2022; Touvron et al., 2023a,b; Luo et al., 2023). More recently, LLMs have also been developed to adapt to processing both textual and visual data, marking a significant advancement. Recent research has focused on constructing versatile multimodal datasets (Yang et al., 2023) from platforms like GPT-4 and GPT-4V (OpenAI, 2023), fine-tuning open-source LMMs, *e.g.*, LLaVA (Liu et al., 2023a) and other innovative projects (Bai et al., 2023; Dai et al., 2023; Wang et al., 2023b). GPT-4o, the latest version of the GPT family, has been released lately. These LMMs have shown excellent emergent abilities in multimodal tasks. In this work, we utilize the representative LMMs, *i.e.*, LLaVA-34B and GPT-4o, as the backbone variants of the LMM agent, to explore the low-resource issue (Lin et al.,

2023b) in harmful meme detection with limited labeled meme examples.

K Examples of Insights

In this section, we exemplify the insights drawn from the MIA strategy with 50-shot meme examples.

HarM by LOREHM(GPT-4o)

- Evaluate if the humor in the meme is appropriate and does not trivialize serious issues, create unnecessary panic, target individuals or groups in a derogatory manner, or undermine critical health measures and policies, especially during ongoing crises like pandemics.
- Consider the broader context and potential negative impact of the meme, especially in sensitive situations such as health crises or political discourse, to determine if it could cause fear, anxiety, spread misinformation, trivialize serious issues, or undermine important public health measures.
- Identify if the meme mocks, trivializes, or undermines public health measures and policies, as this can contribute to harmful behaviors, spread misinformation, and undermine public trust in health directives and governance.
- Ensure the meme does not encourage or normalize behavior that could be harmful in real life, particularly in the context of public health and safety, such as trivializing quarantine, policy effectiveness, or other critical health measures.
- Assess if the meme perpetuates misinformation or undermines public health measures, which can contribute to harmful outcomes.
- Ensure the meme does not inadvertently normalize, trivialize, or encourage harmful behaviors by presenting them humorously, especially in the context of ongoing public health issues and governance policies.
- Recognize that memes leveraging fear-based tactics, even for promoting good behavior, can be harmful and should be scrutinized closely.
- Assess the potential for the meme to incite derogatory or harmful sentiments towards individuals or groups, which can contribute to

a negative, hostile, or unproductive environment.

- Assess if the meme trivializes or mocks behaviors that were critical during crises (*e.g.*, panic buying, hoarding) in a way that could diminish the perceived seriousness of the situation or encourage similar behavior in the future.
- Determine if the meme's absurdity or exaggeration clearly indicates humor, reducing the likelihood of it being taken seriously and causing harm.

FHM by LOREHM(GPT-4o)

- Evaluate the context and possible interpretations of both the image and text together to determine if the meme could be harmful.
- Use a balanced approach, considering the potential for harm while recognizing the context, intent, and humor of the meme. Clearly differentiate between harmful content and satire or light-hearted humor to avoid unnecessary censorship.
- Assess the potential for a meme to be perceived as offensive by considering the broader context and audience, but avoid overemphasizing potential offense if the context and humor suggest a harmless or light-hearted intent.
- Consider the use of exaggeration and satire in memes, recognizing that these elements are often used for humor. Clearly differentiate between satirical humor and content with genuine harmful intent.
- Always consider the potential for text in a meme to be interpreted in a harmful or inappropriate manner. However, distinguish between light-hearted humor and genuinely harmful or disrespectful content.
- Consider the broader social, cultural, and racial implications of the meme's content to assess its potential harm.
- Differentiate between controversial or politically charged content and genuinely harmful content to prevent unnecessary censorship of opinion-based memes.

- When evaluating memes, always consider if the content perpetuates racial, cultural, or ethnic stereotypes, as this can be harmful even if the intent seems humorous.
- Explicitly differentiate between anthropomorphism and harmful stereotypes about people or specific groups, recognizing that portraying animals with human traits for humor is generally acceptable if it does not convey a negative or harmful message.
- When evaluating memes, consider the historical and cultural significance of the references used, ensuring that sensitive topics are not trivialized or used inappropriately.
- Evaluate whether the meme contains humor that could desensitize viewers to serious issues or dangerous situations, and consider the context and intent to determine if it amplifies potential harm.
- Evaluate if the humor or satire presents a conflicting message that diminishes the original intent of empowering or positive content.
- Distinguish between explicit content meant to entertain a mature audience and content genuinely intended to harm or offend, ensuring cultural and contextual humor is appropriately considered.

MAMI by LOREHM(GPT-4o)

- Recognize and account for humor or satire, ensuring they are not mistaken for serious harmful content. Consider the context, intent, and whether the humor directly trivializes serious subjects such as violence, sexual behavior, or perpetuates harmful stereotypes.
- When evaluating humor, assess whether it perpetuates or trivializes harmful stereotypes, biases, serious life-threatening situations, or explicit content inappropriately. Consider the overall impact on the targeted group, while balancing this with the intent and context of the humor.
- When identifying harmful content, evaluate the context, intent behind the language used, and the cultural nuances of humor, especially when dealing with sensitive topics.
- Consider the potential for reinforcing harmful stereotypes, real-world harm, or incitement that the content may cause, but also recognize the difference between satire and genuinely harmful content.
- Avoid overgeneralizing potentially divisive content as inherently harmful without further examination.
- Consider if the meme's humor is based on sensitive or protected characteristics, and the potential negative impact it could have on those groups, especially if it reinforces harmful stereotypes or biases.

- Consider whether the combination of humor and serious messages dilutes the impact of the positive message or reinforces harmful attitudes.

HarM by LOREHM(LLaVA-34B)

- Assess the overall impact of the meme, including whether it is likely to be perceived as light-hearted or offensive, especially in sensitive contexts such as health crises, disasters, social issues, or activism. Consider whether it trivializes serious situations, promotes irresponsible or dangerous behavior, spreads misinformation, or is insensitive to those affected, even if the intent is humorous.
- Consider the context and intent of the meme, especially if it is meant to be humorous or relatable, but also weigh the potential for it to be perceived as insensitive, harmful, or misleading in serious situations.
- Evaluate whether the meme directly or indirectly promotes harmful stereotypes, ideas, or personal attacks, trivializes serious issues or individuals' efforts, spreads misinformation, or could be perceived as insensitive or harmful.
- Evaluate whether the meme trivializes serious situations, issues, or the efforts of activists, which could be harmful or offensive to those affected, and consider the potential for it to spread misinformation, encourage irresponsible behavior, or be perceived as insensitive, especially in the context of public health and safety.

- Evaluate whether the meme could potentially encourage harmful actions or behaviors, especially in the context of public health and safety. Consider the potential for it to downplay the severity of serious issues, spread misinformation, or be perceived as insensitive, even if the intent is humorous.
- Consider the potential for the meme to spread misinformation, harmful stereotypes, or contribute to harmful narratives, even if it is intended as a joke, especially in the context of public health and safety.
- Pay special attention to memes that reference sensitive social, political, or cultural issues, as these are more likely to be harmful.
- Assess the potential for the meme to be interpreted in a harmful or misleading way by different audiences, considering the broader social, political, and cultural context, and the potential for it to spread misinformation or harmful narratives.
- Evaluate the potential for the meme to be misinterpreted by different audiences, leading to harmful consequences, even if the intent is humorous, and consider the broader social and cultural context.
- Consider the potential for the meme to be interpreted as promoting or endorsing harmful actions or behaviors, especially in the context of public health and safety.

FHM by LOREHM(LLaVA-34B)

- Always consider the potential for double meanings, innuendos, and cultural stereotypes in the text of a meme.
- Evaluate the context in which the meme could be shared, its intent, and its potential impact on different audiences.
- Be cautious of memes that could be interpreted as promoting or trivializing inappropriate or harmful behavior.
- Consider the broader social and cultural implications of the meme's content and text, while also weighing the intent, humor, and factual accuracy behind it.

- Be sensitive to cultural stereotypes and avoid memes that could perpetuate or reinforce negative stereotypes about any group.
- Differentiate between provocative or controversial statements and genuinely harmful content, ensuring that the meme's potential to cause real harm is the primary consideration.

MAMI by LOREHM(LLaVA-34B)

- Always consider if the meme perpetuates stereotypes or objectifies individuals based on gender, race, or other attributes.
- Assess if the humor or content used in the meme could be interpreted as demeaning, derogatory, or objectifying towards any group.
- Take into account the broader social context, intent, and potential negative interpretations of the meme, distinguishing between harmful content and harmless humor.
- Evaluate whether the meme could be perceived as trivializing or making light of serious social issues, even if the intent appears humorous.
- Consider the use of satire, exaggeration, and wordplay in the meme and whether it is intended to provoke thought or humor rather than to harm.

L Examples of Retrieved Memes

Figures 10-12 demonstrate the top 5 relative memes retrieved from the RSA strategy with 50-shot meme examples on HarM, FHM, and MAMI, respectively.



Figure 10: Retrieved top 5 relative memes to the target meme sampled from HarM.

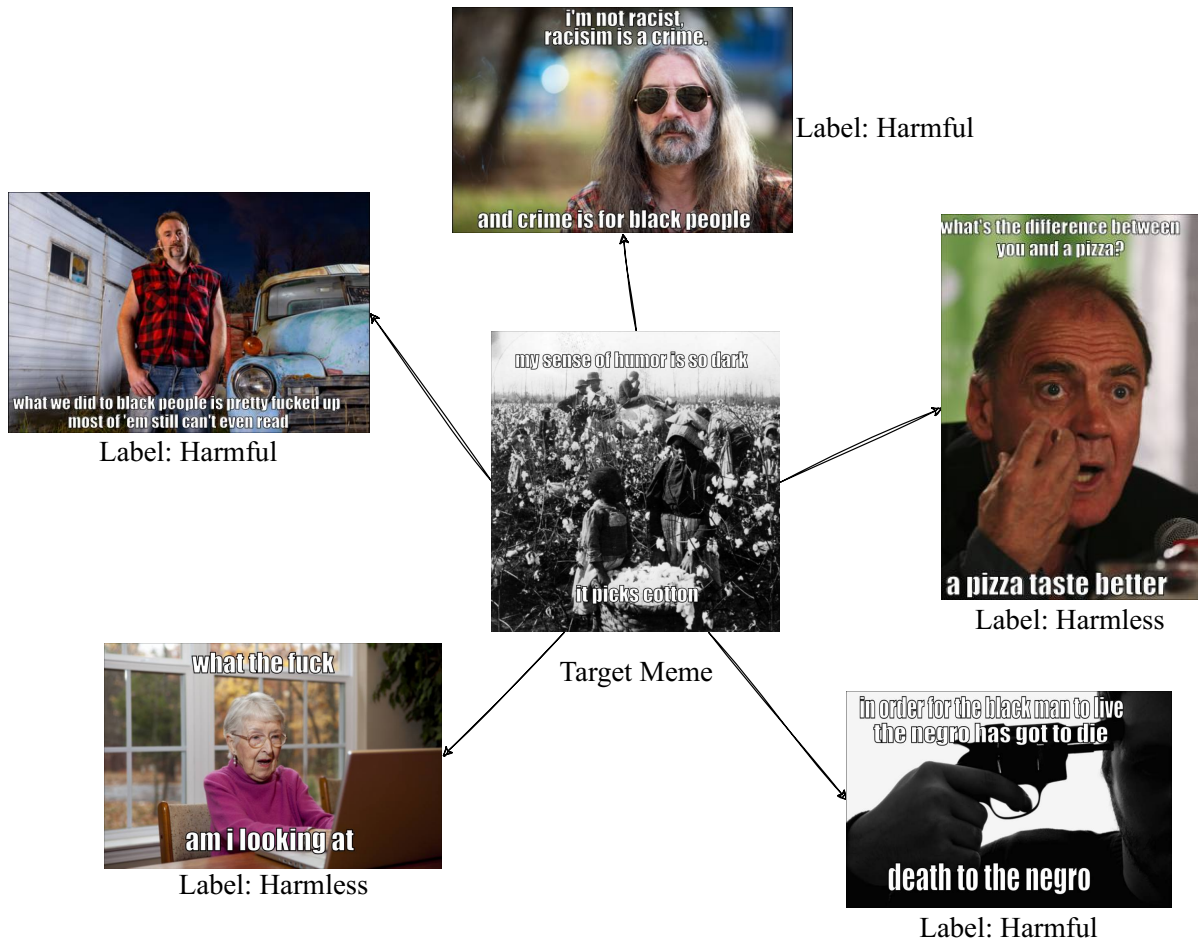


Figure 11: Retrieved top 5 relative memes to the target meme sampled from FHM.

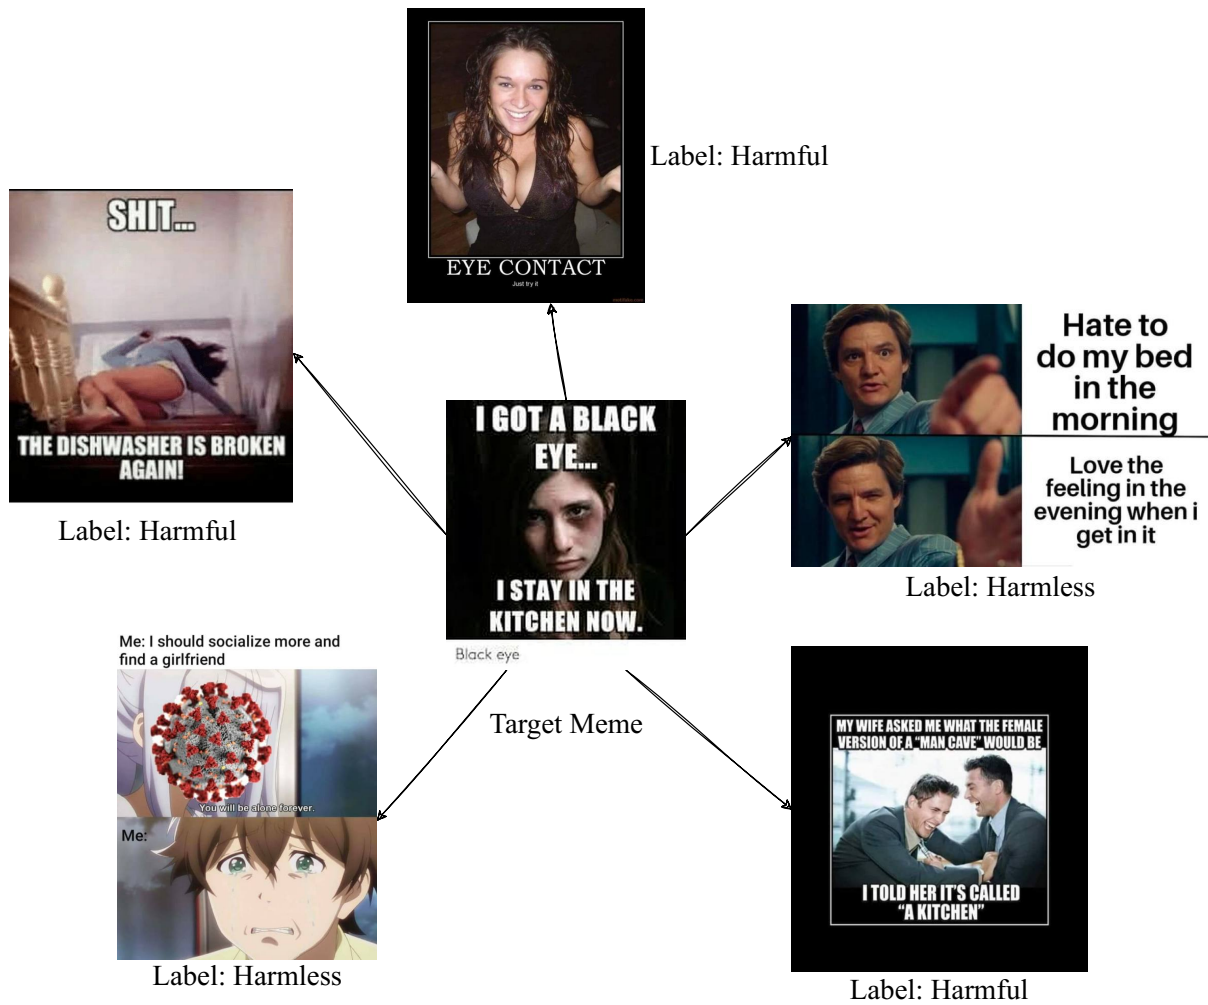


Figure 12: Retrieved top 5 relative memes to the target meme sampled from MAMI.