

Enhancing Meme Emotion Understanding with Multi-Level Modality Enhancement and Dual-Stage Modal Fusion

Yi Shi¹, Wenlong Meng¹, Zhenyuan Guo¹, Chengkun Wei^{1*}, Wenzhi Chen¹

¹Zhejiang University

{shiyi666, mengwl, zhenyuanguo, weichengkun, chenwz}@zju.edu.cn

Abstract

With the rapid rise of social media and Internet culture, memes have become a popular medium for expressing emotional tendencies. This has sparked growing interest in Meme Emotion Understanding (MEU), which aims to classify the emotional intent behind memes by leveraging their multimodal contents. While existing efforts have achieved promising results, two major challenges remain: (1) a lack of fine-grained multimodal fusion strategies, and (2) insufficient mining of memes' implicit meanings and background knowledge. To address these challenges, we propose *MemoDetector*, a novel framework for advancing MEU. First, we introduce a four-step textual enhancement module that utilizes the rich knowledge and reasoning capabilities of Multimodal Large Language Models (MLLMs) to progressively infer and extract implicit and contextual insights from memes. These enhanced texts significantly enrich the original meme contents and provide valuable guidance for downstream classification. Next, we design a dual-stage modal fusion strategy: the first stage performs shallow fusion on raw meme image and text, while the second stage deeply integrates the enhanced visual and textual features. This hierarchical fusion enables the model to better capture nuanced cross-modal emotional cues. Experiments on two datasets, MET-MEME and MOOD, demonstrate that our method consistently outperforms state-of-the-art baselines. Specifically, *MemoDetector* improves F1 scores by 4.3% on MET-MEME and 3.4% on MOOD. Further ablation studies and in-depth analyses validate the effectiveness and robustness of our approach, highlighting its strong potential for advancing MEU.

Code — <https://github.com/singing-cat/MemoDetector>

1 Introduction

In recent years, the rise of Internet culture has led to an increasing use of memes as a medium for expressing personal emotions and intentions (Shifman 2013; Vásquez and Aslan 2021). Memes, characterized by their inherently multimodal nature combining image and text, have become a central form of online communication. As a result, Meme Emotion Understanding (MEU) has emerged as an important and timely research topic. However, unlike conventional multimodal sentiment analysis, memes often rely on metaphorical



(a) When you're on the verge of mental collapse and a minor inconvenience happens. Today's the day!
 (b) You have $\int_a^b f(x) dx$ friends. Calculus can be funny if you understand it.

Figure 1: Examples illustrating the challenges faced by existing methods.

expressions and are highly context-dependent, which poses a crucial challenge for MEU.

Several studies have made promising strides in meme emotion understanding. Xu et al. (2022) constructed a bilingual dataset focused on metaphor-rich memes, where emotion detection was introduced as a subtask. They also proposed a single-stage multimodal fusion method tailored to this task. Similarly, Sharma et al. (2024) created a dataset based on Ekman's six basic emotions (Ekman and Cordaro 2011), specifically designed for meme emotion classification. While these works represent meaningful progress, they still suffer from two major limitations: (1) **a lack of fine-grained multimodal fusion strategies** that can fully capture the intricate interplay between visual and textual elements in memes and (2) **insufficient mining of implicit meanings and background knowledge**, which are crucial for understanding the nuanced emotions embedded in memes.

On one hand, most existing approaches perform only a shallow fusion of the visual and textual modalities in memes, overlooking the subtle and intricate relationships between image and text. Such simplistic fusion strategies often limit the model's ability to capture cross-modal nuances, leading to misinterpretations of the meme's intended emotional tone. As illustrated in Figure 1 (a), the meme combines a cheerful facial expression with sorrowful text, creating a contrast that effectively highlights the sender's underlying sadness. Simple fusion strategies fail to capture this nuanced interplay between modalities, making it difficult for the model to correctly interpret the emotional message. On the other hand, existing methods focus solely on the meme itself, lacking deeper exploration of its implicit meanings and relevant

*Corresponding author

background knowledge. As a result, crucial emotional cues embedded in cultural or contextual references may be overlooked. For example, as illustrated in Figure 1 (b), the sentiment of this meme is embedded in the mathematical result of an integral expression. Without sufficient mathematical knowledge, small models are unlikely to recognize the underlying sadness behind this meme.

To address these challenges, we introduce *MemoDetector*, a novel framework designed to advance MEU. First, we leverage the rich world knowledge and powerful reasoning capabilities of Multimodal Large Language Models (MLLMs) to perform a four-step textual enhancement of memes. These four steps progress from surface to depth, covering image, text, their joint semantics, and the meme’s potential real-world context. Through this multi-level analysis, our approach effectively uncovers the underlying meanings and background knowledge embedded in memes, significantly enriching the original information and providing strong guidance for downstream classification by a small model. Next, to address the limitations of prior approaches that rely on overly simplistic fusion strategies, we propose a dual-stage modal fusion mechanism. This strategy not only performs a shallow fusion of surface-level features but also enables a deeper, bidirectional integration of the enhanced visual and textual representations. Specifically, we first conduct a shallow fusion between original meme image and its accompanying text to obtain an enriched visual representation. We then perform a deep bidirectional fusion between the enriched visual features and the enhanced textual features, allowing both modalities to mutually refine and reinforce each other. This dual-stage fusion empowers the model to develop a more nuanced understanding of inter-modal relationships and capture subtle emotional cues that may be overlooked by conventional single-stage fusion methods.

To validate the effectiveness of our method, we conduct extensive experiments on two benchmark datasets: MET-MEME (Xu et al. 2022) and MOOD (Sharma et al. 2024). The results demonstrate that our model significantly outperforms state-of-the-art (SoTA) baselines across all evaluation metrics on both datasets. Specifically, on MET-MEME, our approach achieves improvements of 4.17% in Accuracy and 4.3% in Macro-F1. On MOOD, the gains are 4.04% and 3.4%, respectively. In addition, we conduct ablation and in-depth analytical experiments to validate the unique contributions of each component in our framework and its superiority over existing methods. These findings further demonstrate the effectiveness and robustness of our approach.

Our main contributions can be summarized as follows:

- We analyze two core challenges in meme emotion understanding, and propose a novel framework, *MemoDetector*, which for the first time incorporates contextual background knowledge and introduces a fine-grained modal fusion strategy to enhance MEU.
- We design a four-step textual enhancement module and a dual-stage modal fusion mechanism. These components hierarchically uncover meme background knowledge and implicit meanings, while enabling a more nuanced fusion of visual and textual features.

- Extensive experiments demonstrate the superiority of our method, achieving state-of-the-art performance across all evaluation metrics in meme emotion understanding.

2 Related Work

Multimodal Emotion Detection

Previous studies have primarily focused on the broader task of multimodal emotion analysis (Hu et al. 2022; Ahmed, Al Aghbari, and Girija 2023; Cheng et al. 2024), rather than specifically targeting memes. Unlike unimodal emotion analysis, multimodal emotion detection requires effective fusion of features from different modalities (Zhu et al. 2023; Gan et al. 2024). Traditionally, multimodal fusion strategies are categorized into three types: data-level fusion (Danapal et al. 2020), early fusion (Pranesh and Shekhar 2020), and late fusion (Pramanick et al. 2021). Data-level fusion merges raw inputs at the input stage, while early fusion integrates modality-specific features. In contrast, late fusion combines separate predictions via decision-level strategies. Although these fusion methods have shown effectiveness in general multimodal emotion analysis, they often fall short when applied to memes, a unique medium that frequently involves metaphorical and context-dependent cues (Younes and Al-takhaineh 2022; Xu et al. 2024b). These simplistic fusion methods lack the capacity to capture the intricate and implicit interplay between meme’s modalities.

Meme Analysis

With the rise of memes, research on meme analysis has increasingly emerged. Existing studies mainly focus on hateful memes. Cao et al. (2023) used a vision-language model to generate enriched meme captions, which were then fed into a language model for hate classification. Lin et al. (2024) proposed an innovative approach by leveraging LLMs to conduct debates. Beyond hate detection, finer-grained tasks like meme emotion detection have also been explored. Xu et al. (2022) introduced a meme understanding dataset that includes emotion analysis as a sub-task. Sharma et al. (2024) proposed MOOD dataset, specifically designed for meme emotion classification. Methods for this task include both unimodal and multimodal approaches. For unimodal feature extraction, models such as BERT (Devlin et al. 2019) and GloVe (Pennington, Socher, and Manning 2014) were used for text, while EfficientNet (Tan and Le 2019) and ResNet (He et al. 2016) served as image encoders. Some studies, such as Singh, Bauwelinck, and Lefever (2020) and Vlad et al. (2020), applied multi-task learning to jointly predict emotion and sentiment. Despite encouraging results, existing methods struggle to capture memes’ nuanced semantics and contextual variability, as small models lack the capacity for deep understanding.

3 Methodology

Task Definition

This study addresses the task of meme emotion understanding (MEU), formulated as a multi-class classification problem. Specifically, given a meme instance $M = (I, T)$, where

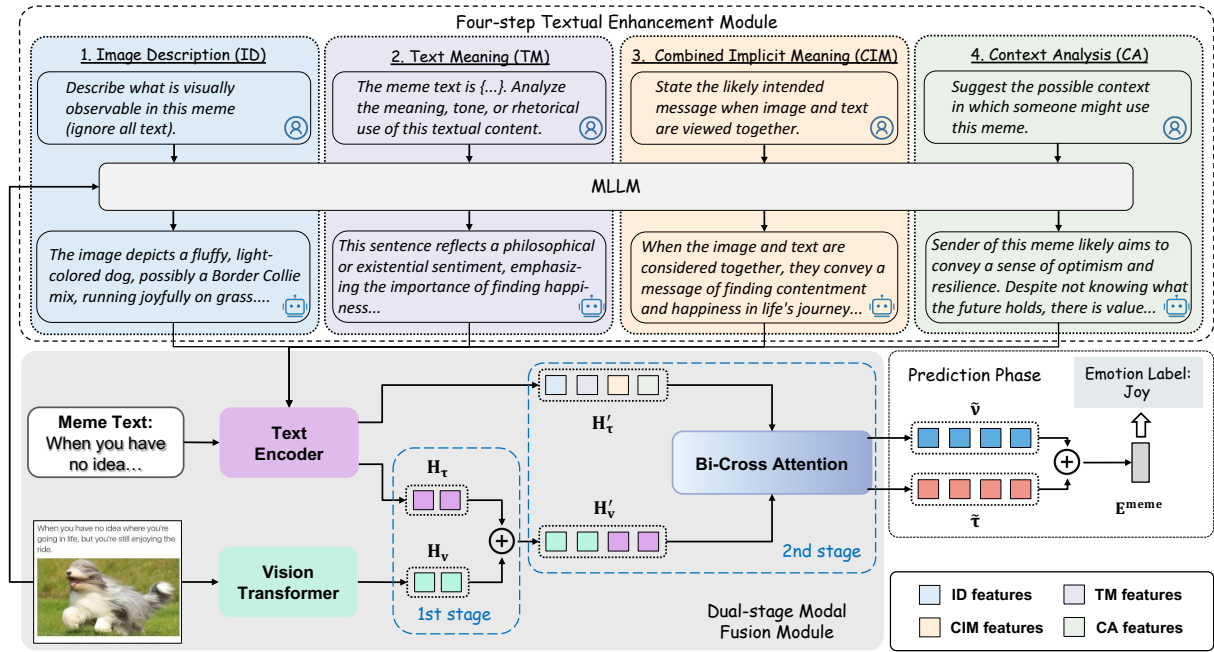


Figure 2: Overview of our framework. We first prompt MLLM to generate multi-level insights for the meme, which are distilled into a small model as auxiliary guidance. The small model then performs dual-stage modal fusion to produce the final prediction.

I denotes the image component and T represents the associated textual content, the objective is to predict the corresponding emotion label y_{emo} that reflects the affective intent likely conveyed by the meme’s sender.

Our core idea is to harness the rich background knowledge and powerful reasoning capabilities of MLLMs (Liang et al. 2024; Wang et al. 2024) to perform a multi-level analysis of memes. This includes shallow perception, deep interpretation, and associative reasoning. The resulting explanations are distilled as auxiliary knowledge to a lightweight classifier, enhancing its capacity for MEU. Subsequently, the classifier employs a carefully designed dual-stage modal fusion strategy to effectively integrate information from both modalities, leading to more accurate predictions. The overview of our framework is shown in Figure 2.

Four-step Text Modality Enhancement

Leveraging the rich knowledge and reasoning abilities of MLLMs, we generate multi-level interpretations of memes to support more accurate emotion prediction. Traditional small-model-based approaches often capture only surface-level features of memes, lacking the ability to understand deeper aspects such as metaphor and context. Our method addresses this by providing enriched semantic insights that enhance the model’s emotional understanding.

Inspired by the progressive nature of human cognition (Bloom et al. 1956), we divide MLLM’s meme understanding process into three hierarchical levels: shallow perception, deep interpretation, and associative reasoning. This tiered structure enables MLLM to gradually build a more nuanced understanding. To further enhance the quality of MLLM-generated textual explanations, we adopt a Chain-of-Thought (CoT) prompting strategy (Wei et al. 2022),

designing a four-step prompt sequence aligned with the three levels of understanding: **Image Description (ID)**, **Text Meaning (TM)**, **Combined Implicit Meaning (CIM)**, and **Context Analysis (CA)**. We provide a detailed explanation of each step as follows.

Step 1: Image Description. To provide MLLM with a holistic understanding of memes’ visual content, we begin by prompting it to describe key visual elements, such as objects, characters, and background scenes. While these visual cues represent surface-level features, they are essential for emotion recognition. For instance, facial expressions often convey strong emotional signals. To eliminate interference from meme text, we design the following prompt p^{ID} :

“Describe what is visually observable in this meme (ignore all text).”

Given meme image I and prompt p^{ID} , MLLM will generate image description T^{ID} as follows:

$$T^{ID} = \text{MLLM}(I, p^{ID}) \quad (1)$$

Step 2: Text Meaning. In addition to visual information, textual modality also plays a crucial role in meme emotion detection. Many memes convey specific emotional intentions through cleverly crafted language and metaphorical expressions. However, meme texts often employ rhetoric or emphatic tones that obscure their underlying emotional cues. To uncover these signals, we instruct MLLM to focus on the semantics, tone, and rhetorical usage of meme text by providing the following prompt p^{TM} :

“The meme text is {...}. Analyze the meaning, tone, or rhetorical use of this textual content.”

Given meme text T and prompt p^{TM} , MLLM will generate text meaning T^{TM} as follows:

$$T^{TM} = \text{MLLM}(T, p^{TM}) \quad (2)$$

Step 3: Combined Implicit Meaning. After separately analyzing each modality, we prompt the MLLM to interpret the combined metaphorical meaning that emerges from the interplay between image and text. This step is crucial, as many memes rely on subtle cross-modal associations to convey nuanced messages that cannot be captured by uni-modal analysis alone. Ignoring such interactions can result in a biased or incomplete understanding. To elicit this joint interpretation, we design the following prompt p^{CIM} :

“State the likely intended message when image and text are viewed together.”

Given meme image I , meme text T and prompt p^{CIM} , MLLM will generate combined implicit meaning T^{CIM} :

$$T^{CIM} = \text{MLLM}(I, T, p^{CIM}) \quad (3)$$

Step 4: Context Analysis. In the final step, after MLLM has formed a relatively comprehensive understanding of the meme, we simulate human cognitive processes by encouraging associative reasoning, prompting the model to infer possible usage contexts based on its prior analysis. Given the context-dependent nature of memes, their emotional connotations often become more salient when grounded in a plausible situational backdrop. To enable the model to capture this contextual nuance, we introduce prompt p^{CA} :

“Suggest the possible context in which someone might use this meme.”

Given meme image I , meme text T and prompt p^{CA} , MLLM will generate context analysis T^{CA} as follows:

$$T^{CA} = \text{MLLM}(I, T, p^{CA}) \quad (4)$$

These four steps constitute a core component of our approach, enabling a multi-level interpretation of memes. Step ID and TM belong to shallow perception, extracting surface information from visual and textual modalities. Step CIM targets deep interpretation, revealing implicit cross-modal meanings. Step CA represents associative reasoning, inferring plausible usage contexts. Together, these four steps significantly enrich the original meme content, empowering downstream classifier to make more accurate predictions.

Dual-stage Modal Fusion

Feature Extraction. After obtaining the enhanced textual modality, we first extract features from both image and text. Following the setup of Lin et al. (2024), we employ ViT (Dosovitskiy et al. 2020) to encode visual features. For text modality, to ensure our model is capable of handling memes in multiple languages, we adopt XLM-R (Conneau et al. 2019) as the text encoder. The overall feature extraction process can be formulated as follows:

$$\mathbf{H}_v = \text{ViT}(I) \quad (5)$$

$$\mathbf{H}_\tau = \text{XLM-R}(T), \quad \mathbf{H}_\tau^* = \text{XLM-R}(T^*) \quad (6)$$

I and T denote meme image and text. T^* denotes the enhanced text. The image is encoded into a sequence of visual patches $\mathbf{H}_v = [v_1, v_2, \dots, v_n]$, $v_i \in \mathbb{R}^d$, and the original and enhanced texts are encoded into token sequences $\mathbf{H}_\tau = [\tau_1, \tau_2, \dots, \tau_m]$, $\tau_i \in \mathbb{R}^d$ and $\mathbf{H}_\tau^* = [\tau_1^*, \tau_2^*, \dots, \tau_m^*]$, $\tau_i^* \in \mathbb{R}^d$, respectively. Here, $*$ \in {ID, TM, CIM, CA} corresponds to the four enhancement steps.

Modal Fusion. The interplay between modalities in memes is often subtle and implicit, making it difficult for traditional single-stage fusion methods to achieve fine-grained alignment. This may lead to misinterpretation of emotional signals. To address this, we propose a dual-stage fusion strategy. In the first stage, we perform a shallow fusion of surface-level features from both modalities, enabling the small model to form an initial, coarse understanding. In the second stage, we deepen this understanding by integrating the shallow features with the enriched, multi-level semantic cues extracted by MLLM. This hierarchical design enables the model to progressively enhance its interpretation of cross-modal signals. We now detail the fusion process.

Stage 1. The primary goal of this stage is to achieve an initial alignment between the shallow features of the visual and textual modalities. Specifically, we concatenate the extracted image patch sequence \mathbf{H}_v with the original text token sequence \mathbf{H}_τ , resulting in an enhanced image patch sequence \mathbf{H}'_v that incorporates information from both modalities.

$$\mathbf{H}'_v = \text{Concat}(\mathbf{H}_v, \mathbf{H}_\tau) = [v_1, \dots, v_n, \tau_1, \dots, \tau_m] \quad (7)$$

The enhanced visual representation $\mathbf{H}'_v \in \mathbb{R}^{N \times d}$ consists of n original image patches and m text tokens. Here, the text tokens can be viewed as special “pseudo-patches” that highlight the textual modality and integrate it into the visual representation. By incorporating these tokens, the fused feature \mathbf{H}'_v embeds initial multimodal information into a unified representation.

Stage 2. After obtaining \mathbf{H}'_v , we proceed with the second-stage fusion by integrating it with the enriched textual tokens derived from MLLM. First, we concatenate the tokens generated from the four-stage enhancement process to form a unified enriched textual representation $\mathbf{H}'_\tau \in \mathbb{R}^{M \times d}$:

$$\mathbf{H}'_\tau = \text{Concat}(\mathbf{H}_\tau^{\text{ID}}, \mathbf{H}_\tau^{\text{TM}}, \mathbf{H}_\tau^{\text{CIM}}, \mathbf{H}_\tau^{\text{CA}}) \quad (8)$$

Next, we employ a bidirectional cross-attention mechanism to deeply fuse the enhanced visual representation \mathbf{H}'_v and enriched textual representation \mathbf{H}'_τ . Specifically, to attend visual representations to textual ones, we compute the attended visual features $\mathbf{H}_v^{\text{att}}$ as follows:

$$\mathbf{H}_v^{\text{att}} = \text{softmax} \left(\frac{Q_\tau K_v^\top}{\sqrt{d_k}} \right) V_v \quad (9)$$

where $\{Q_\tau, K_v, V_v\} = \{\mathbf{H}'_\tau W_Q, \mathbf{H}'_v W_K, \mathbf{H}'_v W_V\}$. Then we fuse $\mathbf{H}_v^{\text{att}}$ with \mathbf{H}'_τ to attain cross-modally enhanced textual representation $\tilde{\tau}$:

$$\tilde{\tau} = \mathbf{H}'_\tau + \mathbf{H}_v^{\text{att}} W_O \quad (10)$$

where W_O denotes linear projection. Similarly, we can attain cross-modally enhanced visual representation \tilde{v} :

$$\tilde{v} = \mathbf{H}'_v + \text{softmax} \left(\frac{Q_v K_\tau^\top}{\sqrt{d_k}} \right) V_\tau W'_O \quad (11)$$

After attaining \tilde{v} and $\tilde{\tau}$, we concatenate these two vectors to obtain the final fused representation of the meme:

$$\mathbf{E}^{\text{meme}} = [\text{mean}(\tilde{v}), \text{mean}(\tilde{\tau})] \quad (12)$$

Finally, we use a linear layer and a softmax classifier for

Modality	Model	MET-MEME				MOOD			
		Accuracy	Precision	Recall	Macro-F1	Accuracy	Precision	Recall	Macro-F1
Unimodal Methods	ResNet-50	29.12	22.87	22.27	21.86	68.23	70.63	66.74	68.35
	ViT	32.65	28.57	26.07	26.39	69.85	74.12	67.78	70.32
	BERT	29.67	26.85	25.45	25.56	68.56	67.71	63.95	65.27
MLLMs	Qwen2.5-VL-7B (zero-shot)	34.79	40.18	35.97	32.30	56.41	60.38	51.48	48.04
	Qwen2.5-VL-7B (sft)	<u>45.40</u>	41.62	<u>41.43</u>	<u>41.03</u>	62.93	65.96	54.36	54.33
	Qwen2.5-VL-32B (zero-shot)	38.58	44.29	40.37	35.30	59.10	59.07	56.58	53.04
	GPT-4.1 (zero-shot)	41.61	<u>44.66</u>	36.92	33.08	72.13	69.09	70.42	67.13
Multimodal Methods	MMBT	26.18	23.50	22.74	22.32	68.00	71.27	66.05	68.00
	VisualBERT	28.92	25.51	24.95	24.92	67.25	79.61	67.25	70.02
	MET_add	38.72	36.23	34.96	35.33	68.19	69.73	65.17	66.93
	MET_cat	40.07	36.37	35.79	36.01	70.32	71.83	67.76	69.51
	Early Fusion	44.27	41.04	39.76	39.94	<u>79.48</u>	81.48	77.11	78.93
	Late Fusion	44.05	40.01	39.94	39.61	<u>79.38</u>	<u>81.77</u>	77.23	79.02
	ALFRED	34.70	34.94	34.70	34.44	79.10	81.63	<u>79.22</u>	<u>79.25</u>
<i>MemoDetector</i> (Ours)		49.57	47.18	44.91	45.33	83.52	84.59	81.19	82.65

Table 1: Meme emotion detection results on MET-MEME and MOOD datasets. The best and second test results are in bold and underlined, respectively.

meme emotion classification:

$$\hat{y} = \text{softmax}(WE^{\text{meme}} + b) \quad (13)$$

Training. The training objective is to minimize the average cross-entropy loss between predictions and ground-truth labels across the dataset. The loss function is as follows:

$$\mathcal{L} = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} L_{CE}(\hat{y}_i, y_i) \quad (14)$$

where \mathcal{D} is the number of samples in the training set, with \hat{y} and y representing predicted label and true label. L_{CE} is the cross-entropy loss function.

4 Experiments

Experimental Setup

Datasets. We evaluate on two public meme datasets: MET-MEME (Xu et al. 2022), a bilingual dataset with 4,000 English and 6,045 Chinese memes labeled across seven emotion classes, and MOOD (Sharma et al. 2024), which contains 10,004 English memes annotated with six emotions.

Baselines. We compare our model against several state-of-the-art approaches for meme emotion detection. These baseline models can be broadly categorized into three groups. (1) Unimodal methods that utilize only visual or textual information, such as ResNet (He et al. 2016), ViT (Dosovitskiy et al. 2020) and BERT (Devlin et al. 2019). (2) MLLM-based methods. We evaluate open-source models including Qwen2.5-VL-7B and Qwen2.5-VL-32B (Bai et al. 2025), as well as the commercial GPT-4.1. (3) Multimodal methods based on small models. These include MMBT (Kiela et al. 2019), VisualBERT (Li et al. 2019), MET_add, MET_cat (Xu et al. 2022), Early Fusion (Pranesh and Shekhar 2020), Late Fusion (Pramanick et al. 2021) and ALFRED (Sharma et al. 2024).

Evaluation Metrics. We report accuracy and macro-F1 as primary metrics, along with macro-averaged precision and recall for completeness. All our scores are the average over 5 runnings with random seeds.

Main Results

Table 1 presents the performance of our method against all baselines on MET-MEME and MOOD. Our approach consistently outperforms all SoTA baselines across all metrics on both datasets, highlighting several key insights. (1) Multimodal methods generally outperform unimodal ones, as they better exploit both visual and textual signals for meme understanding. However, the effectiveness of these models largely depends on the fusion strategy. Poorly designed fusion can lead to inferior performance, as seen with MMBT, which underperforms both ViT and BERT on MET-MEME in terms of Macro-F1. (2) Although MLLMs have the largest number of parameters, their zero-shot performances fall short of expectation, likely due to their inherent limitations and unsuitability for direct classification tasks. However, fine-tuning significantly improves results—Qwen2.5-VL-7B even surpasses GPT-4.1 on MET-MEME, suggesting that MLLMs possess strong potential for deep meme understanding when properly adapted. (3) Both Early Fusion and Late Fusion outperform other baselines across the two datasets, with ALFRED demonstrating notably strong performance on the MOOD dataset. These superior performances can be attributed to the fact that the encoders for both modalities are well pre-trained, enabling them to effectively capture emotion-related features. When combined with relatively effective fusion strategies, these models are able to produce more accurate emotion predictions. (4) Notably, our model significantly surpasses these baseline systems. On MET-MEME, it surpasses the Early Fusion method by 5.3%, 6.14%, 5.15%, and 5.39% in terms of accuracy, precision, recall, and macro-F1, respectively. On MOOD dataset, it achieves improvements of 4.42% in accuracy, 2.96% in precision, 1.97% in recall, and 3.4% in macro-F1 compared to ALFRED. These findings suggest that by incorporating MLLM-generated multi-level meme analysis for textual enrichment, combined with carefully designed dual-stage modal fusion strategy, we achieve substantial performance gains in meme emotion detection.

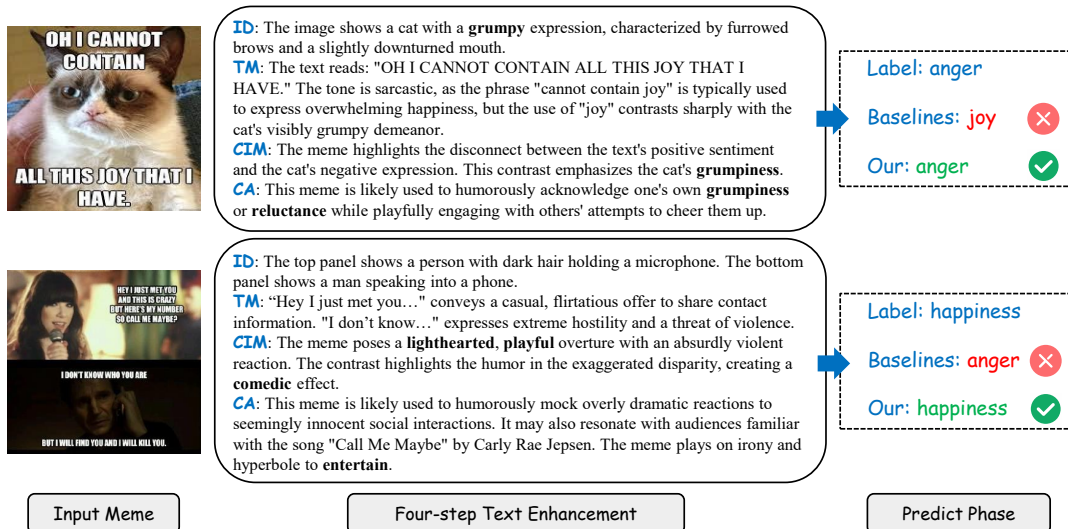


Figure 3: Examples of correctly predicted memes.

Dataset	Method	Accuracy	Precision	Recall	Macro-F1
MET-MEME	Ours	49.57	47.18	44.91	45.33
	w/o DF	48.57	46.51	43.35	43.57
	w/o ID	48.84	45.61	44.23	44.29
	w/o TM	47.82	46.35	43.06	43.78
	w/o CIM	48.63	46.08	43.91	44.32
	w/o CA	48.22	45.16	43.82	43.97
MOOD	Ours	83.52	84.59	81.19	82.65
	w/o DF	82.32	82.93	80.23	81.39
	w/o ID	82.30	82.29	80.56	81.32
	w/o TM	81.62	81.68	79.72	80.54
	w/o CIM	81.73	82.95	79.25	80.75
	w/o CA	81.81	81.85	79.87	80.75

Table 2: Ablation studies by removing components from our proposed framework. "DF" denotes Dual-stage Fusion module, while "ID", "TM", "CIM", and "CA" correspond to the four steps in textual modality enhancement, respectively.

Ablation Study

We perform ablation studies to evaluate the contribution of each component in our model. As shown in Table 2, our full model, without any component removed, achieves the best performance across all metrics, clearly demonstrating the effectiveness of each proposed module. Specifically, when the dual-stage fusion strategy is ablated—replacing it with a simple concatenation of the enhanced textual representation and visual features—the model’s performance drops significantly across both datasets (with a notable 1.76% decrease in F1 on MET-MEME). This highlights the importance of our dual-stage fusion design, which captures the intricate interactions between visual and textual modalities in memes and enables the model to progressively develop a deeper emotional understanding. Moreover, removing any of the four textual enhancement steps leads to a performance drop to varying degrees, which confirms the contribution of each step. Both shallow-level enhancements such as ID and TM, and deeper-level reasoning components like CIM and CA, provide valuable information that boosts the model’s emotion detection performance. Notably, removing the TM

module results in the most significant performance degradation (a 1.55% drop in F1 on MET-MEME and 2.11% on MOOD), highlighting the crucial role of understanding meme text in capturing its emotional implications.

Case Study

To more clearly demonstrate the effectiveness of our method, we present two representative examples, as shown in Figure 3. The first meme presents a sarcastic contrast between the text and the image of an obviously angry cat, which collectively conveys a sense of displeasure rather than genuine joy. Baseline methods, lacking a fine-grained modal fusion mechanism, are misled by the overtly positive textual sentiment and thus fail to predict the correct emotion. In contrast, our method effectively integrates both the surface content and the underlying multimodal cues through a carefully designed dual-stage fusion strategy, enabling it to correctly identify the intended emotional tone. The second meme contrasts the lighthearted lyrics of an English song with an exaggerated overreaction from a person, thereby creating a humorous and playful atmosphere. Baseline methods, lacking knowledge about the song lyrics, fail to capture the intended meaning. In contrast, our method leverages a four-stage text enhancement process, where MLLM contributes rich contextual understanding to uncover the humorous intent behind the lyrics. This enhanced textual representation effectively guides the small model to make the right prediction.

Analyses on Proposed Methods

To gain deeper insights into the effectiveness of our proposed approach, we perform a series of in-depth analyses aimed at addressing the following three key questions.

Q1: What are the advantages of four-step text enhancement? In Table 2, we have already shown that each step of the proposed four-step text enhancement contributes positively to the overall performance. To further investigate its true advantage, we introduce an alternative approach called direct textual enhancement for comparison. In this setting,

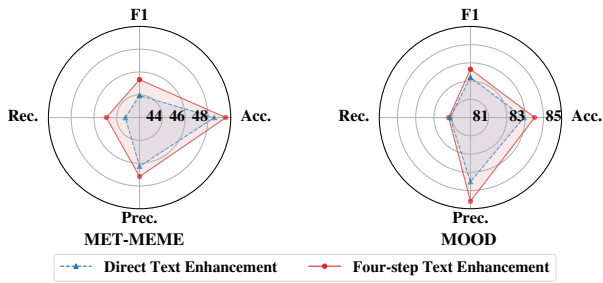


Figure 4: Comparison between four-step text enhancement strategy and direct text enhancement strategy.

Dataset	MET-MEME		MOOD	
Model	Accuracy	Macro-F1	Accuracy	Macro-F1
QwenVL7B	34.79	32.30	56.41	48.04
+ CoT	40.77	36.59	61.05	55.41
+ Ours	45.60	41.22	82.70	81.72
QwenVL32B	38.58	35.30	59.10	53.04
+ CoT	38.18	32.74	55.00	49.31
+ Ours	49.57	45.33	83.52	82.65

Table 3: Effect of MLLM scale and usage paradigm on meme emotion detection performance.

MLLM is prompted to directly infer the meme’s emotional tendency and provide a single-step explanation, which is then used as the enhanced text. As shown in Figure 4, our four-step enhancement strategy consistently outperforms the direct method across all metrics on both datasets. This suggests that our structured, progressive prompting enables the MLLM to develop a deeper and more comprehensive understanding of the meme. The resulting interpretive texts are of higher quality and thus more effective in guiding the small model to make accurate predictions.

Q2: How do the scale of MLLMs and usage paradigms influence the performance of meme emotion detection?

In our framework, QwenVL-32B is adopted by default to perform text enhancement for memes. To investigate how the scale of MLLM and the usage paradigm affect the final performance, we additionally introduce a smaller model, QwenVL-7B, and design three distinct usage paradigms for comparison, as presented in Table 3. Specifically, we explore three paradigms: (1) QwenVL7B/32B: Prompt MLLM directly to generate the emotion label for a given meme. (2) + CoT: Incorporate chain-of-thought prompting to encourage step-by-step reasoning before producing the final answer. (3) + Ours: Apply our proposed framework, where MLLM performs multi-level textual augmentation of memes to help a smaller model in emotion prediction. We can observe that: (1) The direct deployment of both MLLMs struggles since large models are not specifically designed for this classification task (Xu et al. 2024a; Li and Flanigan 2024). (2) CoT prompting significantly improves the performance of QwenVL7B, whereas QwenVL32B exhibits a performance drop. A possible explanation is that smaller models benefit from explicit reasoning structures due to their limited capacity for complex multimodal inference. In contrast, larger

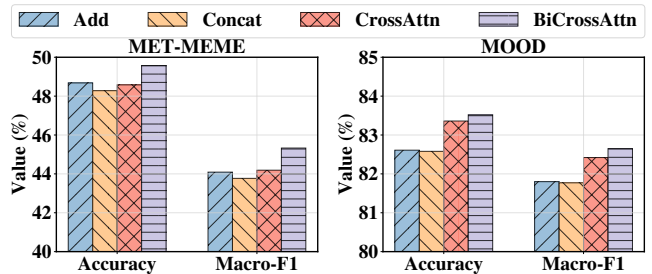


Figure 5: Impact of different second-stage fusion strategies on model performance.

models may already possess sufficient reasoning capabilities, and step-by-step constraints might inhibit their natural inference process, leading to degraded performance. (3) Our approach focuses on training a small model, which not only avoids the impractical cost of fine-tuning large models but also benefits from the insights provided by them, ultimately leading to improved performance. (4) Our method achieves better results when using the 32B model for textual enhancement compared to the 7B variant. This suggests that larger models may generate higher-quality interpretative text, which in turn provides more effective guidance for downstream classification.

Q3: Why do we design a bidirectional cross-attention mechanism in the second stage of modal fusion?

In ablation study, we have demonstrated that our dual-stage fusion strategy outperforms conventional single-stage fusion approaches. Specifically, in the second fusion stage, we employ a bidirectional cross-attention mechanism to integrate the enhanced visual and textual features. To further investigate the superiority of this design, we replace the bidirectional cross-attention module with several commonly used modal fusion strategies (i.e., add, concatenate, and cross-attention). The comparison results are shown in Figure 5. We can observe that our proposed bidirectional cross-attention consistently outperforms other fusion strategies across both datasets, demonstrating the effectiveness and superiority of our design. This mechanism enables the enhanced textual features to attend to the enriched visual representations and vice versa, facilitating a more comprehensive and interactive integration of multimodal information. As a result, the model can better capture the intricate relationships between modalities, ultimately leading to improved performance.

5 Conclusion

In this paper, we address two core challenges in MEU: the lack of fine-grained multimodal fusion strategies and insufficient mining of memes’ implicit meanings and background knowledge. To tackle these issues, we propose *MemoDetector*, a novel framework that leverages MLLMs to deeply interpret meme semantics and contextual cues. In addition, we design a dual-stage modal fusion strategy that performs both shallow and deep integration of visual and textual features, enabling more nuanced emotional understanding. Extensive experiments and thorough analyses demonstrate the effectiveness, robustness, and SoTA performance of our method.

Acknowledgments

This research received support from the National Natural Science Foundation of China under Grant No. 62302441. This work was also supported by the Key Research and Development Program Project of Ningbo Grant No. 2025Z029. The author gratefully acknowledges the support of Zhejiang University Education Foundation Qizhen Scholar Foundation.

References

- Ahmed, N.; Al Aghbari, Z.; and Girija, S. 2023. A systematic survey on multimodal emotion recognition using learning algorithms. *Intelligent Systems with Applications*, 17: 200171.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Bloom, B. S.; et al. 1956. *Taxonomy of educational objectives: The classification of educational goals. Handbook 1: Cognitive domain*. Longmans, Green.
- Cao, R.; Hee, M. S.; Kuek, A.; Chong, W.-H.; Lee, R. K.-W.; and Jiang, J. 2023. Pro-cap: Leveraging a frozen vision-language model for hateful meme detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, 5244–5252.
- Cheng, Z.; Cheng, Z.-Q.; He, J.-Y.; Wang, K.; Lin, Y.; Lian, Z.; Peng, X.; and Hauptmann, A. 2024. Emotion-llama: Multimodal emotion recognition and reasoning with instruction tuning. *Advances in Neural Information Processing Systems*, 37: 110805–110853.
- Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Danapal, G.; Santos, G. A.; da Costa, J. P. C.; Praciano, B. J.; and Pinheiro, G. P. 2020. Sensor fusion of camera and LiDAR raw data for vehicle detection. In *2020 Workshop on Communication Networks and Power Systems (WCNPS)*, 1–6. IEEE.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Ekman, P.; and Cordaro, D. 2011. What is meant by calling emotions basic. *Emotion review*, 3(4): 364–370.
- Gan, C.; Fu, X.; Feng, Q.; Zhu, Q.; Cao, Y.; and Zhu, Y. 2024. A multimodal fusion network with attention mechanisms for visual–textual sentiment analysis. *Expert Systems with Applications*, 242: 122731.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hu, G.; Lin, T.-E.; Zhao, Y.; Lu, G.; Wu, Y.; and Li, Y. 2022. Unimse: Towards unified multimodal sentiment analysis and emotion recognition. *arXiv preprint arXiv:2211.11256*.
- Kiela, D.; Bhooshan, S.; Firooz, H.; Perez, E.; and Teggine, D. 2019. Supervised multimodal bitransformers for classifying images and text. *arXiv preprint arXiv:1909.02950*.
- Li, C.; and Flanigan, J. 2024. Task contamination: Language models may not be few-shot anymore. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18471–18480.
- Li, L. H.; Yatskar, M.; Yin, D.; Hsieh, C.-J.; and Chang, K.-W. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Liang, Z.; Xu, Y.; Hong, Y.; Shang, P.; Wang, Q.; Fu, Q.; and Liu, K. 2024. A Survey of Multimodal Large Language Models. In *Proceedings of the 3rd International Conference on Computer, Artificial Intelligence and Control Engineering*, 405–409.
- Lin, H.; Luo, Z.; Gao, W.; Ma, J.; Wang, B.; and Yang, R. 2024. Towards explainable harmful meme detection through multimodal debate between large language models. In *Proceedings of the ACM Web Conference 2024*, 2359–2370.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Pramanick, S.; Dimitrov, D.; Mukherjee, R.; Sharma, S.; Akhtar, M. S.; Nakov, P.; and Chakraborty, T. 2021. Detecting harmful memes and their targets. *arXiv preprint arXiv:2110.00413*.
- Pranesh, R. R.; and Shekhar, A. 2020. Memesem: a multimodal framework for sentimental analysis of meme via transfer learning. In *4th Lifelong Machine Learning Workshop at ICML 2020*.
- Sharma, S.; Ramaneswaran, S.; Akhtar, M. S.; and Chakraborty, T. 2024. Emotion-Aware Multimodal Fusion for Meme Emotion Detection. *IEEE Transactions on Affective Computing*.
- Shifman, L. 2013. Memes in a digital world: Reconciling with a conceptual troublemaker. *Journal of computer-mediated communication*, 18(3): 362–377.
- Singh, P.; Bauwelinck, N.; and Lefever, E. 2020. Lt3 at semeval-2020 task 8: Multi-modal multi-task learning for memotion analysis. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, 1155–1162.
- Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, 6105–6114. PMLR.
- Vásquez, C.; and Aslan, E. 2021. “Cats be outside, how about meow”: multimodal humor and creativity in an internet meme. *Journal of Pragmatics*, 171: 101–117.

Vlad, G.-A.; Zaharia, G.-E.; Cercel, D.-C.; Chiru, C.-G.; and Trausan-Matu, S. 2020. UPB at SemEval-2020 task 8: Joint textual and visual modeling in a multi-task learning architecture for memotion analysis. *arXiv preprint arXiv:2009.02779*.

Wang, Y.; Chen, W.; Han, X.; Lin, X.; Zhao, H.; Liu, Y.; Zhai, B.; Yuan, J.; You, Q.; and Yang, H. 2024. Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning. *arXiv preprint arXiv:2401.06805*.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.

Xu, B.; Li, T.; Zheng, J.; Naseriparsa, M.; Zhao, Z.; Lin, H.; and Xia, F. 2022. Met-meme: A multimodal meme dataset rich in metaphors. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, 2887–2899.

Xu, H.; Lou, R.; Du, J.; Mahzoon, V.; Talebianaraki, E.; Zhou, Z.; Garrison, E.; Vucetic, S.; and Yin, W. 2024a. LLMs' Classification Performance is Overclaimed. *arXiv preprint arXiv:2406.16203*.

Xu, Y.; Hua, Y.; Li, S.; and Wang, Z. 2024b. Exploring Chain-of-Thought for Multi-modal Metaphor Detection. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 91–101. Bangkok, Thailand: Association for Computational Linguistics.

Younes, A. S.; and Altaxhaineh, A. R. M. 2022. Metaphors and metonymies used in memes to depict COVID-19 in Jordanian social media websites. *Ampersand*, 9: 100087.

Zhu, L.; Zhu, Z.; Zhang, C.; Xu, Y.; and Kong, X. 2023. Multimodal sentiment analysis based on fusion methods: A survey. *Information Fusion*, 95: 306–325.