An Empirical Study on Configuring In-Context Learning Demonstrations for Unleashing MLLMs' Sentimental Perception Capability

Daiqing Wu¹² Dongbao Yang¹ Sicheng Zhao³ Can Ma¹ Yu Zhou⁴

Abstract

The advancements in Multimodal Large Language Models (MLLMs) have enabled various multimodal tasks to be addressed under a zeroshot paradigm. This paradigm sidesteps the cost of model fine-tuning, emerging as a dominant trend in practical application. Nevertheless, Multimodal Sentiment Analysis (MSA), a pivotal challenge in the quest for general artificial intelligence, fails to accommodate this convenience. The zeroshot paradigm exhibits undesirable performance on MSA, casting doubt on whether MLLMs can perceive sentiments as competent as supervised models. By extending the zero-shot paradigm to In-Context Learning (ICL) and conducting an indepth study on configuring demonstrations, we validate that MLLMs indeed possess such capability. Specifically, three key factors that cover demonstrations' retrieval, presentation, and distribution are comprehensively investigated and optimized. A sentimental predictive bias inherent in MLLMs is also discovered and later effectively counteracted. By complementing each other, the devised strategies for three factors result in average accuracy improvements of 15.9% on six MSA datasets against the zero-shot paradigm and 11.2% against the random ICL baseline.

1. Introduction

Equipping models with emotional intelligence has been a fascinating yet vital challenge over the past few decades (Zhao et al., 2022; Zhang et al., 2023). Studies on various facets of emotions and sentiments in numerous do-



Figure 1. Comparison of fully-supervised models, few-shot models, and MLLMs based on average accuracy and annotated data requirement across six MSA datasets. The MLLMs' zero-shot paradigm, although avoiding the laborious annotation, exhibits a substantial performance gap compared to fully-supervised models. With proper demonstration configuration, this gap can be notably narrowed by In-Context Learning (ICL).

mains have flourished. Among them, Multimodal Sentiment Analysis (MSA) aims to classify the sentiment polarity embedded in multimodal data. As corroborated by empirical (Zadeh et al., 2017) and theoretical (Huang et al., 2021) evidence, the synergy between modalities facilitates more comprehensive modeling of sentiment clues compared to unimodal data. This superiority has led to growing interest in MSA from academia and industry (Zhang et al., 2018a; Yue et al., 2019).

In this paper, we focus on MSA with the input form of image-text posts, whose number grows exponentially thanks to the prosperity of social media. The current research landscape can be broadly grouped into two streams: post-level branch (Xu & Mao, 2017), which identifies the overall sentiment of posts, and aspect-level branch (Xu et al., 2019), which probes the sentiment associated with specific aspect terms within the context of posts. Leveraging elaborate modules and a large volume of annotated data, both streams have attained remarkable achievements. Entering the era of Multimodal Large Language Models (MLLMs), various multimodal tasks can be reformulated as filling textual prompts and properly accomplished by MLLMs in a zero-shot manner (Yin et al., 2023). This emerging paradigm avoids timeintensive fine-tuning and task-specific annotation. However, recent surveys (Yang et al., 2023b; Lian et al., 2024) have

¹Institute of Information Engineering, Chineses Academy of Sciences ²School of Cyber Security, University of Chinese Academy of Sciences ³Tsinghua University ⁴VCIP & TMCC & DISSec, College of Computer Science, Nankai University. Correspondence to: Dongbao Yang <yangdongbao@iie.ac.cn>, Yu Zhou <yzhou@nankai.edu.cn>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

Configuring ICL Demonstrations for Unleashing MLLMs' Sentimental Perception Capability



Figure 2. Comparison between MLLMs' zero-shot paradigm and ICL. In addition to the test sample, ICL sequences three demonstrations with inputs and corresponding outputs, facilitating more precise sentiment predictions for MLLMs.

revealed that, under the zero-shot paradigm, MLLMs lag behind supervised models by a significant margin on MSA, as also illustrated in Figure 1. Given the success of MLLMs on tasks such as image captioning (Li et al., 2023) and visual question answering (Driess et al., 2023), it is natural to wonder whether the capabilities of MLLMs to perceive sentiment are yet to be fully explored. If not, how can we fully tap into their potential at a minimal cost?

A feasible answer is In-Context Learning (ICL), which extends MLLMs' zero-shot paradigm to a few-shot scenario by sequencing a series of input-output pairs as demonstrations (Brown et al., 2020). Figure 2 illustrates this process with an example. In this manner, ICL showcases the formulation of the task and the mapping between inputs and outputs (Pan et al., 2023), which has been proven beneficial for MLLMs in both unimodal and multimodal tasks (Chen et al., 2023a; Li et al., 2024). Despite its success, the efficacy of ICL heavily relies on the retrieval (Zhang et al., 2022), presentation (Li et al., 2024), and distribution (Lyu et al., 2023) of demonstrations. Studying the impact of these factors is a prevalent subject in natural language processing, yet there's a notable absence of necessary attention in multimodal settings, particularly for MSA.

To fill this gap and unleash the potential of MLLMs, we systematically investigate the configuration of ICL demonstrations in MSA. Specifically, we delve deeper into three crucial factors underexplored in current studies. (1). The similarity measurement of multimodal data for demonstration retrieval. The similarity between the test sample and demonstration is positively correlated with the effectiveness of ICL (Liu et al., 2022). The mainstream measurement method directly aggregates images' and texts' similarity scores (Yang et al., 2022). Nevertheless, our experiments reveal that this approach overlooks fine-grained aspect terms and fails to weigh the significance of modalities. In response, we refine and customize similarity measurements specifically tailored for MSA. (2). The trade-off between multimodal information presented in the demonstrations. Owing to the disparity in information density, images and texts each have unique pros and cons under different

scenarios. Therefore, image captioning (Yang et al., 2023a) and text-to-image generation (Alimisis et al., 2024) are commonly employed to convert between the two modalities to furnish supportive information in multimodal tasks. Inspired by this, we explore the effects of modality compositions on ICL, ultimately arriving at the most efficacious modality presentation. (3). The biases introduced by sentiment **distribution.** It has been observed on various tasks (Yang et al., 2023c; Li et al., 2024; Baldassini et al., 2024) that MLLMs are prone to be affected by the biases in ICL sequences, a phenomenon we also validate on MSA. Driven by this discovery, we devise various distribution protocols to intentionally incorporate sentiment biases and probe their influences. Comprehensive experiments reveal that appropriate sentiment biases can counterbalance the inherent predictive bias of MLLMs, thereby promoting fair prediction. In summary, our contributions are three-fold:

- Through configuring ICL demonstrations, we unleash the potential of MLLMs on MSA, validating that MLLMs are competent in perceiving sentiment.
- We investigate and optimize three key factors covering the retrieval, presentation, and distribution of demonstrations in ICL on MSA. During the process, a sentimental predictive bias inherent in MLLMs is discovered and mitigated, facilitating fairness in sentiment prediction.
- By complementing each other, the ICL strategies tailored for the three factors improve the accuracy of MLLMs by an average of 15.9% on six MSA datasets compared to the zero-shot paradigm, and 11.2% compared to the random ICL baseline.

2. Related Works

2.1. Multimodal Sentiment Analysis

With the explosion of multimodal data from social media platforms in recent years, MSA has evolved as a pivotal topic in affective computing (Zhu et al., 2024; Yang et al.,



Figure 3. Illustration of the three factors to be investigated and optimized, during which we aim to address the following questions. (a). How do we measure the similarity score between multimodal data? (b). How do we decide which modality should be presented in the input? (c). What kind of impact does the sentiment distribution of demonstrations have?

2024). Image-text posts (Xu & Mao, 2017) and videos (Zadeh et al., 2017) are two major forms of input in MSA. Our investigation centers on image-text posts, as they are more congruent with the input specifications of MLLMs. Within this domain, the majority of research adheres to a similar methodology, typically initiating with the extraction of unimodal features followed by their fusion for sentiment classification. Over the past few years, notable advancements have been made in both post-level (Xu, 2017; Yang et al., 2021b; Li et al., 2022; Wei et al., 2023; Wu et al., 2024a) and aspect-level (Xu et al., 2018; Ling et al., 2022; Zhou et al., 2023; Peng et al., 2024) branches, by refining fusion strategies and learning from thousands of handannotated data. However, as highlighted by Yu et al. (2022), the acquisition of such data is unrealistic in practical settings. Consequently, few-shot MSA has been proposed (Yu et al., 2022) and developed (Yang et al., 2023a). Although it reduces data requirements, the time-intensive fine-tuning persists as an impediment to the cost-effective application of MSA.

The advent of MLLMs presents a viable alternative via the zero-shot paradigm (Lian et al., 2024), which is subsequently augmented by ICL. By harnessing the general knowledge acquired during pre-training, MLLMs are endowed with the proficiency to address a wide array of downstream tasks, requiring merely a minimal set of annotated samples and no gradient updates. Compared with fullysupervised and few-shot MSA, the ICL of MLLMs accords better with practical applications, holding the potential to evolve into a prevailing trend for future development.

2.2. In-Context Learning

The scaling up of model size has empowered Large Language Models (LLMs) to generalize to unseen tasks through analogical learning (Dong et al., 2024), a capability known as ICL (Wang et al., 2024). First identified in GPT-3 (Brown et al., 2020), ICL has sparked numerous studies on LLMs. Among these, a predominant emphasis has been placed on the impact of demonstration configurations. Through exhaustive explorations, various effective strategies have been proposed from distinct angles (Lu et al., 2022; Li & Qiu, 2023; Wu et al., 2023; Levy et al., 2023). Certain ICL strategies of LLMs also demonstrate proficiency in tackling multimodal tasks (Yang et al., 2022; Cai et al., 2023) by converting them into textual format.

To transfer the ICL capability into MLLMs, Flamingo (Alayrac et al., 2022) interleaves cross-attention modules that receive visual input into the pre-trained LLMs. Inspired by its success, a series of open-sourced MLLMs armed with ICL capability are developed, such as Open-Flamingo (Awadalla et al., 2023), IDEFICS (Laurençon et al., 2023) and Kosmos (Huang et al., 2023). Parallel to the case in LLMs, the performance of MLLMs exhibits significant variability in response to diverse demonstrations. A line of studies methodically probes into such effects in image captioning (Yang et al., 2023c) and visual question answering (Li et al., 2024; Chen et al., 2024; Qin et al., 2024), uncovering valuable insights of ICL and important properties of MLLMs. Nonetheless, the unique challenges intrinsic to MSA remain inadequately overcome. We aim to fill this blank through a comprehensive empirical exploration.

3. Configuring ICL Demonstrations

3.1. Task Definition

MSA is a classification task where the target categories are sentiment polarities. Depending on the classification objective, MSA is grouped into post-level MSA and aspectlevel MSA. For post-level MSA, the input comprises an image-text post (i, t), and the target is to identify the overall sentiment. For aspect-level MSA, the input includes an additional aspect a: (i, t, a), and the target shifts to identify the sentiment of the aspect within the context of the post. We wrap the input as \mathcal{I} and the output as \mathcal{O} for uniformity.

Before feeding into the MLLM, the test sample $\hat{\mathcal{I}}$ and n demonstrations $(\mathcal{I}, \mathcal{O})$ are templatized into a multimodal sequence $\mathcal{S} = \{\mathcal{P}; (\mathcal{I}_1, \mathcal{O}_1); (\mathcal{I}_2, \mathcal{O}_2); \ldots; (\mathcal{I}_n, \mathcal{O}_n); \hat{\mathcal{I}}\},$ where \mathcal{P} is a textual prompt including the task description and the set of target categories. A 3-shot case is illustrated in Figure 2 (b). The output $\hat{\mathcal{O}}$ is subsequently generated by MLLM $M(\cdot)$ predicting the next single token:

$$\hat{\mathcal{O}} = \underset{\mathcal{T}}{\arg\max} M(\mathcal{T}|\mathcal{S}).$$
(1)

Achieving desirable outcomes through ICL involves the interplay of various factors that each contribute to the process. Among these, the retrieval, presentation, and distribution of demonstrations have been empirically verified to wield significant impact (Zhang et al., 2022; Li et al., 2024; Lyu et al., 2023). In light of MSA's multimodal and affective nature, we concretize these aspects into three key factors to optimize: similarity measurement, modality presentation, and sentiment distribution. The three factors are briefly illustrated in Figure 3 (a), (b), and (c), respectively, and are elaborated on in the subsequent sections.

3.2. Similarity Measurement

In demonstration retrieval, prior research has reached two primary agreements: demonstrations bearing greater similarity to the test sample are more beneficial for the reasoning process of MLLMs (Liu et al., 2022), and demonstrations should be ordered from lowest to highest similarity in the ICL sequence (Lu et al., 2022). Given adherence to these two fundamental principles, measuring the similarity between multimodal data emerges as the crux of the matter.

Considering the input of a demonstration (i, t, a) and the test sample $(\hat{i}, \hat{t}, \hat{a})$, there are currently three predominant strategies to measure similarity. Here we discuss the scenario of aspect-level MSA due to its generalizability.

Image Based (I) and **Text Based (T)** strategies. Under unimodal settings, the similarity score \mathcal{K} can be straightforwardly obtained by computing the cosine similarity in the embedding space of a pre-trained encoder \mathcal{E} . Indexing by images, this process can be described as:

$$\mathcal{K}_{I} = \frac{\mathcal{E}(\boldsymbol{i}) \odot \mathcal{E}(\hat{\boldsymbol{i}})}{||\mathcal{E}(\boldsymbol{i})|| \cdot ||\mathcal{E}(\hat{\boldsymbol{i}})||}.$$
(2)

For simplicity, we abbreviate it as $\mathcal{K}_I = C(i, \hat{i})$. Likewise, with text indexing, it would be $\mathcal{K}_T = C(t, \hat{t})$.

Table 1. Formulations of six devised strategies for measuring similarity between multimodal data.

Strategy	Formulation
А	$\mathcal{K}_A = C(\boldsymbol{a}, \boldsymbol{\hat{a}})$
IA	$\mathcal{K}_{IA} = \mathcal{K}_I + \mathcal{K}_A$
TA	$\mathcal{K}_{TA} = \mathcal{K}_T + \mathcal{K}_A$
ITA	$\mathcal{K}_{ITA} = \mathcal{K}_I + \mathcal{K}_T + \mathcal{K}_A$
WIT	$\mid \mathcal{K}_{WIT} = oldsymbol{lpha} \cdot \mathcal{K}_I + oldsymbol{eta} \cdot \mathcal{K}_T$
WITA	$\mid \mathcal{K}_{WITA} = \boldsymbol{\alpha} \cdot \mathcal{K}_{I} + \boldsymbol{\beta} \cdot \mathcal{K}_{T} + \boldsymbol{\gamma} \cdot \mathcal{K}_{A}$

Image-Text Based (IT) strategy. It extends the unimodal strategies to a multimodal form by aggregating the unimodal similarity scores: $\mathcal{K}_{IT} = \mathcal{K}_I + \mathcal{K}_T$. This strategy is also referred to as RICES (Yang et al., 2022).

Despite their proven efficacy across various tasks, these strategies fall short in two critical areas when applied to MSA. In particular, they fail to consider the aspect-specific relevance and overlook the unequal significance of modalities. To probe the impact of these shortcomings and optimize accordingly, we further devise six strategies in Table 1. Among these, the **Aspect Based** (A) strategy is an unimodal measurement indexed by aspects. The **Image-Aspect Based** (IA), **Text-Aspect Based** (TA), and **Image-Text-Aspect Based** (ITA) strategies are the aspect-inclusive versions of the I, T, and IT strategies, respectively. The **Weighted Image-Text Based** (WIT) and **Weighted Image-Text-Aspect Based** (WITA) strategies expand upon the IT and ITA strategies by incorporating weights α , β , γ to modulate the impact of each modality.

3.3. Modality Presentation

The density of information embedded in images and texts naturally differs (Chen et al., 2023b; Wu et al., 2024b). Images convey information in a broadly ranged yet abstract manner, whereas text is typically more precise and concise. By converting images to text through image captioning, salient objects are emphasized and the irrelevant finegrained details are omitted (Xu et al., 2015; Yang et al., 2023c). Conversely, transforming texts into images through text-to-image generation can supplement the background and context beyond the description of texts. The extra information is conducive to the elicitation of previously latent emotions (Kosti et al., 2017). As exemplified in Figure 3 (b), the captioned text summarizes the flyer's contents, and the generated image portrays the unmentioned details of the barbecue. Recognizing these benefits, various multimodal studies leverage these techniques to derive auxiliary modalities to promote performance.

This drives us to investigate whether such success can be replicated in multimodal ICL, where the form in which information is presented within the ICL sequence is of critical importance. Specifically, we obtain the captioned texts from

Table 2. Average accuracy across 4,8,16-shot demonstrations retrieved based on varying similarity measurements. R strategy represents the random retrieval.

Strategy	Post-	Level	Aspect-Level			
Suutegy	MVSA-S	MVSA-M	Twitter-15	Twitter-17		
R	49.2	60.8	57.4	56.4		
Ι	56.5	64.9	59.1	56.7		
Т	56.0	66.2	58.7	57.0		
IT	55.7	66.2	61.4	57.6		
А	-	-	61.3	57.4		
IA	-	-	59.6	58.0		
TA	-	-	60.9	57.3		
ITA			61.0	58.1		

the MLLM itself, and the generated images from the diffusion model (Rombach et al., 2022). Subsequently, we reconstruct the inputs of demonstrations and the test sample with combinations of original and auxiliary modalities. An input form combining the captioned text and generated image is depicted at the bottom of Figure 3 (b), and a broader range of combinations is investigated in the experiments.

3.4. Sentiment Distribution

Under ICL, MLLMs have been shown to manifest a shortcut effect (Lyu et al., 2023). Upon encountering a test sample, MLLMs tend to duplicate the output from one of the demonstrations, and this tendency intensifies when multiple demonstrations share identical outputs. This exposes MLLMs' susceptibility to biases presented within the ICL sequence, where divergent biases can lead to markedly different predictions.

Regarding MSA, its classification nature and limited categories inevitably result in certain sentiment biases within the ICL sequence, reflected in the sentiment distribution of demonstrations. To analyze the potential impact of these biases and utilize them to optimize the ICL configuration, we formulate five distribution protocols that regulate the sentiment distribution of demonstrations, thereby injecting distinct sentiment biases. The first two protocols serve to highlight the spectrum of the impact of biases. (1). Ideal protocol: every demonstration aligns with the test sample's sentiment. (2). Contrary to Ideal protocol: every demonstration differs from the test sample's sentiment. Utilizing the ground truth from test samples, they introduce extreme biases into the ICL sequence to establish the theoretical ceiling and floor. In contrast, the following protocols are devised for practical ICL configuration. (3). Unlimited protocol: demonstrations are retrieved without distribution restriction. (4). Category Balanced protocol: each sentiment class carries an identical number of demonstrations. (5). Identical to Support Set protocol: the sentiment distribution of demonstrations mirrors that of the support set.



Figure 4. Average accuracy across 4,8,16-shot demonstrations retrieved based on the WIT and WITA strategies.

4. Experiments

4.1. Datasets and Implementation Details

In exploratory experiments, IDEFICS-9B (Laurençon et al., 2023) is selected as a representative MLLM (Shukor et al., 2024). Datasets MVSA-S, MVSA-M (Niu et al., 2016) are utilized for post-level MSA, and Twitter-15 (Zhang et al., 2018b), Twitter-17 (Lu et al., 2018) are chosen for aspectlevel MSA. For each dataset, demonstrations are retrieved from the support set sampled as Yang et al. (2023a), which accounts for 1% of the training set. MLLMs' performances are evaluated on the original test set, with accuracy serving as the primary metric. The three factors are investigated separately, and the optimal strategies are subsequently combined for a comprehensive assessment. To validate the generalizability, additional experiments are carried out using Open-Flamingo2-9B (Awadalla et al., 2023) and on two other MSA datasets TumEmo (Yang et al., 2021a) and MASAD (Zhou et al., 2021).

Due to practical considerations, we choose the above model scale, which also aligns with the primary research scope of Yang et al. (2023c); Li et al. (2024). By default, we adopt the **IT** strategy for similarity measurement, compose the input by the original image and text, and put no constraints on the distribution. Variations are introduced only to the pertinent settings when investigating a specific factor.

4.2. Results and Analysis

4.2.1. SIMILARITY MEASUREMENT

Table 2 contrasts the impact of similarity measurement. Here we employ CLIP (Radford et al., 2021) as the pre*Table 3.* Average accuracy over 4,8,16-shot settings with the inputs composed of different modalities. "I" abbreviated for Image, T–Text, C–Captioned Text, G–Generated Image.

Modality	Post-	Level	Aspect	Aspect-Level			
	MVSA-S	MVSA-M	Twitter-15	Twitter-17			
Ι	51.7	56.7	57.4	53.4	54.8		
С	43.9	49.5	56.2	52.1	50.4		
I,C	56.9	59.3	54.3	50.1	55.3		
Т	46.3	56.0	61.7	59.0	55.8		
G	46.8	55.5	54.8	55.2	53.1		
T,G	47.7	58.1	54.0	54.8	53.6		
I,T	55.7	66.2	61.4	57.6	60.2		
I,G	54.5	58.8	54.5	53.8	55.4		
C,T	49.7	60.3	62.6	56.7	57.3		
C,G	47.5	53.7	53.1	53.1	51.8		
I,C,T	55.3	64.4	60.3	55.6	58.9		
I,T,G	50.8	62.0	59.7	54.8	56.8		
C,T,G	47.4	58.2	56.5	54.1	54.1		
I,C,G	50.3	57.0	52.9	52.1	53.1		
I,C,T,G	49.2	61.4	58.7	53.1	55.6		

trained encoder for both images and texts following previous studies (Yang et al., 2022; Li et al., 2024). Observably, all measurements surpass random retrieval, revealing their validity in reflecting the relationships among multimodal samples. Moreover, in aspect-level MSA, the consistent improvement achieved by factoring in aspect similarity underscores the aspect's significance. Counterintuitively, the IT or ITA strategies do not always yield optimal results, despite considering the most comprehensive information. Explanations for this necessitate an analysis in conjunction with the results of WIT and WITA strategies. As evidenced in Figure 4, the performance of these strategies fluctuates in response to alterations in the weighting factors. This emphasizes the criticality of maintaining a proper balance among similarity components, which the IT and ITA strategies fail to achieve, thereby leading to their deficiency.

Delving into the impact of weighting factors, we uncover that prioritizing textual similarity ($\alpha < \beta$) yields a pronounced benefit for the WIT strategy in post-level MSA. This benefit peaks at α : $\beta = 2$: 8, where the **WIT** strategy achieves an accuracy of 57.6 on MVSA-S and 66.7 on MVSA-M, outperforming all competing strategies. In aspect-level MSA, as illustrated in Figure 4 (b), the WITA strategy should, instead, accord priority to the similarity of aspects. Interestingly, under this priority, the relationship between image and text weights is reversed. We attribute these to the fact that aspects are most directly correlated with sentiment prediction, and the similarity among aspects tends to eclipse the textual similarity, given that these aspects typically originate from the texts. When $\alpha : \beta = 7 : 3$ and $(\alpha + \beta)$: $\gamma = 2: 8$, the WITA strategy achieves peak accuracy of 61.8 on Twitter-15 and 59.0 on Twitter-17, making it the optimal strategy for aspect-level MSA.



Figure 5. Evaluation of ICL's "Task Learning" effect by progressively incorporating modalities into the inputs.

4.2.2. MODALITY PRESENTATION

Table 3 compares inputs composed of various modalities. When input is confined to single-modal information, presenting the text modality generally leads to superior outcomes, reaffirming the significance of texts in MSA. As more modalities are incorporated, the MLLM can benefit from processing multimodal information, thereby enhancing sentiment prediction in most cases. Noticeably, substituting original modalities with auxiliary ones (e.g. from (I, T) to (C, G)) leads to universal performance degradation. This implies that the potential benefits of modality conversions, which involve highlighting salient objects or providing richer context, are outweighed by the inherent drawbacks. These drawbacks encompass information loss and the generation of extraneous noise during the conversion process. Strangely, augmenting original modalities with auxiliary ones (e.g. from (I, T) to (I, C, T)) still results in a performance drop, despite the pure addition of information without any loss. Drawing from insights in Pan et al. (2023), we associate this with the impairment of the "Task Learning" effect in ICL. Specifically, Pan et al. (2023) decompose ICL's role into "Task Recognition" and "Task Learning". The former prompts the task format for MLLMs to apply their prior knowledge, and the latter aids MLLMs in building mapping relationships between inputs and outputs. Incorporating additional modalities complicates the inputs, making it more challenging for the MLLM to learn these relationships accurately.

To validate this, we design the experiments in Figure 5 (a) to quantitatively evaluate the "Task Learning" effect. Before feeding into the MLLM, each sentiment is substituted with an animal unrelated to the input semantics according to a predefined mapping. This approach ensures that the MLLM can no longer rely on pre-trained prior knowledge but is instead compelled to learn the input-output mapping solely from ICL. As shown in Figure 5 (b), the "Task Learning" effect diminishes as additional modalities are incorporated into the input, supporting our explanation. To this end, employing original images and text to form inputs has empirically



Figure 6. Impact of distribution protocols on MLLM's sentiment predictions. In addition to overall accuracy, precision and recall for each category are studied in 2-16 shot settings. The two selected datasets possess distinctive test set distributions, which are detailed above the accuracy charts. In (a-2,3,4) and (b-2,3,4), the results of the same protocol under various shots are grouped into clusters of three colors, and the results of different protocols under 16-shot are connected in a fixed order.



Figure 7. Illustration of SLR under 6-shot setting. Taking SLR-Neutral as an example, the ICL sequences of the two Neutral test samples contain 4 and 3 Neutral demonstrations respectively, thus SLR-Neutral is computed as (3/6 + 4/6)/2 = 7/12.

proven optimal, as it attains the finest equilibrium between information enrichment and input complexity.

4.2.3. SENTIMENT DISTRIBUTION

When assessing the impact of sentiment biases within the ICL sequence, it is intuitive to assume that the more demonstrations sharing the same sentiment as the test sample, the more likely the MLLM will make the correct prediction. Indeed, this is one of the primary mechanisms by which distribution protocols affect the performance of MLLMs. To quantify this, we introduce a metric termed "Same Label Rate (SLR)". First, the proportion of demonstrations with identical sentiments to each test sample is calculated. Then, these proportions are averaged across a collection of test samples. Depending on whether the collection comprises the entire test set or all test samples of a specific category, SLR can be calculated for all test samples (SLR-Overall) or the corresponding sentiment (e.g., SLR-Positive). Figure 7 provides an example illustrating SLR. Utilizing SLR, Figure 6 evaluates the nuanced impact of distribution protocols.

On MVSA-S, when the number of shots is fixed, overall accuracy exhibits a stable positive correlation with SLR-Overall (Figure 6 (a-1)). Under the extreme sentiment biases imposed by the **Ideal** and **Contrary to Ideal** protocols, MLLM's sentiment prediction tends to adapt correspondingly, with this trend becoming more pronounced as the number of shots increases. It validates the vulnerability of MLLM to sentiment biases within the ICL sequence. Inspecting each sentiment category (Figure 6 (a-2,3,4)),



Figure 8. Confusion matrices of IDEFICS under the Unlimited and Category Balanced protocols on Twitter-17 (16-shot setting).

MLLM performs notably well on the positive samples, while it shows considerably diminished precision on neutral samples and recall on negative samples. This reveals a potential predictive bias in the MLLM, which favors positive and neutral over negative. Among three practical protocols, the Unlimited protocol stands out with superior performance, particularly in positive and negative samples. Compared to the Contrary to Ideal protocol, it attains a higher SLR in these two categories. This indicates that similarity-based retrieval improves SLR upon the original distribution in the support set. Though resulting in the least favorable performance, the Category Balanced protocol exhibits an intriguing phenomenon. It achieves the highest SLR in the rarest category (Neutral) while obtaining the lowest SLR in a more prevalent category (Negative), leading the model to classify samples of the latter as the former.

On Twitter-17, most distribution protocols still adhere to the positive correlation between SLR-Overall and accuracy (Figure 6 (b-1)), except for the **Category Balanced** protocol. It substantiates SLR as a feasible indicator to reflect the efficacy of ICL sequences and drives us to explore the reasons behind this exception. In parallel to MVSA-S, the MLLM shows slightly lower precision than recall on positive and neutral samples (Figure 6 (b-2,3)). Whereas, on negative samples (Figure 6 (b-4)), recall remains remarkably lower than precision, which is intuitively illustrated in Figure 8. This further validates that the predictive bias observed in MVSA-S is irrelevant to datasets or demonstrations, and might be intrinsic to the MLLM itself. Such bias could potentially originate from the curation of the pre-training data, where most negative instances are filtered out, impairing

	Model & Strategy	Support Set		Post-Level			Mean		
	house of serangy	Support Set	MVSA-S	MVSA-M	TumEmo	Twitter-15	Twitter-17	MASAD	
	Zero-Shot Paradigm	_	38.6	56.5	46.6	60.7	54.7	73.0	55.0
	ICL Random 16-shot	1% Training Set	49.9	60.4	52.1	61.1	57.5	77.2	59.7
	ICL RICES 16-shot (Yang et al., 2022)	1% Training Set	57.9	64.2	61.4	61.6	61.5	82.7	64.9
	ICL SQPA 16-shot (Li et al., 2024)	1% Training Set	53.6	62.2	64.6	55.8	50.4	74.2	60.1
MLLM	ICL MMICES 16-shot (Chen et al., 2024)	1% Training Set	59.2	65.3	60.5	61.6	56.6	80.5	63.9
IDEFICS	ICL Ours 16-shot	1% Training Set	66.5	67.7	63.4	67.0	62.0	88.2	69.1
	ICL Random 16-shot	Full Training Set	50.0	60.2	52.1	61.0	57.5	77.3	59.7
	ICL RICES 16-shot (Yang et al., 2022)	Full Training Set	58.3	65.5	61.7	62.7	61.8	83.4	65.6
	ICL SQPA 16-shot (Li et al., 2024)	Full Training Set	55.8	63.2	59.7	57.3	52.1	75.1	60.5
	ICL MMICES 16-shot (Chen et al., 2024)	Full Training Set	63.0	64.6	62.3	59.2	54.5	83.5	64.5
	ICL Ours 16-shot	Full Training Set	68.5	69.5	65.0	69.0	63.7	89.8	70.9
	Zero-Shot Paradigm	—	52.5	59.2	27.8	34.4	47.1	69.6	48.4
	ICL Random 16-shot	1% Training Set	59.2	67.3	34.8	57.6	51.9	75.0	57.6
	ICL RICES 16-shot (Yang et al., 2022)	1% Training Set	61.5	67.9	42.0	61.2	55.8	84.7	62.2
	ICL SQPA 16-shot (Li et al., 2024)	1% Training Set	59.5	68.3	44.9	56.2	53.9	75.5	59.7
MLLM	ICL MMICES 16-shot (Chen et al., 2024)	1% Training Set	63.3	68.0	46.4	56.9	55.4	80.8	61.8
Open-Flamingo	ICL Ours 16-shot	1% Training Set	65.2	68.5	52.7	68.8	61.4	87.6	67.4
open i minigo	ICL Random 16-shot	Full Training Set	59.0	67.4	34.5	57.6	51.7	75.1	57.5
	ICL RICES 16-shot (Yang et al., 2022)	Full Training Set	62.2	68.0	44.3	61.5	56.4	85.5	63.0
	ICL SQPA 16-shot (Li et al., 2024)	Full Training Set	61.4	68.3	43.5	60.3	54.7	76.3	60.8
	ICL MMICES 16-shot (Chen et al., 2024)	Full Training Set	64.7	68.1	47.5	59.8	57.0	82.4	63.3
	ICL Ours 16-shot	Full Training Set	66.3	68. 7	54.1	70.4	62.6	89.3	68.6
SOTA Few-Shot	Models	_	69.8	68.3	58.1	67.3	61.9	84.1	68.3
SOTA Fully-Sup	ervised Models	—	79.2	73.5	66.7	81.1	76.4	95.6	78.8

Table 4. Accuracy comparison of MLLMs, SOTA few-shot models and SOTA fully-supervised models. We report the ICL performance of MLLMs under a 16-shot setting, as it includes sufficient demonstrations while not exceeding MLLMs' input limit.

Table 5. Our final strategies for configuring ICL demonstrations on six MSA datasets. Due to the diversity of task focus and data distribution, the optimal strategies for the three factors of each dataset are different.

Dataset	Retrieval	Presentation	Distribution
MVSA-S	WIT $[\alpha : \beta = 2 : 8]$	Image, Text	Unlimited
MVSA-M	WIT $[\alpha : \beta = 2 : 8]$	Image, Text	Unlimited
TumEmo	WIT $[\alpha : \beta = 2 : 8]$	Image, Text	Unlimited
Twitter-15	WITA $[\alpha : \beta = 7 : 3, (\alpha + \beta) : \gamma = 2 : 8]$	Image, Text	Category Balanced
Twitter-17	WITA $[\alpha : \beta = 7 : 3, (\alpha + \beta) : \gamma = 2 : 8]$	Image, Text	Category Balanced
MASAD	WITA $[\alpha : \beta = 7 : 3, (\alpha + \beta) : \gamma = 2 : 8]$	Image, Text	Unlimited

MLLM's discernment of negative samples. In contrast to MVSA-S, the rarest category on Twitter-17 is negative. Under this disparity, the **Category Balanced** protocol instead inclines the model to predict the negative category over the positive category, offsetting MLLM's predictive bias and attaining the peak performance across both categories.

To summarize, two conclusions have been deduced. Firstly, MLLM exhibits a predictive bias in sentiment analysis, with an inclination to avoid negative predictions. Secondly, in datasets with less prevalent negative samples, the **Category Balanced** protocol emerges as the optimal selection by mitigating MLLM's bias effectively. Otherwise, the **Unlimited** protocol is the preferable alternative.

4.2.4. OVERALL STRATEGY

The overall strategy is crafted by integrating three finely optimized factors, as presented in Table 5. Table 4 compares our devised strategies and other models. Among the ICL baselines we compared, "ICL Random" randomly selects demonstrations from the support set. "ICL RICES" retrieves samples based solely on the similarity measured by the **WIT** strategy. "ICL SQPA" and "ICL MMICES" are originally designed for the Visual Question Answering task. The former assigns pseudo-labels to the samples in the support set, while the latter first retrieves images and then retrieves texts. In our case, pseudo-labels are assigned as random sentiment categories, and "ICL MMICES" is identical to the I strategy.

From Table 4, our devised strategies exhibit consistent and pronounced superiorities against other ICL baselines, extending to datasets and MLLMs unused in the exploration. By expanding the support set to the full training set, our approach enables both MLLMs to outperform the SOTA few-shot models. Notably, configuring demonstrations operates independently of annotations, ensuring that the annotated data required for each test sample remains constant. This renders MLLMs' ICL more efficient and accurate in MSA than few-shot models. However, compared to fullysupervised models, ICL still exhibits a performance gap, particularly noticeable in certain datasets. It suggests that the role of fine-tuning with specific data still cannot be easily replaced. Overall, Table 4 demonstrates the effectiveness and generalizability of our devised ICL strategies, as well as the considerable potential of ICL for further development.

5. Conclusion and Discussion

In this paper, we seek to unleash the sentimental perception capability of MLLMs through ICL. As a pioneering effort that applies ICL to MSA, we conduct an in-depth investigation into three pivotal factors that influence the configuration of ICL demonstrations: similarity measurement, modality presentation, and sentiment distribution. For these factors, we optimize strategies by balancing the similarity components, weighing the richness of information against the complexity of inputs, and offsetting the sentimental predictive bias of MLLMs. Comprehensive experiments on six datasets and two MLLMs demonstrate the pronounced superiority of our strategies against other ICL baselines, validating their effectiveness and generalizability. Our findings confirm that MLLMs can perceive sentiment as competent as supervised models, paving the way for further research and exploration.

A primary limitation of this paper lies in the range of MLLMs investigated. The effectiveness of ICL heavily depends on the MLLMs themselves, however, some advanced MLLMs are beyond the research scope due to practical reasons. Despite our progress, multimodal ICL is still in its infancy compared to ICL on text modality. Other aspects of configuring ICL demonstrations also merit further investigation.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant NO 62376266, 62406318 and 62441614), Key Laboratory of Ethnic Language Intelligent Analysis and Security Governance of MOE, Minzu University of China, Beijing, China.

Impact Statement

This manuscript focuses on configuring ICL demonstrations to unleash MLLMs' sentimental perception capability. By investigating three factors, we aim to advance multimodal ICL in MLLMs and shed light on MLLMs' sentimentrelated properties. During the investigation, we discover a sentimental predictive bias, later mitigating it on the inference level to facilitate fairness in MLLMs' sentiment prediction. We have not delved deeper into the source limitations of this bias, as it lies beyond the primary scope of our research.

However, systematic studies of these limitations can potentially contribute to both MLLMs and MSA. As a further discussion, we attribute these limitations to pretraining data rather than model architecture, which is validated by Xie et al. (2024). By constructing emotion-related data, it enhances MLLMs' zero-shot performance on visual emotion recognition. This success has the potential to be replicated in solving the sentimental predictive bias.

References

Alayrac, J., Donahue, J., Luc, P., Miech, A., Barr, I., and et. al. Flamingo: a visual language model for few-shot learning. In *NeurIPS 2022*, 2022.

- Alimisis, P., Mademlis, I., Radoglou-Grammatikis, P., Sarigiannidis, P., and Papadopoulos, G. T. Advances in diffusion models for image data augmentation: A review of methods, models, evaluation metrics and future research directions, 2024.
- Awadalla, A., Gao, I., Gardner, J., Hessel, J., Hanafy, Y., and et. al. Openflamingo: An open-source framework for training large autoregressive vision-language models, 2023.
- Baldassini, F. B., Shukor, M., Cord, M., Soulier, L., and Piwowarski, B. What makes multimodal in-context learning work? In *CVPR 2024 - Workshops*, pp. 1539–1550, 2024.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., and et. al. Language models are few-shot learners. In *NeurIPS 2020*, 2020.
- Cai, C., Wang, Q., Liang, B., Qin, B., Yang, M., Wong, K., and Xu, R. In-context learning for few-shot multimodal named entity recognition. In *Findings of EMNLP 2023*, pp. 2969–2979, 2023.
- Chen, H., Yang, X., Huang, Y., Wu, Z., Wang, J., and Geng, X. Manipulating the label space for in-context classification, 2023a.
- Chen, S., Han, Z., He, B., Buckley, M., Torr, P., Tresp, V., and Gu, J. Understanding and improving in-context learning on vision-language models. In *ICLR 2024 Workshop*, 2024.
- Chen, Y., Yuan, J., Tian, Y., Geng, S., Li, X., Zhou, D., Metaxas, D. N., and Yang, H. Revisiting multimodal representation in contrastive learning: From patch and token embeddings to finite discrete tokens. In *CVPR* 2023, pp. 15095–15104, 2023b.
- Dong, Q., Li, L., Dai, D., Zheng, C., Ma, J., Li, R., Xia, H., Xu, J., Wu, Z., Chang, B., Sun, X., Li, L., and Sui, Z. A survey on in-context learning, 2024.
- Driess, D., Xia, F., Sajjadi, M. S. M., Lynch, C., Chowdhery, A., and et. al. Palm-e: An embodied multimodal language model. In *ICML 2023*, volume 202 of *Proceedings of Machine Learning Research*, pp. 8469–8488, 2023.
- Huang, S., Dong, L., Wang, W., Hao, Y., Singhal, S., and et. al. Language is not all you need: Aligning perception with language models. In *NeurIPS 2023*, 2023.
- Huang, Y., Du, C., Xue, Z., Chen, X., Zhao, H., and Huang, L. What makes multi-modal learning better than single (provably). In *NeurIPS 2021*, pp. 10944–10956, 2021.

- Kosti, R., Álvarez, J. M., Recasens, A., and Lapedriza, À. Emotion recognition in context. In *CVPR 2017*, pp. 1960– 1968, 2017.
- Laurençon, H., Saulnier, L., Tronchon, L., Bekman, S., Singh, A., and et. al. OBELICS: an open web-scale filtered dataset of interleaved image-text documents. In *NeurIPS 2023*, 2023.
- Levy, I., Bogin, B., and Berant, J. Diverse demonstrations improve in-context compositional generalization. In ACL 2023, pp. 1401–1422, 2023.
- Li, J., Li, D., Savarese, S., and Hoi, S. C. H. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML* 2023, volume 202 of *Proceedings of Machine Learning Research*, pp. 19730–19742, 2023.
- Li, L., Peng, J., Chen, H., Gao, C., and Yang, X. How to configure good in-context sequence for visual question answering. In *CVPR 2024*, pp. 26710–26720, 2024.
- Li, X. and Qiu, X. Finding support examples for in-context learning. In *Findings of EMNLP 2023*, pp. 6219–6235, 2023.
- Li, Z., Xu, B., Zhu, C., and Zhao, T. CLMLF: A contrastive learning and multi-layer fusion method for multimodal sentiment detection. In *Findings of NAACL 2022*, pp. 2282–2294, 2022.
- Lian, Z., Sun, L., Sun, H., Chen, K., Wen, Z., Gu, H., Liu, B., and Tao, J. GPT-4V with emotion: A zeroshot benchmark for generalized emotion recognition. *Inf. Fusion*, 108:102367, 2024.
- Ling, Y., Yu, J., and Xia, R. Vision-language pre-training for multimodal aspect-based sentiment analysis. In ACL 2022, pp. 2149–2159, 2022.
- Liu, J., Shen, D., Zhang, Y., Dolan, B., Carin, L., and Chen, W. What makes good in-context examples for gpt-3? In Agirre, E., Apidianaki, M., and Vulic, I. (eds.), *DeeLIO@ACL 2022*, pp. 100–114, 2022.
- Lu, D., Neves, L., Carvalho, V., Zhang, N., and Ji, H. Visual attention model for name tagging in multimodal social media. In ACL 2018, pp. 1990–1999, 2018.
- Lu, Y., Bartolo, M., Moore, A., Riedel, S., and Stenetorp, P. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In ACL 2022, pp. 8086–8098, 2022.
- Lyu, X., Min, S., Beltagy, I., Zettlemoyer, L., and Hajishirzi, H. Z-ICL: zero-shot in-context learning with pseudodemonstrations. In Rogers, A., Boyd-Graber, J. L., and Okazaki, N. (eds.), ACL 2023, pp. 2304–2317, 2023.

- Niu, T., Zhu, S., Pang, L., and El-Saddik, A. Sentiment analysis on multi-view social data. In *MMM 2016*, volume 9517 of *Lecture Notes in Computer Science*, pp. 15–27, 2016.
- Pan, J., Gao, T., Chen, H., and Chen, D. What in-context learning "learns" in-context: Disentangling task recognition and task learning. In *Findings of ACL 2023*, pp. 8298–8319, 2023.
- Peng, T., Li, Z., Wang, P., Zhang, L., and Zhao, H. A novel energy based model mechanism for multi-modal aspectbased sentiment analysis. In AAAI 2024, pp. 18869– 18878, 2024.
- Qin, L., Chen, Q., Fei, H., Chen, Z., Li, M., and Che, W. What factors affect multi-modal in-context learning? an in-depth exploration. In *NeurIPS 2024*, 2024.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., and et. al. Learning transferable visual models from natural language supervision. In *ICML 2021*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763, 2021.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *CVPR 2022*, pp. 10674–10685, 2022.
- Shukor, M., Rame, A., Dancette, C., and Cord, M. Beyond task performance: evaluating and reducing the flaws of large multimodal models with in-context-learning. In *ICLR 2024*, 2024.
- Wang, Y., Chen, W., Han, X., Lin, X., Zhao, H., Liu, Y., Zhai, B., Yuan, J., You, Q., and Yang, H. Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning, 2024.
- Wei, Y., Yuan, S., Yang, R., Shen, L., Li, Z., Wang, L., and Chen, M. Tackling modality heterogeneity with multiview calibration network for multimodal sentiment detection. In ACL 2023, pp. 5240–5252, 2023.
- Wu, D., Yang, D., Zhou, Y., and Ma, C. Robust multimodal sentiment analysis of image-text pairs by distributionbased feature recovery and fusion. In ACM MM 2024, pp. 5780–5789, 2024a.
- Wu, D., Yang, D., Zhou, Y., and Ma, C. Bridging visual affective gap: Borrowing textual knowledge by learning from noisy image-text pairs. In ACM MM 2024, pp. 602– 611, 2024b.

- Wu, Z., Wang, Y., Ye, J., and Kong, L. Self-adaptive incontext learning: An information compression perspective for in-context example selection and ordering. In *ACL 2023*, pp. 1423–1436, 2023.
- Xie, H., Peng, C., Tseng, Y., Chen, H., Hsu, C., Shuai, H., and Cheng, W. Emovit: Revolutionizing emotion insights with visual instruction tuning. In *CVPR 2024*, pp. 26586–26595, 2024.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A. C., Salakhutdinov, R., Zemel, R. S., and Bengio, Y. Show, attend and tell: Neural image caption generation with visual attention. In *ICML 2015*, volume 37, pp. 2048–2057, 2015.
- Xu, N. Analyzing multimodal public sentiment based on hierarchical semantic attentional network. In *ISI 2017*, pp. 152–154, 2017.
- Xu, N. and Mao, W. Multisentinet: A deep semantic network for multimodal sentiment analysis. In *CIKM 2017*, pp. 2399–2402, 2017.
- Xu, N., Mao, W., and Chen, G. A co-memory network for multimodal sentiment analysis. In *SIGIR 2018*, pp. 929–932, 2018.
- Xu, N., Mao, W., and Chen, G. Multi-interactive memory network for aspect based multimodal sentiment analysis. In AAAI 2019, pp. 371–378, 2019.
- Yang, H., Zhao, Y., Wu, Y., Wang, S., Zheng, T., Zhang, H., Che, W., and Qin, B. Large language models meet text-centric multimodal sentiment analysis: A survey, 2024.
- Yang, X., Feng, S., Wang, D., and Zhang, Y. Image-text multimodal emotion classification via multi-view attentional network. *IEEE Trans. Multim.*, 23:4014–4026, 2021a.
- Yang, X., Feng, S., Zhang, Y., and Wang, D. Multimodal sentiment detection based on multi-channel graph neural networks. In ACL/IJCNLP 2021, pp. 328–339, 2021b.
- Yang, X., Feng, S., Wang, D., Zhang, Y., and Poria, S. Fewshot multimodal sentiment analysis based on multimodal probabilistic fusion prompts. In ACM MM 2023, pp. 6045–6053, 2023a.
- Yang, X., Wu, W., Feng, S., Wang, M., Wang, D., Li, Y., Sun, Q., Zhang, Y., Fu, X., and Poria, S. Mm-bigbench: Evaluating multimodal models on multimodal content comprehension tasks, 2023b.
- Yang, X., Wu, Y., Yang, M., Chen, H., and Geng, X. Exploring diverse in-context configurations for image captioning. In *NeurIPS 2023*, 2023c.

- Yang, Z., Gan, Z., Wang, J., Hu, X., Lu, Y., Liu, Z., and Wang, L. An empirical study of GPT-3 for few-shot knowledge-based VQA. In AAAI 2022, pp. 3081–3089, 2022.
- Yin, S., Fu, C., Zhao, S., Li, K., Sun, X., Xu, T., and Chen, E. A survey on multimodal large language models, 2023.
- Yu, J. and Jiang, J. Adapting BERT for target-oriented multimodal sentiment classification. In Kraus, S. (ed.), *IJCAI 2019*, pp. 5408–5414, 2019.
- Yu, Y., Zhang, D., and Li, S. Unified multi-modal pretraining for few-shot sentiment analysis with promptbased learning. In ACM MM 2022, pp. 189–198, 2022.
- Yue, L., Chen, W., Li, X., Zuo, W., and Yin, M. A survey of sentiment analysis in social media. *Knowl. Inf. Syst.*, 60(2):617–663, 2019.
- Zadeh, A., Chen, M., Poria, S., Cambria, E., and Morency, L. Tensor fusion network for multimodal sentiment analysis. In *EMNLP 2017*, pp. 1103–1114, 2017.
- Zhang, L., Wang, S., and Liu, B. Deep learning for sentiment analysis: A survey. WIREs Data Mining Knowl. Discov., 8(4), 2018a.
- Zhang, Q., Fu, J., Liu, X., and Huang, X. Adaptive coattention network for named entity recognition in tweets. In AAAI 2018, pp. 5674–5681, 2018b.
- Zhang, S., Pan, Y., and Wang, J. Z. Learning emotion representations from verbal and nonverbal communication. In *CVPR 2023*, pp. 18993–19004, 2023.
- Zhang, Y., Feng, S., and Tan, C. Active example selection for in-context learning. In *EMNLP 2022*, pp. 9134–9148, 2022.
- Zhao, S., Yao, X., Yang, J., Jia, G., Ding, G., Chua, T., Schuller, B. W., and Keutzer, K. Affective image content analysis: Two decades review and new perspectives. *Trans. Pattern Anal. Mach. Intell.*, 44(10):6729–6751, 2022.
- Zhou, J., Zhao, J., Huang, J. X., Hu, Q. V., and He, L. MASAD: A large-scale dataset for multimodal aspectbased sentiment analysis. *Neurocomputing*, 455:47–58, 2021.
- Zhou, R., Guo, W., Liu, X., Yu, S., Zhang, Y., and Yuan, X. Aom: Detecting aspect-oriented information for multimodal aspect-based sentiment analysis. In *Findings of ACL 2023*, pp. 8184–8196, 2023.
- Zhu, X., Guo, C., Feng, H., Huang, Y., Feng, Y., Wang, X., and Wang, R. A review of key technologies for emotion analysis using multimodal information. *Cogn. Comput.*, 16(4):1504–1530, 2024.

A. Textual Prompts

The textual prompt \mathcal{P} of multimodal ICL sequence \mathcal{S} aims to provide the task description and the set of target categories.

A.1. Aspect-Level MSA

 \mathcal{P} is "A post contains an image, a text and an aspect. Identify the sentiment of the aspect in the post. The optional categories are [Positive, Neutral, and Negative]. Here are some examples".

A.2. Post-Level MSA

 \mathcal{P} is "A post contains an image and a text. Classify the sentiment of the post into [Positive, Neutral, Negative]. Here are some examples".

A.3. Sensitivity of ICL to Prompt Variations

In the investigation, we experiment with various textual prompts and find that they significantly impact zero-shot performance. However, their impact on ICL is minimal. Since this manuscript primarily focuses on how ICL configurations influence MLLMs' sentiment perception capabilities, we select a set of appropriate textual prompts and keep them fixed throughout the investigation. The performance of IDEFICS under different prompts is reported in Table 6.

For post-level MSA:

#1 Prompt: A post contains an image and a text. Classify the sentiment of the post into [Positive, Neutral, Negative].

#2 Prompt: Please classify the sentiment of the image-text post into [Positive, Neutral, Negative].

#3 Prompt: Here is a post containing an image and a text. The optional categories are [Positive, Neutral, Negative]. What is the overall sentiment of the post?

For aspect-level MSA:

#1 Prompt: A post contains an image, a text and an aspect. Identify the sentiment of the aspect in the post. The optional categories are [Positive, Neutral, Negative].

#2 Prompt: Please classify the sentiment of the aspect in image-text post into [Positive, Neutral, Negative].

#3 Prompt: Here is a post containing an image, a text and an aspect. The optional categories are [Positive, Neutral, Negative]. What is the sentiment of the aspect in the post?

Model & Strategy		Support Set		MVSA-S		Twitter-15			
		~~FF	#1 Prompt	#2 Prompt #3 Prompt		#1 Prompt #2 Prompt		#3 Prompt	
IDEFICS	Zero-Shot Paradigm ICL Ours 16-shot	- 1% Training Set	38.6 66.5	28.2 66.3	50.6 66.4	60.7 67.0	51.9 66.9	19.1 66.7	

Table 6. Influence of prompt variations on accuracy of IDEFICS.

B. Dataset Details

The statistics of the adopted datasets are presented in Table 7.

MVSA-S, MVSA-M (Niu et al., 2016) are labeled on single modalities, where the sentiment categories include *Positive*, *Neutral* and *Negative*. The multimodal sentiment categories are obtained by majority voting following Xu & Mao (2017).

TumEmo (Yang et al., 2021a) is a weakly supervised dataset. Image-text posts are retrieved based on seven emotion keywords: *Love, Happy, Calm, Bored, Sad, Angry, Fear*, and labeled accordingly.

Twitter-15 (Zhang et al., 2018b), Twitter-17 (Lu et al., 2018) are initially proposed for Multimodal Named Entity Recognition. Their named entities are later annotated by Yu & Jiang (2019) based on the sentiment polarities: *Positive, Neutral, Negative*, and utilized for aspect-level MSA.

MASAD (Zhou et al., 2021) extends textual aspects to visual aspects, and includes posts from more diverse domains. The aspects are labeled by sentiment polarities: *Positive* and *Negative*.

	Dataset	Train	Test
	MVSA-S (Niu et al., 2016)	3608	452
Post-Level	MVSA-M (Niu et al., 2016)	13618	1703
	TumEmo (Yang et al., 2021a)	156217	19524
	Twitter-15 (Zhang et al., 2018b)	3179	1037
Aspect-Level	Twitter-15 (Lu et al., 2018)	3562	1234
	MASAD (Zhou et al., 2021)	14868	4935

Table 7.	Statistics	of datasets.
----------	------------	--------------

C. Computational Overheads

In the optimized configuration, presenting and distributing demonstrations do not introduce additional computational overhead. The extra costs originate from demonstration retrieval and the expanded input sequence for MLLMs. The former scales with the size of the support set, as each test sample needs to be compared against all support set samples, while the latter is inherent to ICL. We report the average time overhead (ms) of processing an image-text sample under two support set scales on a single NVIDIA GeForce RTX 4090 GPU.

Table 6. Comparison of time costs.											
Mo	del & Strategy	Support Set	Time	Time Overhead (ms)							
			Retrieval	Inference	Total						
	Zero-Shot Paradigm	-	0	78.1	78.1						
	ICL Random 4-shot	136 / 1562 Samples	0	134.5	134.5						
IDEFICE	ICL Ours 4-shot	136 Samples	36.4	134.5	170.9						
IDEFICS	ICL Ours 4-shot	1562 Samples	64.2	134.5	198.7						
	ICL Random 16-shot	136 / 1562 Samples	0	346.1	346.1						
	ICL Ours 16-shot	136 Samples	36.4	346.1	382.5						
	ICL Ours 16-shot	1562 Samples	64.2	346.1	410.3						

Table 8. Comparison of time costs

D. Complete Results

In the main paper, we simplify the reported results to emphasize the key findings. Here we present the complete results in numerical form.

Table 2 reports the average accuracy across 4,8,16-shot demonstrations retrieved based on varying similarity measurements. Table 9 is its complete version.

Figure 4 reports the average accuracy across 4,8,16-shot demonstrations retrieved based on the **WIT** and **WITA** strategies. Table 10 is its complete version.

Table 3 reports the average accuracy across 4,8,16-shot settings with the inputs composed of different modalities. Table 11 is its complete version.

Figure 5 evaluates ICL's "Task Learning" effect by progressively incorporating modalities into the inputs. Table 12 is its complete version.

		4.4	shot			8	shot		16 shot			
Strategy	Post-	-Level	Aspec	t-Level	Post-	Post-Level		Aspect-Level		Level	Aspect-Level	
	MVSA-S	MVSA-M	Twitter-15	Twitter-17	MVSA-S	MVSA-M	Twitter-15	Twitter-17	MVSA-S	MVSA-M	Twitter-15	Twitter-17
R	45.2	59.7	56.1	55.2	50.8	61.2	58.8	56.9	51.4	61.6	57.3	57.1
Ι	51.2	64.3	56.5	56.8	59.2	65.2	59.2	56.7	59.2	65.3	61.6	56.6
Т	49.6	64.6	55.0	55.2	58.3	67.2	58.6	56.8	60.0	66.8	62.5	58.9
IT	49.5	64.8	60.5	58.7	57.7	66.3	62.2	56.2	60.1	67.3	61.6	57.9
А	-	-	60.2	57.0	-	-	60.5	57.6	-	-	63.4	57.7
IA	-	-	57.1	58.2	-	-	59.5	57.8	-	-	62.0	57.9
TA	-	-	58.7	56.6	-	-	61.6	57.3	-	-	62.4	58.0
ITA	-	-	59.3	58.3	-	-	61.6	57.5	-	-	62.2	58.5

Table 9. Complete results of Table 2.

Table 10. Complete results of Figure 4.

WIT Strategy	4-:	shot	8-5	shot	16-	shot	WITA Strategy	4-s	hot	8-s	8-shot		shot
$\boldsymbol{lpha}: \boldsymbol{eta}$	MVSA-S	MVSA-M	MVSA-S	MVSA-M	MVSA-S	MVSA-M	$\boldsymbol{\alpha}:\boldsymbol{\beta};(\boldsymbol{\alpha}+\boldsymbol{\beta}):\boldsymbol{\gamma}$	Twitter-15	Twitter-17	Twitter-15	Twitter-17	Twitter-15	Twitter-17
0:10	49.6	64.6	58.3	67.2	60.0	66.8	1:9; 8:2	56.7	56.9	59.0	57.4	62.8	59.0
1:9	51.7	65.5	58.5	67.0	61.4	66.3	1:9; 5:5	58.6	56.4	62.5	57.5	63.2	58.5
2:8	49.0	65.3	57.2	67.1	66.5	67.7	1:9; 2:8	58.9	58.2	60.8	56.5	63.0	57.9
3:7	50.2	65.4	57.2	67.7	64.9	66.9	3:7; 8:2	56.6	55.8	60.5	55.9	62.9	59.1
4:6	49.0	64.9	57.7	67.2	62.7	67.2	3:7; 5:5	59.7	58.2	61.7	57.3	62.8	58.4
5:5	49.5	64.4	57.5	66.6	60.1	67.3	3:7; 2:8	60.2	58.6	59.8	57.4	63.2	57.8
6:4	49.5	64.4	58.1	67.2	59.4	66.7	5:5; 8:2	57.8	57.1	62.5	56.5	64.0	58.1
7:3	51.2	64.2	57.7	66.8	59.2	66.9	5:5; 5:5	58.8	57.9	60.2	56.8	62.0	58.2
8:2	51.2	64.6	58.8	67.0	57.7	66.7	5:5; 2:8	59.8	59.7	60.2	57.5	63.7	58.7
9:1	50.8	64.8	57.9	65.6	57.2	66.2	7:3; 8:2	55.7	57.2	58.3	58.2	62.4	58.6
10:0	51.2	64.3	59.2	65.2	59.2	65.3	7:3; 5:5	58.4	57.5	59.5	56.4	62.1	57.9
-	-	-	-	-	-	-	7:3; 2:8	60.5	58.9	60.5	59.3	64.4	58.7
-	-	-	-	-	-	-	9:1; 8:2	55.0	58.0	59.7	57.7	63.0	58.5
-	-	-	-	-	-	-	9:1; 5:5	56.8	58.3	59.4	57.1	61.9	58.8
-	-	-	-	-	-	-	9:1; 2:8	60.5	58.8	60.6	57.9	63.8	59.0

Table 11. Complete results of Table 3.

		4-9	shot			8-shot				16-shot			
Modality	Post-	Level	Aspect	-Level	Post-	Level	Aspect	-Level	Post-	Level	Aspect	-Level	
	MVSA-S	MVSA-M	Twitter-15	Twitter-17	MVSA-S	MVSA-M	Twitter-15	Twitter-17	MVSA-S	MVSA-M	Twitter-15	Twitter-17	
Ι	33.5	45.7	55.2	51.7	55.4	57.8	58.5	53.8	66.3	66.5	58.4	54.7	
С	26.2	35.1	53.3	49.8	47.2	52.3	56.8	54.1	58.3	61.1	58.5	52.4	
I, C	48.8	54.9	52.3	50.1	60.1	60.9	53.4	50.1	61.9	62.0	57.1	50.1	
Т	29.7	47.6	60.8	58.9	48.6	56.3	61.5	59.1	60.8	64.2	62.9	59.0	
G	31.7	44.0	54.3	54.6	48.3	57.8	54.9	55.4	60.3	64.9	55.2	55.7	
T, G	34.2	47.5	49.0	54.1	50.6	61.0	55.4	55.0	58.3	65.7	57.6	55.4	
I, T	49.5	64.8	60.5	58.7	57.7	66.3	62.2	56.2	60.1	67.3	61.6	57.9	
I, G	47.3	54.2	53.1	51.8	56.4	58.9	55.0	53.4	59.8	63.3	55.4	56.2	
С, Т	32.2	54.5	61.5	56.4	55.7	61.8	63.9	57.3	61.2	64.6	62.4	56.3	
C, G	33.7	41.0	51.6	51.5	49.9	56.8	52.5	54.3	59.0	63.2	55.2	53.5	
I, C, T	47.7	63.3	56.5	54.9	57.9	64.0	61.0	56.2	60.3	65.8	63.5	55.8	
I, T, G	45.0	58.6	58.0	53.9	53.9	62.4	60.3	55.2	53.4	64.8	61.0	55.2	
C, T, G	36.1	50.5	51.0	53.0	48.6	59.9	57.8	55.4	57.7	64.2	60.9	54.0	
I, C, G	41.0	52.3	52.8	51.1	54.1	57.5	51.1	52.4	55.9	61.2	54.7	52.9	
I, C, T, G	41.7	59.4	54.7	52.4	51.2	61.4	59.4	53.6	54.8	63.5	62.2	53.4	

Table 12. Complete results of Figure 5.

Modality	4-shot		8-shot		16-shot	
	Twitter-15	Twitter-17	Twitter-15	Twitter-17	Twitter-15	Twitter-17
Т	33.7	43.0	39.6	44.9	52.7	48.5
+I (I, T)	30.4	43.5	36.7	44.7	48.2	46.7
+C (I, C, T)	30.3	43.7	33.4	43.3	49.2	45.7
$+G\left(I,C,T,G\right)$	23.2	39.7	35.8	44.4	49.9	45.8