

# 000 CH-CEMS: A CHINESE MULTI-CONCEPT BENCHMARK 001 DATASET TOWARDS EXPLAINABLE MULTI-MODAL SENTIMENT 002 ANALYSIS

003  
004 **Anonymous authors**  
005

006 Paper under double-blind review  
007  
008  
009  
010  
011  
012

## ABSTRACT

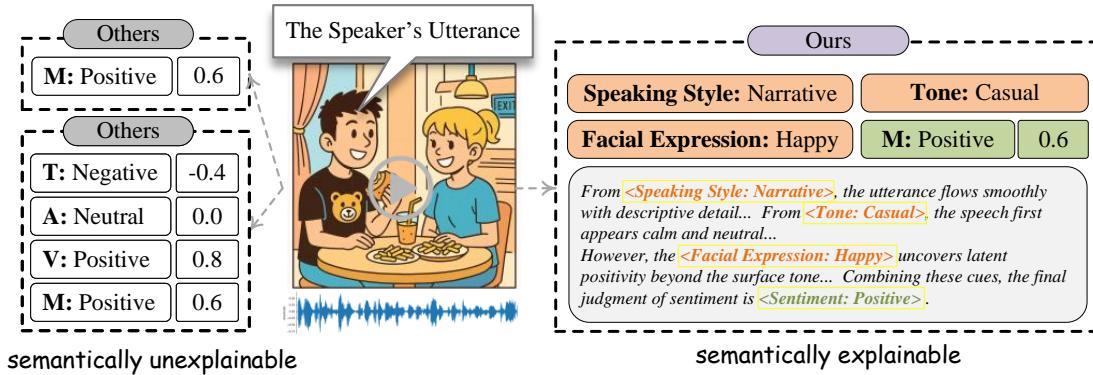
013 Explainable Multimodal Sentiment Analysis (EMSA) is a booming research area aimed at advanc-  
014 ing robust and faithful multimodal language understanding. Recent explainable datasets and meth-  
015 ods based on multimodal large language models (MLLMs) have introduced a new paradigm that  
016 produces chain-of-thought-style explanations within affective computing. However, high-quality  
017 data resources for EMSA remain scarce, largely because annotating reliable reasoning cues is  
018 costly and difficult. To address this gap, we introduce CH-CEMS, the first multimodal senti-  
019 ment dataset for explainable multimodal sentiment analysis. It contains 3,715 curated video seg-  
020 ments with polarity and intensity annotations. In addition, we annotate three semantic concepts  
021 for each sample (i.e., speaking style, tone of voice, and facial expression), which serve as ex-  
022 plicit reasoning cues to enable process-level supervision. To fully leverage these concept cues,  
023 we propose a concept-guided reinforcement learning framework with Group Relative Policy Opti-  
024 mization (GRPO) for MLLMs, in which concept-level supervision explicitly constrains cross-  
025 modal semantic relations and guides the model to infer sentiment from verifiable concepts. We  
026 further establish baselines with state-of-the-art multimodal machine learning methods and MLLMs  
027 via zero-shot inference and supervised fine-tuning. Experiments show that MLLMs outperform  
028 feature-based methods, typically by 4–12% in accuracy for three-class sentiment analysis, and  
029 that our concept-guided GRPO yields a further 8.5% improvement, even surpassing closed-source  
030 model such as GPT-5. We believe CH-CEMS and the benchmark will facilitate future research  
031 on explainable multimodal sentiment analysis. The dataset and codes are available for use at  
032 <https://anonymous.4open.science/r/CH-CEMS-C34F>.

## 033 1 INTRODUCTION

034 Multimodal Sentiment Analysis (MSA), which integrates text, speech, and vision to overcome the limitations of uni-  
035 modal text by leveraging complementary multimodal information for more robust sentiment understanding, holds sig-  
036 nificant importance within the multimodal language understanding. The progress of MSA has been largely driven by  
037 benchmark datasets such as CMU-MOSI (Zadeh et al., 2016), and CMU-MOSEI (Zadeh et al., 2018), which provide  
038 essential resources for the training and systematic evaluation of multimodal models. Furthermore, in order to achieve  
039 friendly representations for multimodal fusion, datasets with modality-specific sentiment annotations are proposed  
040 such as CH-SIMS (Yu et al., 2020) and CH-SIMS2.0 (Yu et al., 2023). Building upon these datasets, prior studies  
041 have mainly emphasized intra-modal representation learning and inter-modal fusion. Though these approaches have  
042 achieved strong performance, they remain confined to feature-level modeling and have limited capacity to capturing  
043 high-level and diverse multimodal semantics, making it difficult to reason over explainable and verifiable cues under  
044 traditional modeling frameworks and task formulations. In real-world scenarios, sentiment polarities across modalities  
045 often diverge, and multimodal sentiment cannot be reduced to a simple aggregation of unimodal predictions. Existing  
046 methods typically address such modality inconsistency through feature weighting or selection (Li et al., 2024; 2025;  
047 Zhao et al., 2025b), but these strategies are still heuristic adjustments that lack semantically grounded modeling of  
048 inter-modal relations and consistency. To move beyond such limitations, a promising direction is to construct multi-  
049 modal datasets enriched with explainable and verifiable semantic cues, and to leverage these cues as context-dependent  
050 references for reasoning.

051 Recently, multimodal large language models (MLLMs) have been applied to MSA (Jim et al., 2024; Yang et al.,  
052 2024b; Zhang et al., 2025a). A prevailing paradigm formulates the task in an autoregressive manner with MLLMs,  
053 where multimodal inputs are implicitly fused through alignment or attention in latent feature spaces, and optimized  
054 via supervised fine-tuning (Luo et al., 2025; Zhang et al., 2025b). Nevertheless, such approaches rarely impose se-

054 mantic constraints or consistency mechanisms, leaving the reasoning process untraceable. Large-scale MLLMs have  
 055 demonstrated the capability to extract high-level semantic cues across modalities and underpin reliable multimodal  
 056 reasoning (Wang et al., 2025b). While recent multi-stage or agent-style pipelines somewhat instruct models to un-  
 057 derstand such concepts to help reason (Huang et al., 2024; Fei et al., 2023; Zhang et al., 2025c), explicit training of  
 058 the reasoning capability is underexplored. Meanwhile, following the advent of DeepSeek R1 (DeepSeek-AI, 2025),  
 059 reinforcement learning (RL)-based training, exemplified by Group Relative Policy Optimization (GRPO) (Shao et al.,  
 060 2024), has been shown to improve explanatory behavior in large language models and substantially enhance semantic-  
 061 level reasoning in multimodal settings (Wang et al., 2025b). In parallel, researchers have also made effective attempts  
 062 toward more explainable affective computing. The EMER dataset provides explainable reasoning data for emotion  
 063 recognition (Liu et al., 2023a), and the MERR dataset extends this line by offering multimodal emotion descrip-  
 064 tion and benchmarks (Wu et al., 2023). Building on these resources, Emotion-LLaMA (Zhang et al., 2024) adopts  
 065 multimodal instruction tuning to supervise explanation generation, demonstrating that explicit reasoning supervision  
 066 can significantly improve both recognition accuracy and explainability. Furthermore, R1-Omni (Zhao et al., 2025a)  
 067 optimizes an omni-modal model via reinforcement learning with verifiable rewards (RLVR) and GRPO, thereby sub-  
 068 stantially improving the model’s reasoning capability and explainability. However, most explainable MSA approaches  
 069 instantiate explanations as chain-of-thought (CoT) rationales, whose faithfulness is debated and which are generally  
 070 unsupervised at the process level, leaving hallucinated reasoning uncorrected. Moreover, to the best of our knowledge,  
 071 explainable multimodal datasets for sentiment analysis with process-level supervision and verifiable reasoning remain  
 072 scarce.



086 Figure 1: Comparison between conventional datasets and ours. Others provide only modality-level scores (semanti-  
 087 cally unexplainable), while ours adds concept-level annotations and reasoning chains, making sentiment judgments  
 088 semantically explainable.

091 To advance research on explainable multimodal sentiment analysis and enable supervision of the reasoning process, we  
 092 introduce the first Chinese multimodal dataset, CH-CEMS, with multi-concept annotations for explainable multimodal  
 093 sentiment analysis. Prior datasets provide only sentiment polarity and intensity annotations of multi- or uni-modalities,  
 094 without any reasoning cues. To fill in this blank, we further provide annotations of three semantic concepts (speaking  
 095 style, tone of voice, and facial expression) and two-stage, post-processed chain-of-thought (CoT) traces as the rea-  
 096 soning process. A comparison with prior datasets is shown in Figure 1. Building on this, we propose an explainable  
 097 multimodal sentiment analysis framework for MLLMs trained with reinforcement learning. By incorporating concept-  
 098 level supervision as verifiable rewards, our method explicitly models multimodal semantic relations, guiding the model  
 099 to reason about speakers’ sentiment through verifiable concepts and their interactions, and to produce final sentiment  
 100 judgments grounded in traceable multimodal evidence. We further establish benchmark for both a regression task using  
 101 five state-of-the-art, feature-based multimodal methods and a generative classification task with mainstream MLLMs.  
 102 Specifically, we benchmark three closed-source and four open-source MLLMs under zero-shot inference and super-  
 103 vised fine-tuning (SFT) settings. The results reveal a substantial accuracy gap of approximately 2-12% for three-class  
 104 sentiment analysis between feature-based methods and MLLMs under the generative classification paradigm in the  
 105 zero-shot setting. While conventional post-training methods improve MLLMs’ performance, our approach not only  
 106 achieves state-of-the-art results on CH-CEMS, surpassing strong closed-source MLLMs such as GPT-5 and Gemini  
 107 2.5-Pro, but also produces sentiment predictions with an explainable reasoning process, demonstrating that concept-  
 108 level supervision within this framework facilitates more faithful and explainable reasoning trajectories and enhances  
 109 overall performance.

**Contributions.** (1) This paper presents CH-CEMS, the first Chinese dataset for explainable multimodal sentiment analysis. It contains 3,715 high-quality curated samples with not only sentiment polarity and intensity annotations but also reasoning process and cues to facilitate high-level semantics understanding. (2) We develop a reinforcement learning framework for MLLMs that implements concept-level supervision as verifiable rewards, the first successful attempt to leverage reasoning cues for multimodal sentiment analysis. (3) We build a comprehensive benchmark on CH-CEMS. Extensive experiments show that our method achieves state-of-the-art results with substantial improvements over state-of-the-art MLLMs and GPT-5, demonstrating the effectiveness of reasoning concepts for analyzing multimodal sentiment. The CH-CEMS dataset provides a new resource in this area, offering a solid basis in further research.

## 2 RELATED WORK

### 2.1 BENCHMARK DATASETS

**Multimodal Sentiment Analysis Benchmark Dataset.** Early work on multimodal sentiment analysis (MSA) was enabled by video–text–audio benchmark datasets such as CMU-MOSI (Zadeh et al., 2016), CMU-MOSEI (Zadeh et al., 2018), and ICT-MMMO (Wöllmer et al., 2013). These corpora catalyzed research on unimodal representation learning and multimodal fusion, and they remain widely used for training and evaluation. Recent resources enrich supervision beyond utterance-level sentiment. For example, CH-SIMS (Yu et al., 2020) provides independent unimodal annotations in addition to multimodal sentiment (Yu et al., 2020), and CH-SIMS v2.0 (Yu et al., 2023) expands the corpus with a greater focus on non-verbal cues. CMU-MOSEAS extends MSA to multiple non-English languages and incorporates additional emotion and attribute labels (Zadeh et al., 2020). These datasets collectively broaden the scope of MSA and provide finer-grained supervision. Nevertheless, their additional annotations remain largely at the feature or physical signal level, without modeling higher-level semantic concepts that can serve as explicit reasoning clues.

**Explainable Affective Computing Benchmark Datasets.** Recently, several datasets have advanced explainable affective computing. EMER curates multimodal emotion–reasoning pairs (Liu et al., 2023a) and MERR provides a multimodal emotion description and reasoning benchmark (Zhang et al., 2024). PanoSent introduces a multimodal conversational ABSA benchmark for panoptic sextuple extraction and sentiment flipping with causal rationales (Luo et al., 2024). However, to the best of our knowledge, these resources primarily target emotion recognition or aspect-based sentiment analysis rather than user-centric sentiment reasoning. To fill this gap, we present *CH-CEMS*, a concept-centric, reasoning-ready dataset tailored for explainable multimodal sentiment analysis. A comparison between CH-CEMS and other benchmark MSA datasets is presented in Table 1.

### 2.2 MULTIMODAL SENTIMENT ANALYSIS

Multimodal sentiment analysis has advanced through feature-based methods and pre-trained model approaches. These methods can be broadly categorized into early fusion, late fusion, and end-to-end sequence modeling within conventional architectures, as well as methods leveraging pre-trained and large-scale models for more complex MSA tasks.

**Feature-based Methods.** Feature-based methods typically rely on early and late fusion strategies. Early fusion methods, such as MISA (Hazarika et al., 2020), combine modality-specific and modality-invariant features to enhance multimodal affective state prediction. Building on this, DLF (Wang et al., 2025a) introduces a disentangled language-focused fusion framework to improve language-targeted feature extraction. Similarly, MAG-BERT (Rahman et al., 2020a) applies a Multimodal Adaptation Gate to BERT, improving multimodal fusion at the feature level. In contrast, late fusion methods independently predict sentiment for each modality before combining them at the decision level, enhancing robustness. For example, MCIS (Yang et al., 2024a) utilizes counterfactual reasoning to mitigate biases and make robust predictions without additional training. Models such as MuLT (Tsai et al., 2019) use an end-to-end approach, employing a multimodal transformer with cross-modal attention to model interactions between modalities without explicit alignment. ALMT (Zhang et al., 2023) enhances model robustness under noisy and imbalanced conditions by guiding multimodal interactions through language.

**Pre-trained Models for MSA.** The rise of pre-trained models has shifted focus from handcrafted feature extraction to leveraging large pre-trained language and vision models. For example, VLP2MSA (Yi et al., 2024) integrates CLIP and TimeSformer to enhance textual prompts from video content for sentiment analysis. Other models, such as those by Yu et al. (Yu et al., 2022) and Yang et al. (Yang et al., 2023), fine-tune pre-trained language models using prompts to improve multimodal alignment for sentiment tasks.

The use of large-scale models has further propelled the performance of MSA. Wang et al. (2024) utilizes large language models to generate contextual information for sentiment analysis, while Feng et al. (2024) integrates vision-language models to reduce noise from images. Shangguan et al. (2025) propose a chain-of-thought reasoning distillation method for large models under resource constraints. Several large-scale models, such as the GPT series and mainstream open-source MLLMs, have been evaluated for zero-shot reasoning, supervised fine-tuning, and instruction tuning in MSA tasks (Song, 2024; Luo et al., 2025; Zhang et al., 2025b), demonstrating their enhanced cognitive capabilities in both multimodal and multilingual contexts. Furthermore, reinforcement learning-based methods for explainable affective computing have emerged as a growing field. R1-Omni (Zhao et al., 2025a) represents a significant step toward explainable emotion recognition. However, MLLM-based EMSA methods remain unexplored due to resource limitations.

Overall, the trend in MSA research reflects a shift from feature-level fusion to methods leveraging pre-trained and large language models, enhancing the performance of multimodal sentiment analysis. To further support the development of explainable multimodal sentiment analysis (EMSA), we introduce the benchmark dataset, CH-CEMS.

### 3 THE CH-CEMS BENCHMARK DATASET

Table 1: Multimodal and Textual Datasets Comparison.  $\checkmark$  indicates availability of the attribute,  $\times$  indicates absence.

Dataset	Samples	Speakers	Modalities	Task	Additional Annotation	Languages	Duration	Domains
YouTube	300	50	l, v, a	SP, SI	$\times$	EN	00:29	diverse
MOUD	400	101	l, v, a	SP, SI	$\times$	ES	00:59	review
ICT-MMMO	340	200	l, v, a	SP, SI	$\times$	EN	13:58	movie
MOSI	2,199	98	l, v, a	SP, SI	Subjectivity	EN	02:36	diverse
MOSEI	23,453	1,000	l, v, a	SP, SI	$\times$	EN	65:53	diverse
CH-SIMS	2281	474	l, v, a	SP, SI	unimodal sentiment	ZH	02:19	adverts
CMU-MOSEAS	40,000	1,645	l, v, a	SP, SI	Attributes, Subjectivity	Diverse	68:49	diverse
CH-SIMS v2.0	14563	–	l, v, a	SP, SI	unimodal sentiment	ZH	–	diverse
EMER	332	–	l, v, a	SP, SI, EC	Visual Clues, CoT	ZH	–	adverts
MERR	33,105	–	l, v, a	EC	Attributes, Description	ZH	–	adverts
<b>CH-CEMS (ours)</b>	<b>3715</b>	<b>1866</b>	<b>l, v, a</b>	<b>SP, SI</b>	<b>Concept, CoT</b>	<b>ZH</b>	<b>07:54</b>	<b>diverse</b>

#### 3.1 DATA COLLECTION

To approximate real-world scenarios while maintaining diversity, we collect raw videos from six scenarios: vox pops, variety shows, melodramas, formal interviews, vlogs, and science-broadcasting programs. We retain the source resolution, up to 1080p and mostly 720p, during the collection process and use an editing pipeline to segment videos into clips of appropriate length to support utterance-level analysis and predominant on-screen presence of the speaker. Compared with unimodal datasets and multimodal datasets built under earlier conventions, we place greater emphasis on fine-grained cross-modal signals to provide additional supervision. Since these signals depend on reliable audio-visual evidence, we specify three principles for clip acquisition and curation. First, we select clips that contain a complete utterance and are predominantly in Mandarin. Second, we constrain clip length to a proper range and require that the speaker exhibits a consistent sentiment without abrupt shifts. Third, to support multimodal reasoning, the speaker’s face is visible for most of the clip and the speech is intelligible without dominant background noise.

#### 3.2 DATA ANNOTATION

**Concept Definition.** Additional annotations in prior datasets primarily target unimodal features or sentiment labels to support modality-specific representation learning in late-fusion settings and representation learning across modalities in early-fusion settings. In the generative reasoning paradigm, feature representations are largely determined by large-scale pre-training in MLLMs. To strengthen reasoning during post-training, richer semantic annotations are required as explainable cues. Thus We accordingly annotate three concepts: speaking style {Informational Interaction, General Narrative Exposition} (Biber, 1989), facial expression {happy, sad, surprised, neutral, disgusted, angry}, and tone {casual, happy, sad, wistful, neutral, surprised, disgusted, angry, excited, intimate, nervous, curious, authoritative,

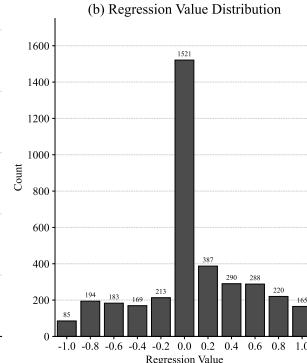
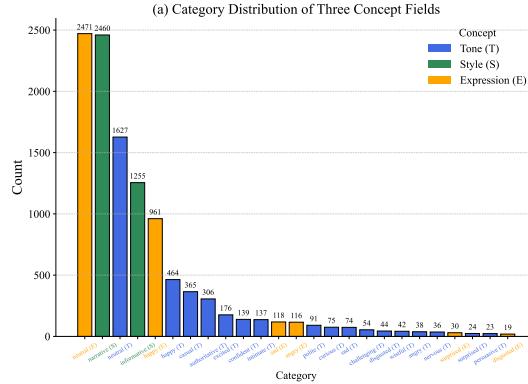


Table 2: Statistics for CH-CEMS

Metric	Value
Average utterance length	34.98
Max/min utterance length	156/2
Average clip duration(s)	7.66
Max/min clip duration(s)	36.46/0.97
Male/Female	1734/1981
kappa_Tone	0.54
kappa_Expression	0.54
kappa_Speaking Style	0.48
kappa_sentiment	0.46

Figure 2: Distribution of three concept categories (tone, style, expression) and sentiment regression values.

confident, polite, persuasive, challenging, disappointed, indifferent, expectant, sarcastic}. Details of the concepts are illustrated in Appendix B.

**Annotation Process.** The annotation procedure was carried out by a team of five annotators. Each multimodal segment was labeled with three semantic concepts and sentiment label. For concepts, annotators independently assigned labels, and the final decision was determined by a 5-out-of-3 voting scheme. If fewer than three annotators agreed, the instance was sent back for re-annotation until consensus was reached. For sentiment, annotators assigned one of three categories {positive, neutral, negative}, which were mapped to {1, 0, -1} respectively. The average score across annotators was then taken as the sentiment value, providing both a categorical polarity and a intensity score for regression-based evaluation. To align with the paradigm of sentiment classification tasks, which are better suited for large language models, we adopted a mapping approach from (Yu et al., 2020; Liu et al., 2022) to map regression values to three class (negative, neutral, positive) or five class (strong negative, weak negative, neutral, weak positive, strong positive) sentiment labels. The full workflow and more details are illustrated in Appendix A (see Figure A)

**Annotation Results.** We ultimately collected 3,715 high-quality clips to build CH-CEMS. For each sample, we provide three concept categories, sentiment annotations, and a reasoning process, as described in Section 4.1. The dataset contains six scenarios and 1,866 speakers, with three interaction types: single-person, multi-person (no interaction), and multi-person (interactive). The dataset’s label distributions vary across scenarios and interaction types, which in turn affect the performance of multimodal methods. We analyze these effects in Appendices C and H.1. Neutral sentiment accounts for the largest share (40.1%), followed by positive (36.3%) and negative (22.7%), consistent with the dataset’s predominantly real-world sources. Additionally, the category distributions for the three concepts and the sentiment regression values exhibit a long-tailed pattern, as shown in Figure 2. To assess annotation quality, we compute Fleiss’ kappa (McHugh, 2012) for the three concepts and for sentiment, and the statistics are reported in Table 2.

### 3.3 BASELINES

We establish benchmark on two tracks: (i) conventional feature-based MSA methods under a regression objective, and (ii) multimodal large language models (MLLMs) under a generative classification objective.

**Traditional regression baselines.** We evaluate five Recent state-of-the-art feature-based MSA methods under the regression paradigm commonly adopted in prior work, and follow the same data splits and evaluation procedures as on CH-SIMS and MOSEI when reporting results on CH-CEMS (see Figure 3). Specifically, **MuLT** (Tsai et al., 2019) is a multimodal transformer that applies directional pairwise cross-modal attention to capture interactions across time; **BERT-MAG** (Rahman et al., 2020b) augments a BERT backbone with a multimodal adaptation gate injected at multiple layers to integrate non-text modalities; **MISA** (Hazarika et al., 2020) jointly learns modality-invariant and modality-specific factors via distributional similarity, orthogonality constraints, reconstruction, and task losses; **DLF** (Liu et al., 2023b) disentangles shared and modality-specific representations with geometric regularization, enhances language features with a Language-Focused Attractor, and performs hierarchical prediction; and **ALMT** (Yu

et al., 2021) is an adaptive language-guided transformer that suppresses irrelevant or conflicting visual/audio signals via an Adaptive Hyper-modality Learning module and fuses the resulting hyper-modality representations.

**Classification baselines.** We evaluate four open-source and three closed-source multimodal large language models in the zero-shot setting. For open-source models, we select the most recent mainstream open-source MLLMs for evaluation, include *MiniCPM-V-4.0* (Yao et al., 2024), *MiniCPM-o-2.6* (Yao et al., 2024), *Qwen2.5-VL-7B* (Bai et al., 2025), and *Qwen2.5-Omni-7B* (Xu et al., 2025). We further optimize *Qwen2.5-Omni-7B* with our concept-guided GRPO. The closed-source baselines are *GPT-4o*, *GPT-5*, and *Gemini 2.5-Pro*. To enable direct comparison with feature-based regression baselines, we map their continuous outputs to three-class or five-class sentiment labels using the same dataset-level mapping and we evaluate all systems with the same classification metrics. Unless otherwise noted, we use a common prompt template, tag schema, input modalities, and decoding settings across models. Results are reported in Table 4 and the significance analysis is illustrated in Appendix G.

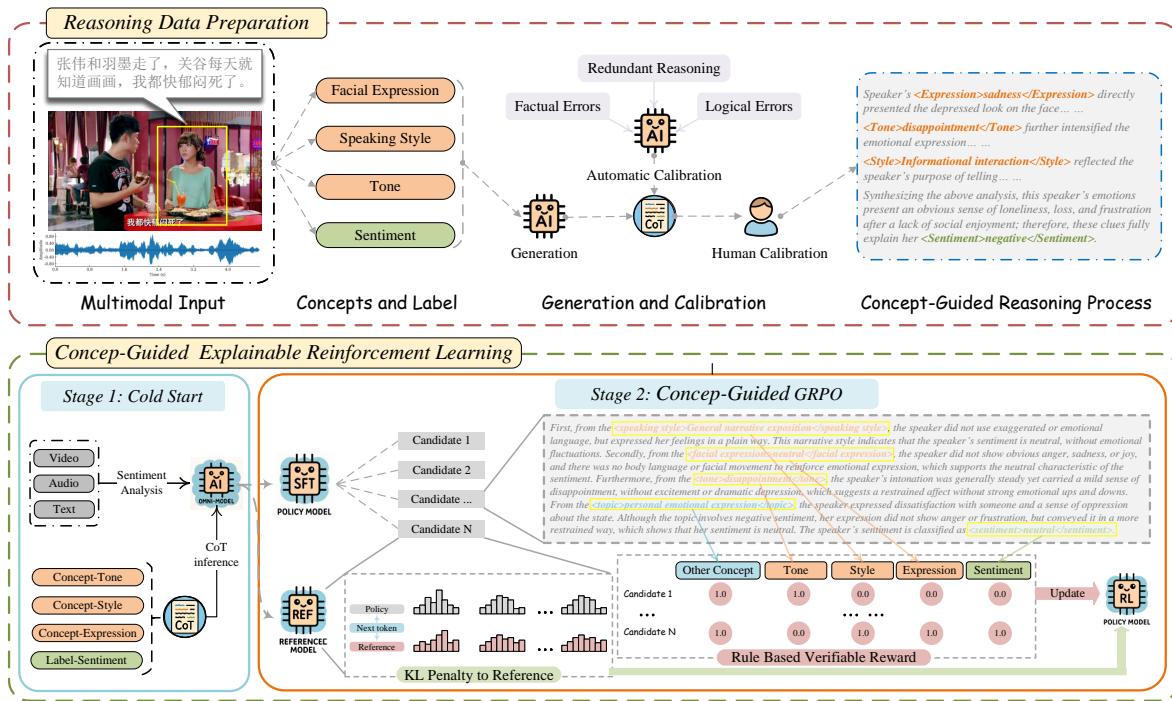


Figure 3: Overview of the Concept-Guided Explainable Reinforcement Learning Framework

Note: Although the training process used Chinese data, we provide a literal English translation here for clarity. Detailed case-study examples in the original Chinese are shown in Table 6.

## 4 CONCEPT-GUIDED EXPLAINABLE REINFORCEMENT LEARNING

This section present a practice of explainable multimodal sentiment analysis based on CH-CEMS by Concept-Guided Explainable Reinforcement Learning Framework, illustrated in Figure 3.

### 4.1 REASONING DATA CONSTRUCTION

To obtain explainable and verifiable sentiment analysis with a structured chain-of-thought, we cold-start training via supervised fine-tuning. We first build a concept-guided reasoning dataset from CH-CEMS with concept-level annotations. Each sample includes multimodal inputs, concept annotations, and a sentiment label, and is submitted to a closed-source multimodal large language model to generate a stepwise reasoning trace. In our setup, we employ GPT-4o and Gemini 2.5-Pro. The model first identifies and summarizes concept-level cues from the multimodal evidence, then analyzes how these cues relate to one another and contribute to the target sentiment, and finally produces the sentiment judgment. To make the reasoning auditable, the analysis is constrained to three predefined concepts that we

324 annotate: the speaker’s facial expression, tone of voice, and speaking style. This ensures that the chain-of-thought is  
 325 grounded in explainable concepts and follows a consistent template. In the same structured format, we also encourage  
 326 the model to initiate analyses from other summarizable concepts when supported by the evidence. The prompt used in  
 327 our approach is detailed in Appendix E.1.

328 By synthesizing evidence from all modalities and concept categories, the reasoning trace explains why a certain sentiment  
 329 (positive, negative, or neutral) is perceived. However, due to the lack of audio modality in GPT-4o and Gemini  
 330 2.5-Pro and hallucination issues during inference, we apply manual calibration. We polish the reasoning process from  
 331 three perspectives: (1) Correct factual errors: missing modality may influence the model to infer using non-existent or  
 332 incorrect facts, such as speaker attribution errors. (2) Reduce redundant reasoning: in an effort to follow instructions,  
 333 the model may generate concepts and analyses that contribute little or repeat similar arguments. (3) Correct logical  
 334 errors in the reasoning process: some relationships between concepts are incorrect or strained and should be corrected.  
 335 The upper panel of Figure 3 illustrates an example of the process.

## 337 4.2 COLD START WITH SUPERVISED FINE-TUNING

338 In the first stage of training, we cold-start the model by supervised fine-tuning (SFT) on concept-annotated reasoning  
 339 traces, adapting a pretrained multimodal large language model via LoRA for parameter efficiency. Concretely, LoRA  
 340 freezes the pretrained weights and learns a low-rank update on top: for a frozen matrix  $W_0 \in \mathbb{R}^{d \times k}$ , we introduce  
 341 trainable  $B \in \mathbb{R}^{d \times r}$  and  $A \in \mathbb{R}^{r \times k}$  with  $r \ll \min(d, k)$ ,

$$343 \quad W = W_0 + \Delta W, \quad \Delta W = BA, \quad h = W_0 x + BAx, \quad (1)$$

344 so that only the low-rank parameters  $(A, B)$  are updated while  $W_0$  remains fixed.

345 Under this parameterization, we optimize the model parameters  $\theta$  to maximize the likelihood of the annotated reasoning  
 346 sequence  $y$ , which is achieved by minimizing the cross-entropy between the model output distribution and the  
 347 reference explanation:

$$349 \quad \mathcal{L}_{\text{SFT}}(\theta) = - \sum_{t=1}^T \log P_\theta(y_t | y_{<t}, x), \quad (2)$$

351 where  $y_t$  denotes the  $t$ -th token in the ground-truth reasoning sequence and  $P_\theta(y_t | y_{<t}, x)$  is the model-assigned  
 352 probability given the preceding tokens and the input.

353 This training objective aligns the model’s behavior with the concept-annotated traces: the model learns to generate  
 354 a explanatory reasoning process, with the correct `<Style>`, `<Tone>`, and `<Expression>` tags, culminating in  
 355 the final `<Sentiment>` label. The supervision covers not only the final decision but the entire reasoning path,  
 356 encouraging the model to internalize how semantic concepts map to sentiment outcomes. After this supervised fine-  
 357 tuning (SFT) stage, the model can already produce structured explanations for sentiment, effectively acquiring the  
 358 reasoning template defined by our annotations. This cold-start initialization is critical for stabilizing the next phase,  
 359 as it provides the policy model  $\pi_\theta$  with a strong prior for concept-grounded reasoning. Consequently, subsequent  
 360 reinforcement learning refines a sensible initial policy rather than starting from scratch and mitigates reward hacking  
 361 on format-based rewards.

## 363 4.3 CONCEPT-GUIDED GRPO

364 After supervised fine-tuning, we further adapt the model with reinforcement learning to explicitly optimize explanation  
 365 quality and accuracy. We adopt a Concept-Guided Group Relative Policy Optimization algorithm, a policy-gradient  
 366 method that refines the model using task-specific reward signals. For each input, the policy  $\pi_\theta$  generates  $G$  candidate  
 367 reasoning outputs, which are then evaluated by a composite reward function comprising three components. The format  
 368 reward  $R_{\text{format}}$  ensures structural correctness of the explanation, returning 1 only if all required tags appear in the proper  
 369 structure and 0 otherwise. The answer accuracy reward  $R_{\text{acc}}$  encourages correct sentiment classification, assigning 1  
 370 when the sentiment inside the `<Sentiment>` tag matches the ground truth. The concept reward  $R_{\text{concept}}$  measures  
 371 the correctness of predicted semantic concepts, defined as

$$373 \quad R_{\text{concept}} = \frac{1}{K} \sum_{k=1}^K \mathbf{1}\{\hat{c}_k = c_k^{GT}\}, \quad (3)$$

376 where  $c_k^{GT}$  is the ground-truth label for concept  $k$  and  $\hat{c}_k$  the predicted label. The total reward is

$$377 \quad R_{\text{total}} = R_{\text{format}} + R_{\text{acc}} + R_{\text{concept}}, \quad (4)$$

378 with a maximum score when all criteria are satisfied.

379  
380 During optimization, each sampled output  $y_i$  is assigned  $r_i = R_{\text{total}}(x, y_i)$ , and group-relative advantages are com-  
381 puted as

$$382 \quad A_i = \frac{r_i - \mu_r}{\sigma_r}, \quad (5)$$

383 where  $\mu_r$  and  $\sigma_r$  are the mean and standard deviation of the group rewards. These normalized advantages serve as  
384 relative-quality signals, guiding the policy update to increase the likelihood of high-reward generations while sup-  
385 pressing low-reward ones. To stabilize training, we include a KL-regularization term that constrains  $\pi_\theta$  from deviating  
386 excessively from the SFT-initialized reference policy  $\pi_{\text{ref}}$ . The overall optimization objective is therefore  
387

$$388 \quad \max_{\theta} \mathbb{E}_{y \sim \pi_\theta} [R_{\text{total}}(x, y)] - \beta \mathbb{E}_x [\text{KL}(\pi_\theta(\cdot|x) \parallel \pi_{\text{ref}}(\cdot|x))]. \quad (6)$$

389 This procedure aligns the model with verifiable rewards covering structural validity, sentiment accuracy, and semantic  
390 consistency, thereby encouraging generations that are not only accurate but also accompanied by faithful and inter-  
391 pretable reasoning traces grounded in multimodal evidence, as illustrated in our framework diagram.  
392

## 394 5 EXPERIMENTS AND DISCUSSION

### 395 5.1 IMPLEMENTATION DETAILS

396 We follow the dataset partition shown in Tabel 6, maintaining the predefined train/validation/test splits for both the  
397 3-class and 5-class settings. For conventional feature-based MSA baselines, we adopt the official implementations  
398 and run all experiments on a single Tesla V100 GPU with 32 GB memory. For multimodal large language models,  
399 we evaluate zero-shot inference, and we train our reinforcement learning framework on eight H20 GPUs with 96 GB  
400 memory each. For conventional baselines and MLLM zero-shot inference, we report results over five random seeds  
401 (mean and standard deviation) to assess statistical significance in G. The reinforcement learning stage is implemented  
402 with the Swift framework (Zhao et al., 2024) using the GRPO backend, with concept-guided rewards as described in  
403 Section 4. We initialize the policy from Qwen2.5-Omni-7B, and train with a learning rate of  $1 \times 10^{-6}$ , a batch size  
404 of 8 per device with gradient accumulation of 4, and 3 epochs. The optimizer is AdamW with weight decay, using  
405 bfloat16 precision, the `flash_attn` implementation, and the ZeRO-2 distributed strategy. We set  $\beta = 0.04$  for  
406 KL regularization, employ 4 sampled generations per prompt with temperature 1.1 and top- $p$  0.93, and apply gradient  
407 clipping at 0.1. Additional training hyperparameters to support reproducibility are available in our repository.  
408

### 410 5.2 EVALUATION METRICS

411 We report results under two complementary evaluation paradigms. For conventional regression-based methods, we  
412 adopt the standard metrics in MSA: binary accuracy (Acc2), F1-score, weak accuracy, correlation (Corr), coefficient of  
413 determination ( $R^2$ ), and mean absolute error (MAE). These metrics capture both classification-oriented and regression-  
414 oriented perspectives, enabling fair comparison with prior work. For MLLMs, we follow the recent trend of casting  
415 MSA as a classification problem. In this setting, we report overall accuracy (ACC), weighted F1 (WF1), weighted  
416 precision (WP), recall (R), and precision (P), under both the 3-class and 5-class settings.  
417

418 Table 3: Results for feature-based methods on regression task.  
419

420 <b>Models</b>	421 <b>Acc2</b> ( $\uparrow$ )	422 <b>F1_score</b> ( $\uparrow$ )	423 <b>Acc2_weak</b> ( $\uparrow$ )	424 <b>Corr</b> ( $\uparrow$ )	425 <b>MAE</b> ( $\downarrow$ )
426 MAG-BERT	61.67/73.13	61.78/75.02	50.67	56.04	28.37
427 MISA	65.19/69.34	65.73/71.78	55.45	56.70	28.50
428 MuLT	62.53/71.74	62.92/73.89	52.18	56.00	28.47
429 ALMT	62.96/72.65	63.29/74.70	52.37	56.09	28.95
430 DLF	63.61/71.09	63.94/73.23	53.72	56.14	28.54

### 431 5.3 RESULTS ON CH-CEMS

432 **Regression track.** On CH-CEMS, prior feature-based MSA models trained under the regression paradigm exhibit  
433 modest absolute scores with no single method dominating all metrics and performance varies across systems. For

Table 4: Classification results on two tasks (3-class and 5-class).

Models	3-class						5-class					
	ACC	WF1	WP	F1	R	P	ACC	WF1	WP	F1	R	P
MAG-BERT	55.96	55.29	56.04	56.34	57.55	56.47	42.77	39.90	43.94	31.51	34.44	39.06
MISA	57.79	57.18	58.50	57.95	60.12	57.77	43.72	42.80	47.90	36.03	37.12	43.88
MuLT	56.42	55.77	56.68	56.75	58.49	56.56	42.61	40.84	46.01	33.20	35.10	42.52
ALMT	55.59	54.85	55.47	55.95	57.62	55.58	41.94	41.15	44.68	36.10	36.48	41.78
DLF	57.55	57.12	58.08	58.10	59.71	58.06	44.09	42.44	45.92	35.11	36.62	42.25
MiniCPM-V-4.0 (ZS)	52.68	46.31	56.09	49.37	56.25	56.14	28.42	22.80	28.23	18.98	35.48	17.85
MiniCPM-o-2.6 (ZS)	62.85	62.76	64.75	62.93	62.20	65.54	40.89	41.55	44.24	39.13	41.95	38.20
Qwen2.5-VL-7B (ZS)	61.91	61.91	70.69	61.08	65.58	65.58	31.22	31.22	48.74	31.35	38.82	38.82
Qwen2.5-Omni-7B (ZS)	64.98	64.89	66.48	64.61	63.12	68.38	46.92	46.65	50.78	42.04	45.67	44.88
GPT-4o (ZS)	67.43	67.43	67.48	68.29	68.08	68.54	44.01	44.32	48.99	41.43	45.59	43.30
Gemini2.5-Pro (ZS)	68.51	66.85	71.66	68.50	71.07	71.45	37.28	37.03	53.17	37.87	49.24	42.39
GPT-5 (ZS)	70.12	69.21	71.69	70.45	71.86	71.94	45.49	44.64	54.60	47.40	<b>54.32</b>	48.66
MiniCPM-V-4.0 (SFT)	62.45	62.42	63.61	62.26	60.92	65.23	45.49	38.38	43.83	27.95	32.71	34.09
MiniCPM-o-2.6 (SFT)	65.01	64.97	65.50	65.63	65.83	65.95	49.53	49.54	50.16	32.70	32.93	33.22
Qwen2.5-VL-7B (SFT)	71.60	71.53	71.62	71.66	72.13	71.34	55.59	54.79	<b>55.89</b>	48.22	46.69	<b>53.43</b>
Qwen2.5-Omni-7B(SFT)	70.12	70.14	70.38	70.42	70.14	70.91	52.89	50.12	51.96	47.79	50.48	50.93
Qwen2.5-Omni-7B (GRPO)	69.18	69.44	70.09	69.07	69.21	69.56	51.68	46.31	48.27	47.30	50.91	52.05
CD-GRPO (ours)	<b>73.62</b>	<b>73.61</b>	<b>75.56</b>	<b>74.06</b>	<b>74.77</b>	<b>75.03</b>	<b>56.39</b>	<b>54.86</b>	<b>55.60</b>	<b>51.75</b>	<b>53.76</b>	51.83

example, some models favor correlation and error while others are relatively stronger on Acc2 or F1, indicating the increased difficulty and distribution shift in CH-CEMS. Representative numbers are reported in Table 3.

**Classification track.** In the 3-class setting, a clear gap emerges: small feature-based models trail strong MLLM baselines by about 10 percentage points in accuracy. Among off-the-shelf MLLMs, GPT-5 is the strongest closed-source baseline. Building on an open-source backbone, our concept-guided post-training on Qwen2.5-Omni-7B achieves 73.62% accuracy on the three-class sentiment task (Table 4), exceeding GPT-5 by 3.50 points and open-source SFT baselines by 2.0–10.0 points, with consistent gains in WF1, WP, F1, recall, and precision. In the 5-class setting, the gap narrows: several traditional baselines match or exceed closed-source models on some metrics, and GPT-5 attains the strongest accuracy among evaluated MLLMs under zero-shot inference. CD-GRPO achieves performance comparable to Qwen2.5-VL-7B under SFT and surpasses most other methods, attaining top scores on three metrics. Unlike label-only predictors, it also outputs concept-grounded, verifiable reasoning traces which is demonstrated in Figure 6. Overall, large models excel at polarity recognition (3-class), while fine-grained intensity (5-class) remains challenging. The performance of CD-GRPO demonstrates the effectiveness of concept-level semantic cues from CH-CEMS under a reinforcement learning paradigm and offers a solid basis for further research on explainable multimodal sentiment analysis.

## 6 CONCLUSIONS

Motivated by the growing interest in explainable affective computing and the fact that while explainable datasets are plentiful for emotion recognition they remain scarce for sentiment analysis, we introduce CH-CEMS, a multimodal sentiment dataset with concept-level annotations and accompanying reasoning traces towards EMSA. Building on CH-CEMS, we propose a training pipeline that cold-starts the model via supervised fine-tuning on structured explanations and then applies rule-based reinforcement learning with concept-level rewards, providing verifiable supervision over the reasoning process. To support research on CH-CEMS, we benchmark both feature-based regression baselines and generative classification with four open- and three closed-source multimodal large language models. On our benchmark, MLLMs exhibit strong overall performance, and our concept-guided GRPO on Qwen2.5-Omni-7B surpasses strong closed-source baselines such as GPT-5 while producing structured reasoning traces under a unified tag schema, highlighting the effectiveness of semantic cues and a promising direction for explainable multimodal sentiment analysis in the generative paradigm.

486 REFERENCES  
487

488 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun  
489 Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren  
490 Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and  
491 Junyang Lin. Qwen2.5-v1 technical report. *arXiv preprint arXiv:2502.13923*, 2025.

492 Douglas Biber. *Variation Across Speech and Writing*. Cambridge University Press, Cambridge, UK, 1988.  
493

494 Douglas Biber. A typology of english texts. *Linguistics*, 27:3–43, 1989.  
495

496 Penelope Brown and Stephen C. Levinson. *Politeness: Some Universals in Language Usage*. Cambridge University  
497 Press, Cambridge, UK, 1987.

498 DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL  
499 <https://arxiv.org/abs/2501.12948>.  
500

501 Paul Ekman. An argument for basic emotions. *Cognition and Emotion*, 6(3-4):169–200, 1992. doi: 10.1080/  
502 02699939208411068.

503 Paul Ekman, Wallace V. Friesen, and Joseph C. Hager. *Facial Action Coding System (FACS): Manual and Investiga-  
504 tor's Guide*. Research Nexus, Salt Lake City, UT, 2002. 2nd edition.  
505

506 Cunhang Fan, Kang Zhu, Jianhua Tao, Guofeng Yi, Jun Xue, and Zhao Lv. Multi-level contrastive learning: Hierar-  
507 chical alleviation of heterogeneity in multimodal sentiment analysis. *IEEE Transactions on Affective Computing*,  
508 16(1):207–222, 2024.

509 Hao Fei, Bobo Li, Qian Liu, Lidong Bing, Fei Li, and Tat-Seng Chua. Reasoning implicit sentiment with chain-of-  
510 thought prompting. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp.  
511 1171–1182, 2023.  
512

513 Junjia Feng, Mingqian Lin, Lin Shang, and Xiaoying Gao. Autonomous aspect-image instruction a2ii: Q-former  
514 guided multimodal sentiment classification. In *Proceedings of the 2024 joint international conference on computa-  
515 tional linguistics, language resources and evaluation (LREC-COLING 2024)*, pp. 1996–2005, 2024.  
516

517 Devamanyu Hazarika, Roger Zimmermann, and Soujanya Poria. Misa: Modality-invariant and -specific representa-  
518 tions for multimodal sentiment analysis. In *Proceedings of the 28th ACM International Conference on Multi-  
519 media*, pp. 1122–1131, 2020. doi: 10.1145/3394171.3413678. URL [https://doi.org/10.1145/3394171.  
520 3413678](https://doi.org/10.1145/3394171.3413678).

521 Guimin Hu, Ting-En Lin, Yi Zhao, Guangming Lu, Yuchuan Wu, and Yongbin Li. Unimse: Towards unified multi-  
522 modal sentiment analysis and emotion recognition. *arXiv preprint arXiv:2211.11256*, 2022.  
523

524 Qi Huang, Lin Zhao, Li Li, Wei Zhang, and Yi Chen. Multimodal pear chain-of-thought reasoning for msa. In  
525 *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, pp. 15, 2024. URL  
526 <https://doi.org/10.1145/3542614>.  
527

528 Jamin Rahman Jim, Md Apon Riaz Talukder, Partha Malakar, Md Mohsin Kabir, Kamruddin Nur, and M.F. Mridha.  
Recent advancements and challenges of nlp-based sentiment analysis: A state-of-the-art review. *Natural Language  
529 Processing Journal*, 6:100059, 2024. ISSN 2949-7191. doi: <https://doi.org/10.1016/j.nlp.2024.100059>. URL  
530 <https://www.sciencedirect.com/science/article/pii/S2949719124000074>.  
531

532 Zaid Khan and Yun Fu. Exploiting bert for multimodal target sentiment classification through input space translation.  
In *Proceedings of the 29th ACM international conference on multimedia*, pp. 3034–3042, 2021.  
533

534 Jing Li, Wei Zhang, Zhihong Yang, and Huixiao Li. Disentangling bias by modeling intra- and inter-modal causal  
535 attention for multimodal sentiment analysis. *arXiv preprint arXiv:2506.04567*, 2025. URL <https://arxiv.org/abs/2506.04567>.  
536

537 Shuang Li, Xiang Zhang, Zhiqiang Liu, and Wen Wang. Infoenh: Towards multimodal sentiment analysis via informa-  
538 tion bottleneck filter and optimal transport alignment. *Proceedings of the 2024 Language Resources and Evaluation  
539 Conference (LREC)*, 2024. URL <https://aclanthology.org/2024.lrec-1.128>.

540 Yan Ling, Jianfei Yu, and Rui Xia. Vision-language pre-training for multimodal aspect-based sentiment analysis. In  
 541 *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,  
 542 pp. 2149–2159, 2022.

543 Wenjie Liu, Shoushan Li, Qinghua Zhang, Qian Chen, and Guodong Zhou. Explainable multimodal emotion reasoning. In *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 1412–1420. ACM, 2023a. doi:  
 544 10.1145/3581783.3612255.

545 Yihe Liu, Ziqi Yuan, Huisheng Mao, et al. Make acoustic and visual cues matter: CH-SIMS v2.0 dataset and av-mixup  
 546 consistent module. In *ICMI*, 2022.

547 Yihe Liu, Ziqi Yuan, Huisheng Mao, and Tiejun Zhao. A dual-level fusion framework for multimodal sentiment  
 548 analysis. *IEEE Transactions on Affective Computing*, 2023b. doi: 10.1109/TAFFC.2023.3245678.

549 Zhiwei Liu, Lingfei Qian, Qianqian Xie, Jimin Huang, Kailai Yang, and Sophia Ananiadou. Mmaffben: A multilingual  
 550 and multimodal affective analysis benchmark for evaluating llms and vlms. *arXiv preprint arXiv:2505.24423*, 2025.

551 Patrick Lucey, Jeffrey F. Cohn, Takeo Kanade, Jason Saragih, Zara Ambadar, and Iain Matthews. The extended cohn–  
 552 kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *2010 IEEE Computer  
 553 Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 94–101, 2010. doi:  
 554 10.1109/CVPRW.2010.5543262.

555 Meng Luo, Hao Fei, Bobo Li, Shengqiong Wu, Qian Liu, Soujanya Poria, Erik Cambria, Mong-Li Lee, and Wynne  
 556 Hsu. Panosent: A panoptic sextuple extraction benchmark for multimodal conversational aspect-based sentiment  
 557 analysis. In *Proceedings of ACM Multimedia (ACM MM)*, 2024.

558 Miaosen Luo, Jiesen Long, Zequn Li, Yunying Yang, Yuncheng Jiang, and Sijie Mai. Multimodal large language  
 559 models for end-to-end affective computing: Benchmarking and boosting with generative knowledge prompting.  
 560 *arXiv preprint arXiv:2508.02429*, 2025. URL <https://arxiv.org/abs/2508.02429>.

561 Mary L McHugh. Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3):276–282, 2012.

562 Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, AmirAli Bagher Zadeh, Chengfeng Mao, Louis-Philippe Morency,  
 563 and Ehsan Hoque. Integrating multimodal information in large pretrained transformers. In *Proceedings of the 58th  
 564 Annual Meeting of the Association for Computational Linguistics*, pp. 2359–2369, 2020a. doi: 10.18653/v1/2020.  
 565 acl-main.214. URL <https://aclanthology.org/2020.acl-main.214/>.

566 Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, et al. Integrating multimodal information in large pretrained  
 567 transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*,  
 568 pp. 2359–2369, 2020b.

569 James A. Russell. A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6):1161–1178,  
 570 1980. doi: 10.1037/h0077714.

571 Klaus R. Scherer. Vocal communication of emotion: A review of research paradigms. *Speech Communication*, 40  
 572 (1-2):227–256, 2003. doi: 10.1016/S0167-6393(02)00084-5.

573 Haonan Shangguan, Xiaocui Yang, Shi Feng, Daling Wang, Yifei Zhang, and Ge Yu. Resource-limited joint multi-  
 574 modal sentiment reasoning and classification via chain-of-thought enhancement and distillation. *arXiv preprint  
 575 arXiv:2508.05234*, 2025.

576 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y.K. Li, Yang Wu, and Daya  
 577 Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.

578 Shezheng Song. Exploring large language models for multimodal sentiment analysis: Challenges, benchmarks, and  
 579 future directions. *arXiv preprint arXiv:2411.15408*, 2024.

580 Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhut-  
 581 dinov. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference.  
 582 Association for computational linguistics. Meeting*, volume 2019, pp. 6558, 2019.

583 Pan Wang, Qiang Zhou, Yawen Wu, Tianlong Chen, and Jingtong Hu. Dlf: Disentangled-language-focused multi-  
 584 modal sentiment analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp.  
 585 21180–21188, 2025a.

594 Wenbin Wang, Liang Ding, Li Shen, Yong Luo, Han Hu, and Dacheng Tao. Wisdom: Improving multimodal sentiment  
 595 analysis by fusing contextual world knowledge. In *Proceedings of the 32nd ACM international conference on*  
 596 *multimedia*, pp. 2282–2291, 2024.

597

598 Yaoting Wang, Shengqiong Wu, Yuecheng Zhang, Shuicheng Yan, Ziwei Liu, Jiebo Luo, and Hao Fei. Multimodal  
 599 chain-of-thought reasoning: A comprehensive survey. *arXiv preprint arXiv:2503.12605*, 2025b.

600

601 Martin Wöllmer, Felix Weninger, Lorenz Rausch, and Björn Schuller. Youtube movie reviews: Sentiment analysis by  
 602 fusion of BoW, AV features and hidden markov models. In *ICASSP*, 2013. ICT-MMMO dataset.

603

604 Zhiwei Wu, Shoushan Li, Qinghua Zhang, Qian Chen, and Guodong Zhou. Multimodal emotion recognition and  
 605 reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, pp.  
 606 5678–5690. ACL, 2023. doi: 10.18653/v1/2023.acl-long.400.

607

608 Luwei Xiao, Rui Mao, Shuai Zhao, Qika Lin, Yanhao Jia, Liang He, and Erik Cambria. Exploring cognitive and  
 609 aesthetic causality for multimodal aspect-based sentiment analysis. *IEEE Transactions on Affective Computing*,  
 610 2025.

611

612 Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai  
 613 Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. Qwen2.5-omni technical report. *arXiv preprint*  
 614 *arXiv:2503.20215*, 2025.

615

616 Dingkang Yang, Mingcheng Li, Dongling Xiao, Yang Liu, Kun Yang, Zhaoyu Chen, Yuzheng Wang, Peng Zhai, Ke Li,  
 617 and Lihua Zhang. Towards multimodal sentiment analysis debiasing via bias purification. In *European Conference*  
 618 *on Computer Vision*, pp. 464–481. Springer, 2024a.

619

620 Hao Yang, Yanyan Zhao, Yang Wu, Shilong Wang, Tian Zheng, Hongbo Zhang, Zongyang Ma, Wanxiang Che, and  
 621 Bing Qin. Large language models meet text-centric multimodal sentiment analysis: A survey. *arXiv preprint*  
 622 *arXiv:2406.08068*, 2024b. URL <https://arxiv.org/abs/2406.08068>.

623

624 Xiaocui Yang, Shi Feng, Daling Wang, Qi Sun, Wenfang Wu, Yifei Zhang, Pengfei Hong, and Soujanya Poria. Few-  
 625 shot joint multimodal aspect-sentiment analysis based on generative multimodal prompt. In *The 61st Annual Meet-  
 626 ing Of The Association For Computational Linguistics*, 2023.

627

628 Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui  
 629 He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024.

630

631 Junjie Ye, Jie Zhou, Junfeng Tian, Rui Wang, Jingyi Zhou, Tao Gui, Qi Zhang, and Xuanjing Huang. Sentiment-aware  
 632 multimodal pre-training for multimodal sentiment analysis. *Knowledge-Based Systems*, 258:110021, 2022.

633

634 Guofeng Yi, Cunhang Fan, Kang Zhu, Zhao Lv, Shan Liang, Zhengqi Wen, Guanxiong Pei, Taihao Li, and Jianhua Tao.  
 635 Vlp2msa: expanding vision-language pre-training to multimodal sentiment analysis. *Knowledge-Based Systems*,  
 636 283:111136, 2024.

637

638 Fei Yu, Jinchao Xu, Zixing Zhang, Bo Chen, Shizhe Yu, and Tiejun Zhao. Ch-sims: A chinese multimodal senti-  
 639 ment analysis dataset with fine-grained annotation of modality. In *Proceedings of the 58th Annual Meeting of the*  
 640 *Association for Computational Linguistics (ACL)*, pp. 3718–3727, 2020. doi: 10.18653/v1/2020.acl-main.342.

641

642 Fei Yu, Jinchao Xu, Ziqi Yuan, and Tiejun Zhao. Learning modality-specific representations with adversarial training  
 643 for multimodal sentiment analysis. In *Proceedings of the 59th Annual Meeting of the Association for Computational*  
 644 *Linguistics (ACL)*, pp. 4820–4832, 2021.

645

646 Fei Yu, Zixing Zhang, Jinchao Xu, Bo Chen, and Tiejun Zhao. Ch-sims 2.0: A large-scale chinese multimodal senti-  
 647 ment analysis dataset with fine-grained annotations. In *Findings of the Association for Computational Linguistics:*  
 648 *ACL 2023*, pp. 1455–1468, 2023. doi: 10.18653/v1/2023.findings-acl.114.

649

650 Yang Yu and Dong Zhang. Few-shot multi-modal sentiment analysis with prompt-based vision-aware language mod-  
 651 eling. In *2022 IEEE international conference on multimedia and expo (ICME)*, pp. 1–6. IEEE, 2022.

652

653 Yang Yu, Dong Zhang, and Shoushan Li. Unified multi-modal pre-training for few-shot sentiment analysis with  
 654 prompt-based learning. In *Proceedings of the 30th ACM international conference on multimedia*, pp. 189–198,  
 655 2022.

648 Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Multimodal sentiment  
 649 intensity analysis in videos: Facial gestures and verbal messages. In *2016 IEEE International Conference on*  
 650 *Multimedia and Expo (ICME)*, pp. 1–6. IEEE, 2016. doi: 10.1109/ICME.2016.7553005.

651 Amir Zadeh, Paul Pu Liang, Sahisnu Mazumder, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Multi-  
 652 modal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *Proceedings*  
 653 *of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 2236–2246, 2018. doi:  
 654 10.18653/v1/P18-1208.

655 AmirAli Bagher Zadeh, Paul Pu Liang, et al. Cmu-moseas: A multimodal language dataset for spanish, portuguese,  
 656 german and french. In *EMNLP*, 2020.

657 Fan Zhang, Zebang Cheng, Chong Deng, Haoxuan Li, Zheng Lian, Qian Chen, Huadai Liu, Wen Wang, Yi-Fan  
 658 Zhang, Zhang Renrui, Ziyu Guo, Zhihong Zhu, Hao Wu, Haixin Wang, Yefeng Zheng, Xiaojiang Peng, Xian  
 659 Wu, Kun Wang, Xiangang Li, Jieping Ye, and Pheng-Ann Heng. Mme-emotion: A holistic evaluation benchmark  
 660 for emotional intelligence in multimodal large language models. *arXiv preprint arXiv:2508.09210*, 2025a. URL  
 661 <https://arxiv.org/abs/2508.09210>.

662 Hanlei Zhang, Zhuohang Li, Yeshuang Zhu, Hua Xu, Peiwu Wang, Haige Zhu, Jie Zhou, and Jinchao Zhang. Can  
 663 large language models help multimodal language analysis? mmla: A comprehensive benchmark. *arXiv preprint*  
 664 *arXiv:2504.16427*, 2025b.

665 Haoyu Zhang, Yu Wang, Guanghao Yin, Kejun Liu, Yuanyuan Liu, and Tianshu Yu. Learning language-guided  
 666 adaptive hyper-modality representation for multimodal sentiment analysis. *arXiv preprint arXiv:2310.05804*, 2023.

667 Hongyu Zhang, Xue Zhang, Shuang Li, and Bing Liu. Hierarchical reasoning enhanced few-shot multimodal sen-  
 668 timent analysis. *Neurocomputing*, 400:123–134, 2025c. URL <https://doi.org/10.1016/j.neucom.2024.06.018>.

669 Rui Zhang, Shoushan Li, Qian Chen, and Guodong Zhou. Emotion-llama: Multimodal emotion recognition and  
 670 reasoning with instruction tuning. *arXiv preprint arXiv:2402.14532*, 2024. doi: 10.48550/arXiv.2402.14532.

671 Jiawei Zhao, Yang Chen, Shoushan Li, and Guodong Zhou. R1-omni: Explainable omni-multimodal emotion recog-  
 672 nition with reinforcement learning. *arXiv preprint arXiv:2503.05379*, 2025a. doi: 10.48550/arXiv.2503.05379.

673 Xianbing Zhao, Xuejiao Li, Ronghuan Jiang, and Buzhou Tang. Decoupled cross-attribute correlation network  
 674 for multimodal sentiment analysis. *Information Fusion*, 117:102897, 2025b. ISSN 1566-2535. doi: <https://doi.org/10.1016/j.inffus.2024.102897>. URL <https://www.sciencedirect.com/science/article/pii/S1566253524006754>.

675 Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang, Zhikai Wu, Baole  
 676 Ai, Ang Wang, Wenmeng Zhou, and Yingda Chen. Swift:a scalable lightweight infrastructure for fine-tuning, 2024.  
 677 URL <https://arxiv.org/abs/2408.05517>.

## 686 A DATASET CONSTRUCTION PIPELINE

### 687 A.1 DATA SOURCE

688 Our dataset was collected from Bilibili and comprises unedited video clips drawn from vox pops, variety shows, melo-  
 689 dramas, formal interviews, vlogs, and other real-world scenarios. These videos capture rich, multimodal emotional  
 690 expressions in natural contexts, offering high authenticity and diversity. More than 97% of the videos have a reso-  
 691 lution of at least 720p, ensuring visual clarity and providing high-quality audiovisual inputs for fine-grained sentiment  
 692 analysis.

### 693 A.2 DATA PROCESSING

694 Due to variations in shooting angles, distances, and lighting conditions in the original videos, the environments in  
 695 which the speakers are situated exhibit considerable complexity. In addition, the audio may contain minor noise, such  
 696 as background music or overlapping monologues, and the speakers' speech rate and tone may also vary. To construct  
 697 a high-quality multimodal sentiment analysis dataset, we designed a systematic data processing pipeline consisting of  
 698 two key steps: video filtering and speech utteranceion.

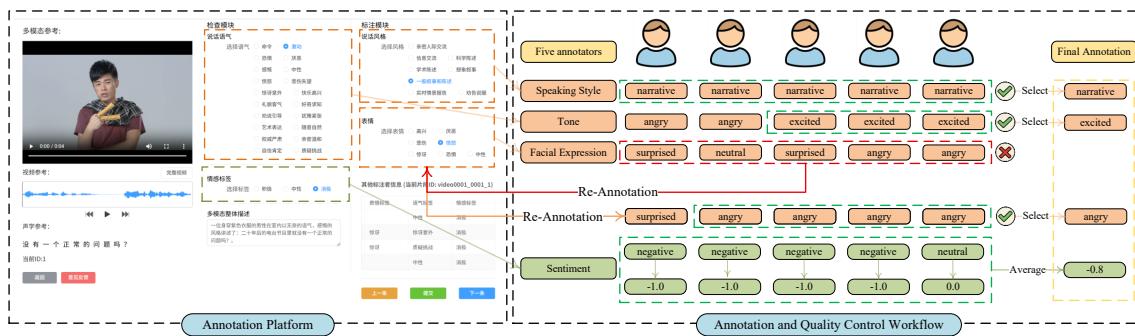
702 **Video Clip Filtering.** All video clips were required to retain their original resolution and be stored in MP4 format.  
 703 Using dedicated cropping tools, the original videos were trimmed into segments no longer than two minutes, which  
 704 were subsequently divided into shorter utterances. A semantic filtering procedure was then applied to automatically  
 705 remove clips lacking substantive semantic content.

706 Building on this step, we introduced a manual screening mechanism to further enhance data quality. The criteria for this  
 707 second stage were as follows: for the visual modality, the video must clearly display the speaker’s face, and the speaker  
 708 must remain in the frame for more than half of the segment’s duration to ensure effective visual information capture;  
 709 for the auditory modality, the speaker’s voice must be clear, and the intensity of background music or noise must not  
 710 interfere with the speech signal, thereby ensuring the recognizability and consistency of the audio. Two professionally  
 711 trained annotators participated in this stage, performing strict screening from the perspective of multimodal reliability,  
 712 which ultimately yielded high-quality sample instances.

713 **Text utteranceion.** To obtain textual data that accurately aligns with the audio-visual content, we employed iFlytek’s  
 714 closed-source speech recognition API for initial utteranceion. All transcribed texts were then manually proofread to  
 715 correct typos, sentence segmentation, and punctuation errors. Personally identifiable information (such as names) was  
 716 anonymized. This process substantially improved the accuracy, consistency, and security of the textual data, providing  
 717 a reliable linguistic foundation for multimodal sentiment analysis.

### 719 A.3 DATA ANNOTATION

721 In the data annotation phase, we independently developed an efficient multimodal sentiment annotation platform and  
 722 established a unified database for both concept and sentiment labels, thereby significantly enhancing annotation effi-  
 723 ciency and data quality. The platform was designed to enable seamless interaction between annotators and multimodal  
 724 data, featuring an intuitive and user-friendly interface, as illustrated in Figure 4.



737 Figure 4: Workflow for dataset annotation: item selection, multi-annotator labeling, quality assurance review,  
 738 re-annotation, and label finalization.

740 The platform supports two core annotation tasks: concept labeling and sentiment labeling. Concept labels span three  
 741 dimensions: the speaker’s tone, speaking style, and facial expressions. By comprehensively watching videos, listening  
 742 to audio, and reading text, annotators extract information from the visual, auditory, and textual modalities to complete  
 743 concept annotations, and subsequently determine the overall sentiment label through multimodal information fusion.  
 744 The annotation process is accomplished through simple point-and-click interactions, which greatly reduce operational  
 745 complexity and learning costs, while significantly accelerating the annotation workflow. All annotation results are  
 746 stored in the database in real time for subsequent statistical analysis. In addition, descriptive text was generated  
 747 for each multimodal sample, providing a valuable resource for researchers conducting further studies in multimodal  
 748 sentiment analysis.

749 We also present a representative case of the constructed CH-CEMS. A total of five professional annotators participated  
 750 in labeling both concept and sentiment categories. The platform’s user-friendly interface allows annotators to  
 751 conveniently compare video and text content, thereby ensuring annotation accuracy and consistency.

752 During the concept labeling stage, we introduced a “five-choose-three” voting mechanism, requiring that each semantic  
 753 concept be selected by at least three annotators to be confirmed. To address the sparsity of samples in certain speaking-  
 754 style categories, we merged and optimized them: “intimate interpersonal communication” and “information exchange”  
 755 were unified into “information exchange,” while other styles were grouped into “narrative.” Labels that failed to reach

756 consensus were returned for re-discussion and re-annotation, ensuring both high quality and consistency of the concept  
 757 labels.

758 Sentiment labeling was built upon the foundation of concept annotations. We adopted a polarity-based sentiment  
 759 quantification scheme, categorizing sentiments into positive (1.0), neutral (0.0), and negative (-1.0). This approach is  
 760 consistent with the sentiment labeling methods employed in CH-SIMS (Yu et al., 2020), CH-SIMS v2.0 (Yu et al.,  
 761 2023), CMU-MOSI (Zadeh et al., 2016), and CMU-MOSEI (Zadeh et al., 2018). By summing and averaging the  
 762 scores assigned by each annotator, we derived continuous sentiment regression values, thereby providing a fine-grained  
 763 depiction of the speaker’s sentiment state.

764 This systematic annotation process not only guarantees the reliability and consistency of the annotated data—laying a  
 765 solid foundation for building robust, high-performance multimodal sentiment analysis models—but also facilitates ef-  
 766 ficient data collection and serves as a valuable tool for exploring the complex relationships between semantic concepts  
 767 and sentiments.

## 769 B CONCEPT DEFINITIONS

770 **Speaking Style label taxonomy.** We adopt Biber’s data-driven register framework because it defines speaking style  
 771 by statistically induced functional dimensions rather than preset genres. Concretely, large balanced corpora are anno-  
 772 tated with dozens of lexical–grammatical features; factor analysis over feature co-occurrence yields five continuous  
 773 dimensions (Involved vs. Informational production, Narrative vs. non-narrative concerns, Elaborated vs. situation-  
 774 dependent reference, Overt expression of persuasion, Abstract vs. non-abstract style). Clustering texts in this space  
 775 gives eight empirically grounded prototypes as illustrated in Table 5. This framework is attractive for concept anno-  
 776 tation because it: (i) ties labels to measurable linguistic cues, (ii) provides interpretable axes and prototype labels, and  
 777 (iii) reduces subjective, task-specific labeling (Biber, 1988; 1989).

778 In our corpus, many prototype classes are too sparse for reliable supervision. We therefore focus our speaking style  
 779 annotation on two frequent and complementary prototypes: Informational interaction and General narrative exposition.  
 780 These labels capture the dominant interactional–informational and mixed narrative–expository styles in our data, while  
 781 keeping annotation consistent with the underlying dimensional view of style.

782 Table 5: Eight prototype text types in Biber’s typology and concise definitions

783 <b>ID</b>	784 <b>Prototype name</b>	785 <b>Concise definition and explanation</b>
786 1	787 Intimate interpersonal interaction	788 Concerned primarily with the immediate interpersonal interaction.
789 2	790 Informational interaction	791 Has a primary informational emphasis.
790 3	791 Scientific exposition	792 Extremely informational, elaborated in reference, and technical and abstract 793 in style and content.
792 4	793 Learned exposition	794 Similar to Scientific exposition except that it is markedly less abstract and 795 less technical in style.
794 5	795 Imaginative narrative	796 A relatively involved text type having a primary narrative focus.
795 6	796 General narrative exposition	797 A very general text type that combines narrative forms with expository, 798 informational elaboration.
797 7	798 Situated reportage	799 Reporting events actually in progress.
798 8	799 Involved persuasion	800 Primarily distinguished by persuasive and argumentative emphases, typically 801 combined with an involved (often interactive) style.

802 **Facial expression label taxonomy.** We annotate six facial expression categories (happy, sad, surprised, neutral,  
 803 disgusted, and angry) because they align with the widely used basic-emotion taxonomy and its operationalization via  
 804 the Facial Action Coding System (FACS), which ties labels to observable action units and supports reproducible coding  
 805 (Ekman, 1992; Ekman et al., 2002). The same categorical set (with a neutral baseline) is standard in mainstream FER  
 806 benchmarks, enabling direct reuse of models and fair comparison across datasets (Lucey et al., 2010). Practically, we  
 807 replace fear with neutral to improve annotation reliability and class balance for limited fear sample.

808 **Tone label taxonomy.** We ground our tone labels in two complementary traditions. First, affective tone is anchored  
 809 in classic emotion theory: discrete emotions (happy, sad, angry, disgusted, surprised) and a neutral baseline, together

with valence–arousal nuances (excited, nervous, wistful) that are well captured by dimensional models and known to surface in prosody and voice quality (Ekman, 1992; Russell, 1980; Scherer, 2003). Second, interpersonal tone (casual, intimate, polite, persuasive, challenging, authoritative, confident, curious) draws on pragmatics of facework and stance as well as register/style research, linking these labels to recognizable linguistic resources (e.g., mitigation, commitment displays, formality cues) (Brown & Levinson, 1987; Biber, 1988). Practically, we open-coded candidate tone tags on a seed set and consolidated them into the above inventory under the constraints that (i) core valence–arousal quadrants are covered, and (ii) key interpersonal functions in dialogue are represented, yielding a compact, theory-aligned taxonomy for downstream modeling and evaluation.

## C DETAILED INFORMATION OF DATASET

We collected raw videos from various scenarios, including vox pops, variety shows, melodramas, formal interviews, vlogs, and science-broadcasting programs. We observe that the sentiment distribution varies across scenarios, indicating a scenario-level distribution shift. Besides, prior multimodal sentiment datasets often restrict clips to a single on-screen speaker to avoid confounds from additional people. Accordingly, we annotate the number of on-screen participants and their interaction relations. Sentiment distributions conditioned on both factors (scenario and speaker count/interaction) are shown in Figure 5. For training and evaluation on CH-CEMS, we split the dataset into train/dev/test sets using a 6:2:2 ratio and perform a stratified split to maintain the sentiment-label distribution across splits. The resulting distributions are reported in Table 6

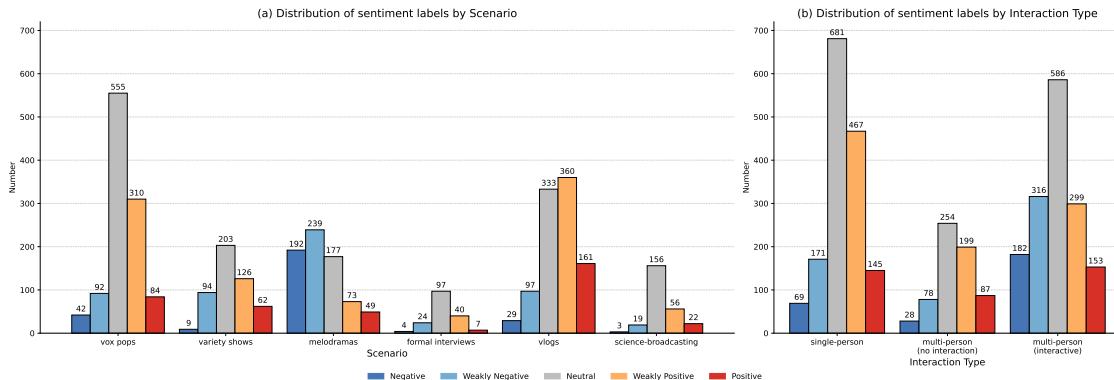


Figure 5: Distribution of five-level sentiment labels (negative, weakly negative, neutral, weakly positive, positive) across (a) six scenarios (vox pops, variety shows, melodramas, formal interviews, vlogs, and science-broadcasting programs) and (b) three interaction types (single-person, multi-person without interaction, and multi-person with interaction)

Item	Total	5-class					3-class		
		NEG	WNEG	NEU	WPOS	POS	NEG	NEU	POS
#Train	2229	167 (7.5%)	339 (15.2%)	913 (41.0%)	579 (26.0%)	231 (10.4%)	506 (22.7%)	913 (41.0%)	810 (36.3%)
#Valid	743	56 (7.5%)	113 (15.2%)	304 (40.9%)	193 (26.0%)	77 (10.4%)	169 (22.7%)	304 (40.9%)	270 (36.3%)
#Test	743	56 (7.5%)	113 (15.2%)	304 (40.9%)	193 (26.0%)	77 (10.4%)	169 (22.7%)	304 (40.9%)	270 (36.3%)

Table 6: Label distribution of the dataset in training, validation, and test sets. Numbers in parentheses indicate the percentage of each class. For 5-class classification: NEG (Negative), WNEG (Weak Negative), NEU (Neutral), WPOS (Weak Positive), POS (Positive); for 3-class classification: NEG (Negative), NEU (Neutral), POS (Positive).

## D ADDITIONAL RELATED WORK

### D.1 FEATURE-BASED MSA METHODS

Research on multimodal sentiment analysis (MSA) mainly focuses on two directions: early fusion and late fusion. Early fusion methods aim to capture cross-modal interactions and consistency at the feature level. Among them,

864 MISA (Hazarika et al., 2020) proposes modality-invariant and modality-specific representations to improve multimodal  
 865 fusion for affective state prediction. Building on the idea of disentangled representations, DLF (Wang et al., 2025a)  
 866 proposes a Disentangled-Language-Focused (DLF) multimodal representation learning framework to further enhance  
 867 language-targeted features. Similarly, MAG-BERT (Rahman et al., 2020a) enhances the multimodal fusion capability  
 868 of pre-trained models through the Multimodal Adaptation Gate, highlighting the effectiveness of early fusion at the  
 869 feature level.

870 Beyond conventional early fusion methods, some works focus on end-to-end modeling of multimodal sequences.  
 871 For instance, Mult (Tsai et al., 2019) is an end-to-end multimodal Transformer that models interactions and adap-  
 872 tations between multimodal sequences without explicit alignment through directional pairwise cross-modal attention.  
 873 ALMT (Zhang et al., 2023) presents the Adaptive Language-guided Multimodal Transformer to enhance robustness  
 874 under noise and modality imbalance, emphasizing the importance of model robustness. Moreover, UniMSE (Hu et al.,  
 875 2022) leverages joint modeling of MSA and ERC combined with contrastive learning to improve cross-task perfor-  
 876 mance, demonstrating the benefit of exploiting task similarities.

877 Semi-supervised learning and modality heterogeneity are also actively explored. MC-Teacher utilizes a teacher-student  
 878 framework augmented with consistency-based pseudo-labels to perform semi-supervised multimodal sentiment analy-  
 879 sis, effectively addressing the challenge of limited annotated data. To handle modality heterogeneity, MCL-MCF (Fan  
 880 et al., 2024) is a progressive multimodal fusion method that alleviates modality heterogeneity at multiple levels through  
 881 multi-level contrastive learning and tensor convolution fusion, enabling continuous feature-level integration across  
 882 unimodal, bimodal, and higher-level fused modalities.

883 In contrast, late fusion combines independent predictions of each modality at the decision level to enhance robustness  
 884 and flexibility. Finally, MCIS (Yang et al., 2024a) leverages counterfactual reasoning during inference to mitigate bi-  
 885 ases and achieve robust decision-making without additional training, providing a complementary approach to handling  
 886 dataset biases in MSA.

## 891 D.2 PRE-TRAINED MODELS FOR MSA

894 In recent years, with the rapid development of pre-trained models, the research focus of multimodal sentiment analysis  
 895 has shifted from traditional cross-modal feature interaction to methods based on pre-trained models. For instance, (Yi  
 896 et al., 2024) combined CLIP and Timesformer architectures to generate discriminative textual prompts from video  
 897 content, thereby enhancing textual representations and enabling end-to-end sentiment analysis. (Yu et al., 2022)  
 898 employed prompt phrases to fine-tune frozen language models, obtaining coherent vision-language representations  
 899 to bridge the semantic gap and reduce parameter dependency. (Yu & Zhang, 2022) integrated visual information  
 900 into pre-trained language models and utilized prompt tuning to minimize the discrepancy between masked language  
 901 modeling and sentiment analysis tasks. (Yang et al., 2023) proposed a generative prompt model that constructs  
 902 multimodal prompts through an encoder-decoder architecture, reducing the reliance on annotated data. (Khan &  
 903 Fu, 2021) designed a dual-stream pre-trained model that converts images into auxiliary sentences via target-aware  
 904 transformation and single-pass non-autoregressive generation, thereby injecting multimodal information. (Ling et al.,  
 905 2022) introduced language-, vision-, and multimodal-specific pre-training tasks within an encoder-decoder framework  
 906 to facilitate cross-modal alignment. (Ye et al., 2022) proposed a sentiment-aware pre-training framework that captures  
 907 fine-grained emotional signals from data through cross-modal contrastive learning and additional sentiment objectives.

908 With the rise of large-scale models, the performance of multimodal sentiment analysis has been further significantly  
 909 improved. (Wang et al., 2024) leveraged large language models to generate rich contextual information for enhanced  
 910 sentiment understanding. (Feng et al., 2024) utilized large-scale vision-language models to integrate textual and vi-  
 911 sual information, mitigating the interference of image noise in sentiment classification. (Xiao et al., 2025) extracted  
 912 information from visual features and relied on large models to generate emotional causes and impressions to assist  
 913 analysis. (Shangguan et al., 2025) proposed a multimodal chain-of-thought reasoning distillation method to allevi-  
 914 ate model training challenges under resource-constrained conditions. Moreover, several studies have systematically  
 915 evaluated the performance of large models in multimodal sentiment analysis tasks: (Song, 2024) pioneered the use of  
 916 large text models and GPT-3.5 Turbo for zero-shot reasoning. (Zhang et al., 2025b) further expanded the evaluation  
 917 methods to include zero-shot reasoning, instruction tuning, and supervised fine-tuning. (Liu et al., 2025) established  
 918 extensive benchmarks in multilingual and multimodal scenarios, thoroughly validating the powerful cognitive and  
 919 reasoning capabilities of large models in this field.

918 **E PROMPT USED**919 **E.1 PROMPT FOR THE CONSTRUCTION OF REASONING PROCESS DATA**

920 During construction of the reasoning dataset, we employ two strong closed-source multimodal models, namely GPT-  
 921 4o and Gemini 2.5-Pro, to generate chain-of-thought rationales conditioned on multimodal inputs together with the  
 922 ground-truth concept labels and the sentiment label. Specifically, we use the prompt in Listing 1 to instruct the models  
 923 to analyze the final sentiment on the basis of the provided concept cues.

924 **Listing 1: Prompt for Sentiment CoT Generation**

925 You are an expert in human sentiment analysis. You will receive a video, the  
 926 speaker's utterance, some concept clues about the speaker, and the speaker's final  
 927 sentiment. I need you to analyze, based on the above, from different concept clues  
 928 why the speaker's sentiment is like this.

929 In addition to the given concept clues, you also need to summarize other concepts  
 930 that you consider helpful for analyzing the speaker's sentiment and, based on the  
 931 concept clues, perform the analysis. Finally, you need to output your thinking  
 932 process.

933 The given concept clues and the newly discovered concept clues need to be marked  
 934 with <concept>...</concept>.

935 For example: <think>The speaker's facial expression is  
 936 <facial\_expression>xxx</facial\_expression> ... The speaker's speaking style is  
 937 <speaking\_style>xxx</speaking\_style>, combined with semantic analysis it ... The  
 938 tone is <tone>xxx</tone>, etc.</think> <sentiment>xxx</sentiment>

939 The speaker's utterance: '<insert utterance>'.

940 Concept clues: <speaking\_style>xxx</speaking\_style>,  
 941 <facial\_expression>xxx</facial\_expression>, <tone>xxx</tone>.

942 Sentiment label: <sentiment>xxx</sentiment>.

943 After collecting the raw reasoning traces, we manually inspected and corrected them. Because instruction-following  
 944 behavior can sometimes produce strained explanations, and because the audio modality is occasionally missing and the  
 945 processing of visual information can be inaccurate, the raw rationales contained many issues. We therefore corrected  
 946 the data along the three perspectives described in Section 4.1. The correction pipeline proceeds in two stages: we first  
 947 apply an automatic correction with a closed-source model, using the prompt in Listing 2 to propose edits, and then  
 948 human annotators verify and amend the outputs.

949 **Listing 2: Prompt for CoT refinement**

950 Please correct and refine a chain-of-thought reasoning sample according to the  
 951 following requirements:

952 1) Correct factual errors. Compare all conceptual and factual statements against  
 953 the ground-truth label(s). If discrepancies exist, fix them and revise the  
 954 reasoning steps accordingly.

955 2) Remove redundancy. Where the analysis is repetitive, streamline it while  
 956 preserving the completeness and coherence of the reasoning.

957 3) Repair reasoning logic. Where the reasoning contains logical errors or overly  
 958 tenuous arguments, revise it to be sound while maintaining the overall reasoning  
 959 structure. You must not invent new concepts or add content that is not supported  
 960 by the provided information or labels.

961 **E.2 PROMPT FOR METHODS**

962 For the benchmark on CH-CEMS for the classification setting, we follow prior work (Zhang et al., 2025a; Luo et al.,  
 963 2025; Zhang et al., 2025b) and design standardized zero-shot inference and supervised fine-tuning prompts, such as  
 964 the five-class sentiment classification task shown in Listing 3, to evaluate both open-source and closed-source models.

972 Furthermore, we design a new prompt in Listing 4 based on the previous prompt for concept-guided reinforcement  
 973 learning, as comprehensively detailed in Sections 4.2 and Sections 4.3.  
 974

975 **Listing 3: Prompt for Zero-Shot Evaluation and SFT of MLLMs in Five-Class Sentiment Classification**

976 You are an expert in human sentiment analysis. You will be given a short video  
 977 represented by four images, along with the speaker's utterance and audio. Analyze  
 978 the speaker's sentiment based on these inputs.  
 979  
 980 At the end, output exactly one label from the set: [Positive, Weakly Positive,  
 981 Neutral, Weakly Negative, Negative].  
 982  
 983 The content spoken in the video is: <utterance>

984 **Listing 4: Prompt for Explainable Reasoning in Five-Class Sentiment (SFT Cold Start & Concept-Guided GRPO)**

985 You are an expert in human sentiment analysis. You will receive a short video  
 986 together with the speaker's utterance and audio.  
 987 Reason over all concepts you consider useful for inferring the speaker's sentiment.  
 988  
 989 At the end, choose exactly one label from [Positive, Weakly Positive, Neutral,  
 990 Weakly Negative, Negative].  
 991 Output your reasoning inside <think></think>, and mark the final sentiment using  
 992 <sentiment></sentiment>.  
 993  
 994 For example: <think>The speaker's facial expression is  
 995 <facial\_expression>xxx</facial\_expression> ... The speaking style is  
 996 <speaking\_style>xxx</speaking\_style> ... The tone is <tone>xxx</tone> ...</think>  
 997  
 998 Candidate labels for tone: [<insert tone labels>].  
 999 Candidate labels for facial expression: [<insert expression labels>].  
 1000 Candidate labels for speaking style: [<insert style labels>].  
 1001 Other useful concepts have no predefined list; analyze them using multimodal cues.  
 1002  
 1003 The utterance is: '<insert text>'

1004 **F CASE STUDIES**

1005 **G SIGNIFICANCE ANALYSIS**

1006  
 1007 Table 7: regression.  
 1008  
 1009

1011 <b>Models</b>	1012 <b>Acc2 (↑)</b>	1013 <b>F1_score (↑)</b>	1014 <b>Acc2_weak (↑)</b>	1015 <b>Corr (↑)</b>	1016 <b>MAE (↓)</b>
1017 MAG-BERT	61.67±2.82/73.13±3.75	61.78±3.16/75.02±3.31	50.67±3.68	56.04±1.50	28.37±0.44
1018 MISA	65.19±1.65/69.34±1.59	65.73±1.71/71.78±1.42	55.45±2.36	56.70±0.55	28.50±0.43
1019 MulT	62.53±1.80/71.74±3.23	62.92±1.98/73.89±2.86	52.18±2.79	56.00±0.75	28.47±0.48
1020 ALMT	62.96±1.68/72.65±2.35	63.29±1.89/74.70±2.04	52.37±2.77	56.09±0.21	28.95±0.15
1021 DLF	63.61±3.88/71.09±5.33	63.94±4.05/73.23±4.84	53.72±6.14	56.14±0.93	28.54±0.83

1022 **H IMPACT OF DATASET DISTRIBUTION ON RESULTS**

1023 **H.1 SCENARIO-WISE LABEL DISTRIBUTION AND MODEL PERFORMANCE**

1024 In multimodal sentiment analysis, emotional expression is strongly influenced by the interplay and integration of mul-  
 1025 tiple modalities. Due to the specific encoding conventions and presentation styles of different video types, identical  
 emotional expressions may convey distinct semantic meanings, which substantially increases model discriminative



Figure 6: Case study of comparison between MISA and CD-GRPO.

difficulty and induces performance variability. In this section, we present the accuracies of conventional SOTA approaches and multimodal large language models (MLLMs), specifically Qwen2.5-Omni-7B and MiniCPM-o-2.6, for three-class and five-class sentiment recognition across six scenarios: melodramas, vox pops, science broadcasting,

Table 8: 3-class.

Models	Acc	WF1	WP	F1	R	P
MAG-BERT	55.96 $\pm$ 1.03	55.29 $\pm$ 1.20	56.04 $\pm$ 0.90	56.34 $\pm$ 1.23	57.55 $\pm$ 1.34	56.47 $\pm$ 1.01
MISA	57.79 $\pm$ 1.23	57.18 $\pm$ 1.28	58.50 $\pm$ 1.27	57.95 $\pm$ 1.18	60.12 $\pm$ 0.88	57.77 $\pm$ 1.27
MuLT	56.42 $\pm$ 1.35	55.77 $\pm$ 1.47	56.68 $\pm$ 1.15	56.75 $\pm$ 1.49	58.49 $\pm$ 0.87	56.56 $\pm$ 1.58
ALMT	55.59 $\pm$ 1.21	54.85 $\pm$ 1.33	55.47 $\pm$ 1.34	55.95 $\pm$ 1.15	57.62 $\pm$ 1.22	55.58 $\pm$ 1.06
DLF	57.55 $\pm$ 1.69	57.12 $\pm$ 1.86	58.08 $\pm$ 2.13	58.10 $\pm$ 1.63	59.71 $\pm$ 1.63	58.06 $\pm$ 1.85
MiniCPM-V-4.0	52.68 $\pm$ 0.51	46.31 $\pm$ 0.69	56.09 $\pm$ 1.21	49.37 $\pm$ 0.60	56.25 $\pm$ 0.52	56.14 $\pm$ 0.95
Qwen2.5-VL-7B	61.91 $\pm$ 0.00	61.91 $\pm$ 3.00	70.69 $\pm$ 4.28	61.08 $\pm$ 0.00	65.58 $\pm$ 0.20	65.58 $\pm$ 0.20
MiniCPM-o-2.6	62.85 $\pm$ 0.90	62.76 $\pm$ 0.94	64.75 $\pm$ 0.87	62.93 $\pm$ 1.02	62.20 $\pm$ 1.04	65.54 $\pm$ 0.95
Qwen2.5-Omni-7B	64.98 $\pm$ 0.43	64.89 $\pm$ 0.44	66.48 $\pm$ 0.35	64.61 $\pm$ 0.50	63.12 $\pm$ 0.53	68.38 $\pm$ 0.39

Table 9: 5-class.

Models	Acc	WF1	WP	F1	R	P
MAG-BERT	42.77 $\pm$ 0.87	39.90 $\pm$ 1.32	43.94 $\pm$ 3.15	31.51 $\pm$ 3.20	34.44 $\pm$ 1.73	39.06 $\pm$ 7.02
MISA	43.72 $\pm$ 1.18	42.80 $\pm$ 0.97	47.90 $\pm$ 1.42	36.03 $\pm$ 0.96	37.12 $\pm$ 0.77	43.88 $\pm$ 2.07
MuLT	42.61 $\pm$ 1.81	40.84 $\pm$ 1.59	46.01 $\pm$ 1.74	33.20 $\pm$ 1.75	35.10 $\pm$ 0.99	42.52 $\pm$ 3.51
ALMT	41.94 $\pm$ 1.21	41.15 $\pm$ 1.24	44.68 $\pm$ 1.28	36.10 $\pm$ 1.18	36.48 $\pm$ 1.05	41.78 $\pm$ 1.12
DLF	44.09 $\pm$ 1.62	42.44 $\pm$ 1.43	45.92 $\pm$ 1.59	35.11 $\pm$ 1.93	36.62 $\pm$ 1.75	42.25 $\pm$ 1.54
MiniCPM-V-4.0	28.42 $\pm$ 0.20	22.80 $\pm$ 0.29	28.23 $\pm$ 0.30	18.98 $\pm$ 1.53	35.48 $\pm$ 3.11	17.85 $\pm$ 1.44
Qwen2.5-VL-7B	31.22 $\pm$ 0.00	31.22 $\pm$ 1.31	48.74 $\pm$ 1.86	31.35 $\pm$ 0.00	38.82 $\pm$ 0.43	38.82 $\pm$ 0.43
MiniCPM-o-2.6	40.89 $\pm$ 0.84	41.55 $\pm$ 0.82	44.24 $\pm$ 0.90	39.13 $\pm$ 3.55	41.95 $\pm$ 3.89	38.20 $\pm$ 3.33
Qwen2.5-Omni-7B	46.92 $\pm$ 0.79	46.65 $\pm$ 0.84	50.78 $\pm$ 0.93	42.04 $\pm$ 2.87	45.67 $\pm$ 3.24	44.88 $\pm$ 3.05

formal interviews, variety shows, and vlogs. The models were evaluated under random seeds 0–4, and the results are illustrated in Figure 7.

Based on the boxplot results, models perform particularly well in melodramas and vlogs, showing higher median accuracies, narrower interquartile ranges, and more concentrated data distributions, indicating strong robustness in these scenarios. In contrast, in vox pops, although the interquartile ranges are relatively narrow, the high number of outliers suggests limited generalization capability and reduced stability. In science-broadcasting programs, formal interviews, and variety shows, the diversity of content themes and the richness of emotional expressions contribute to greater fluctuations in prediction results, accompanied by markedly expanded box widths, reflecting increased task complexity. Importantly, in variety shows—where filming captures spontaneous reactions from multiple participants—emotional signals are highly heterogeneous and dispersed, resulting in the lowest accuracy among all scenarios. In formal interviews, limited sample sizes lead to the greatest variability in model performance. Across all scenarios, five-class tasks generally achieve lower accuracy than three-class tasks, primarily due to finer-grained categories, sparse sample distributions, ambiguous inter-class boundaries, and higher discriminative demands, which further amplify the overall complexity of sentiment analysis.

Overall, the type of video scenario exerts a substantial impact on model performance, and enhancing model stability and cross-scenario generalization remains a critical challenge in current research.

## H.2 EFFECT OF NUMBER OF ON-SCREEN SPEAKERS: SINGLE VS. MULTIPLE

In multimodal sentiment analysis, the interaction patterns among speakers constitute a critical factor affecting both task complexity and model performance. As the complexity of interactions increases, the difficulty of the task rises markedly. In this section, we report the accuracies of conventional SOTA approaches and multimodal large language models—specifically Qwen2.5-Omni-7B and MiniCPM-o-2.6—on three-class and five-class sentiment recognition tasks across single-person, multi-person without interaction, and multi-person with interaction scenarios. The evaluations were conducted under random seeds 0–4, as shown in Figure 8.

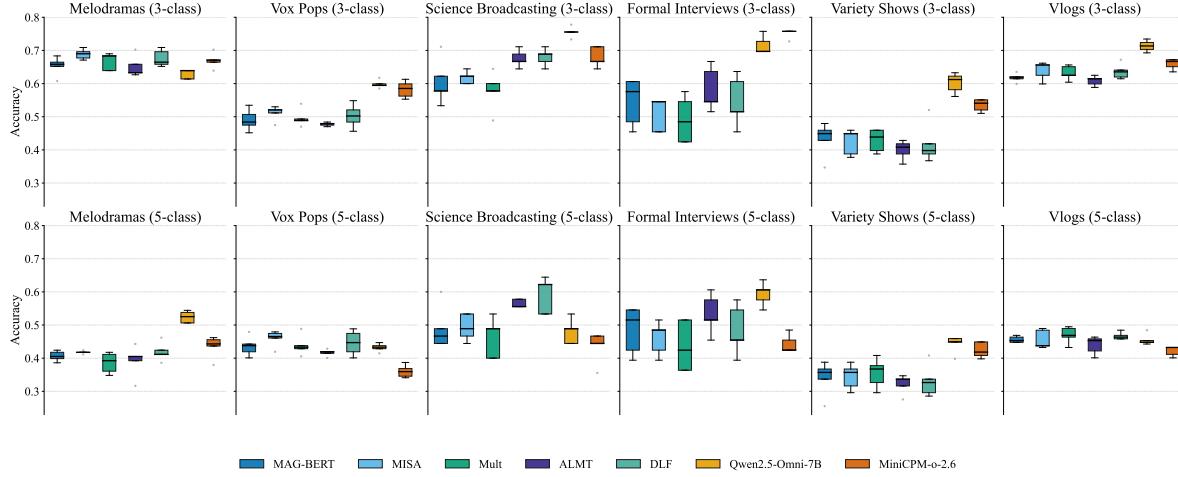


Figure 7: Accuracies of conventional SOTA approaches and multimodal large language models, specifically Qwen2.5-Omni-7B and MiniCPM-o-2.6, for three-class and five-class sentiment recognition across six scenarios—melodramas, vox pops, science broadcasting, formal interviews, variety shows, and vlogs—evaluated under random seeds 0–4.

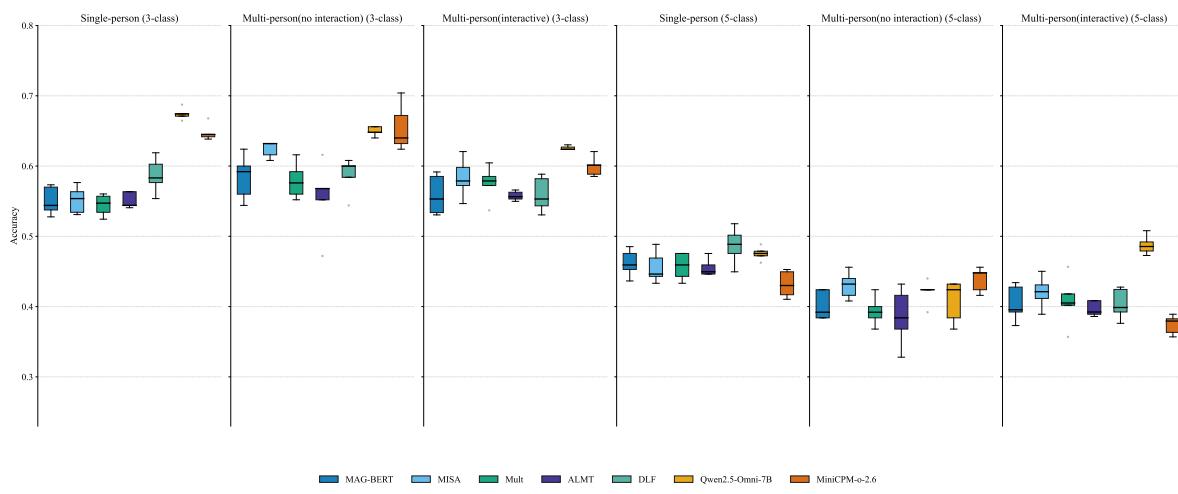


Figure 8: Accuracies of conventional SOTA approaches and multimodal large language models, specifically Qwen2.5-Omni-7B and MiniCPM-o-2.6, are reported for three-class and five-class sentiment recognition across single-person, multi-person without interaction, and multi-person with interaction scenarios, evaluated under random seeds 0–4.

In single-person scenarios, where emotional cues are clear and modal signals are consistent, model recognition is comparatively straightforward, resulting in higher overall accuracy. Small models demonstrate compact accuracy distributions without outliers, reflecting strong stability and low sensitivity to random seeds. By contrast, Qwen2.5-Omni-7B, while generally exhibiting concentrated accuracy distributions, generates a relatively high number of outliers, indicating potential performance variability under certain conditions. In multi-person scenarios without interaction, the model is required to process multiple independent emotional signals simultaneously, thereby increasing recognition difficulty, reducing overall accuracy, and resulting in more dispersed distributions with elongated box ranges. Both the ALMT and DLF methods display outliers, reflecting heightened sensitivity to random seeds, greater training variability, and relatively lower stability. In multi-person scenarios with interaction, the model must capture the dynamic transmission, conflict, and integration of emotions. This requirement imposes higher demands on contextual awareness and relational reasoning, making the task highly complex. In this context, the box range expands significantly. The MULT method exhibits prominent outliers in five-class tasks, accompanied by substantial performance fluctuations, further indicating reduced stability under complex interactive conditions. Nevertheless, Qwen2.5-Omni-7B maintains relatively high accuracy in this scenario, demonstrating its robustness and effectiveness in handling complex contexts.

1188 Furthermore, across different speaker configurations and for both three-class and five-class tasks, MAG-BERT and  
 1189 MISA exhibit no outliers, demonstrating strong robustness, generalization capability, and stability. Overall, three-  
 1190 class tasks consistently achieve higher accuracy than five-class tasks. In three-class settings, large models generally  
 1191 outperform small models, exhibiting higher accuracy distributions. In contrast, in five-class tasks, the advantage of  
 1192 large models diminishes due to finer-grained categories and more ambiguous inter-class boundaries, and in certain  
 1193 cases, they may perform worse than small models.

1194 In summary, with increasing interaction complexity and a greater number of classification categories, model perfor-  
 1195 mance exhibits more pronounced fluctuations, accompanied by heightened sensitivity to random seeds. These obser-  
 1196 vations offer valuable insights for model optimization and methodological selection in complex speaker scenarios.  
 1197

## 1198 I DATA PRIVACY AND CONTENT CONSIDERATIONS

1200 Our dataset has undergone systematic, meticulous screening and includes characters and dialogues from six distinct do-  
 1201 mains: vox pops, variety shows, melodramas, formal interviews, vlogs, and science-broadcasting programs. Although  
 1202 some material originates from real-world settings, we applied comprehensive anonymization and de-identification pro-  
 1203 cedures to all content involving personal identities to eliminate the risk of privacy breaches and to ensure compliance  
 1204 with applicable privacy laws and regulations. Furthermore, we performed a thorough content review and removed any  
 1205 samples that could be offensive or controversial, thereby guaranteeing social appropriateness. The dataset is intended  
 1206 for the study of character dialogues and interactions, aiming to enable comprehensive, multimodal sentiment analysis  
 1207 through holistic character understanding. The data are used solely for scientific research and do not violate legal rights.  
 1208 All data-processing and research procedures adhere strictly to established ethical guidelines.

## 1209 J LLM USAGE

1210 In our work, both the reasoning process data contribution and methods research use large language model as illustrated  
 1211 in Section 4. In addition to this, we use LLM such as GPT-5 to polish grammar and generate draft icons for figures.  
 1212 All reported experimental results are obtained independently of these editorial aids.

## 1213 K REPRODUCIBILITY

1214 To ensure the reproducibility of the benchmark studies on our dataset, we have provided comprehensive documentation  
 1215 and resources. The data sources, along with the complete data processing, filtering pipeline, and detailed annotation  
 1216 guidelines, are detailed in Appendix A. The experimental setup for the main evaluations, including model implemen-  
 1217 tations, hyperparameters, and evaluation protocols, is elaborated in Section 5. Furthermore, the prompts used for large  
 1218 language model-based experiments are fully listed in Appendix E for reference. The code for data loading, benchmark  
 1219 models, and evaluation, alongside a sample of the processed dataset, is available via an anonymized repository (link  
 1220 provided in the abstract). To account for randomness, we report the standard deviation across multiple random seeds  
 1221 in Appendix G. Upon acceptance of this paper, we commit to releasing the full dataset publicly to further support  
 1222 research in the community.