

# Enhancing Multimodal Sentiment Analysis for Missing Modality through Self-Distillation and Unified Modality Cross-Attention

Yuzhe Weng<sup>\*†</sup>, Haotian Wang<sup>\*</sup>, Tian Gao<sup>†</sup>, Kewei Li<sup>\*</sup>, Shutong Niu<sup>\*</sup>, Jun Du<sup>\*§‡</sup>

<sup>\*</sup>NERC-SLIP, University of Science and Technology of China (USTC), Hefei, China

<sup>§</sup>MoE Key Laboratory of Brain-inspired Intelligent Perception and Cognition, USTC, Hefei, China

<sup>†</sup>iFlytek Research, Hefei, China

{yzweng, az1522702192, keweili12, niust}@mail.ustc.edu.cn, tiangao5@iflytek.com, jundu@ustc.edu.cn

**Abstract**—In multimodal sentiment analysis, collecting text data is often more challenging than video or audio due to higher annotation costs and inconsistent automatic speech recognition (ASR) quality. To address this challenge, our study has developed a robust model that effectively integrates multimodal sentiment information, even in the absence of text modality. Specifically, we have developed a Double-Flow Self-Distillation Framework, including Unified Modality Cross-Attention (UMCA) and Modality Imagination Autoencoder (MIA), which excels at processing both scenarios with complete modalities and those with missing text modality. In detail, when the text modality is missing, our framework uses the LLM-based model to simulate the text representation from the audio modality, while the MIA module supplements information from the other two modalities to make the simulated text representation similar to the real text representation. To further align the simulated and real representations, and to enable the model to capture the continuous nature of sample orders in sentiment valence regression tasks, we have also introduced the Rank-N Contrast (RNC) loss function. When testing on the CMU-MOSEI, our model achieved outstanding performance on MAE and significantly outperformed other models when text modality is missing. The code is available at: <https://github.com/WarmCongee/SDUMC>.

**Index Terms**—Multimodal Fusion, Sentiment Analysis, Self Distillation, Missing Modality

## I. INTRODUCTION

Multimodal sentiment analysis has garnered significant attention [1]–[4], with numerous studies focusing on the representation [5]–[8], alignment [9], [10], and fusion [11]–[14] of sentiment information from various modalities. The absence of certain modalities during training and inference significantly influences the alignment and fusion of multimodal information, prompting increased research efforts [15]–[19] aimed at enhancing robustness under these conditions. Traditionally, research [15] has classified the absence of text, audio, and visual modalities as either complete or frame-level missing, often employing partial word masks to mimic text frame-level absence. However, this simulated absence differs markedly from typical real-world scenarios where text is either completely missing or substituted with low-quality ASR transcriptions. Given the text modality’s superior performance [13] in certain tasks,

previous studies [20], [21] have relied on manually annotated transcriptions or the development of ASR/AVSR systems. We contend that it is unnecessary to reconstruct the final text when the text modality is missing. Instead, the model can simply use the hidden states derived from the audio or visual modalities to approximate the semantic space of the missing text modality.

Theoretically, textual information can be highly abstracted from audio and visual modalities. With advancements in Large Language Models (LLM), some studies [22], [23] have employed LLMs as decoders to address ASR tasks. For instance, SLAM-ASR [22] projects audio representations and inputs them with prompts into an LLM to generate ASR results. Consequently, it is natural to infer that simulated text representations can be derived from the projected audio representations via LLMs, facilitating unified multimodal sentiment analysis.

Our research proposes a Double-Flow Self-Distillation Framework, which includes the Modality Imagination Autoencoder (MIA) and Unified Modality Cross-Attention (UMCA). This unified model can address both complete and text-missing multimodal sentiment analysis scenarios by simply activating or deactivating the MIA during inference. The main contributions of our work are: (1) We design UMCA for comprehensive multimodal fusion to extract sentiment-related information. (2) Our method uses the LLM-based model Vicuna [24] to generate simulated text representations from audio representations’ projection when text modality is missing, with MIA complementing the missing information. (3) We train a unified multimodal sentiment analysis network using self-distillation combined with MIA activation and deactivation while achieving distance-space alignment of complete and missing modality representations through MKD Loss and RNC Loss [25]. (4) Our method achieves optimal performance on the CMU-MOSEI [26] across multiple metrics for complete modality inference, and significantly improving performance in text-missing scenarios compared to other methods.

## II. METHOD

In this section, we discuss our proposed Double-Flow Self-Distillation Framework for the multimodal sentiment analysis. The overall architecture of our system is depicted in Fig. 1.

<sup>‡</sup> Corresponding author

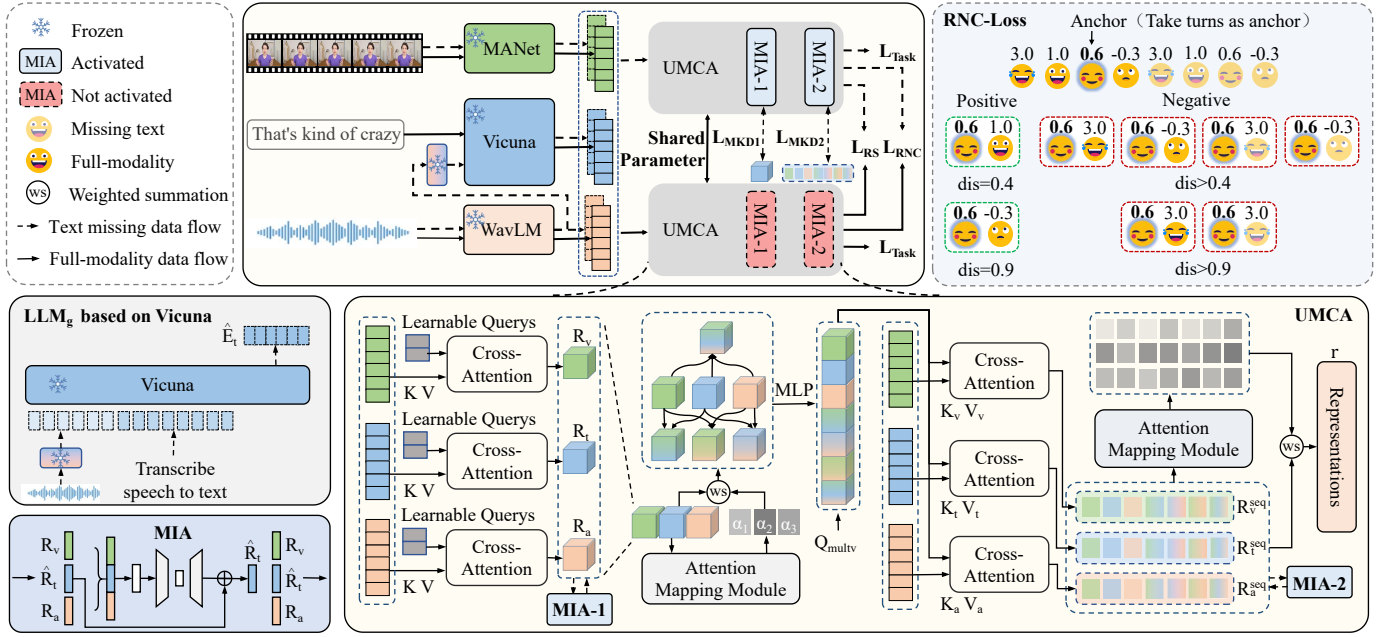


Fig. 1. Overall architecture of Double-Flow Self-Distillation Framework for modality missing. In the figure, the middle part of the upper row shows the training process of the framework. The lower row shows the specific network structure of each module.

### A. Unified Modality Cross-Attention

As shown in the bottom right corner of the Fig. 1, the UMCA model is divided into two main stages. And we set up the cross-attention module, which is applied in both stages. The formulas are as follows:

$$K_m = \text{Tanh}(\mathbf{W}V_m + \mathbf{b}) \quad (1)$$

$$R_m = \text{Softmax}\left(\frac{QK_m^T}{\tau}\right)V_m \quad (2)$$

$m \in \{a, v, t\}$ , where a, v, and t refer to the three modalities of audio, vision, and text.  $\tau$  denotes the scaling parameter that ensures the stability of the gradient of softmax.  $K_m$  and  $V_m$  are the key and value sequences from the  $m$  modality,  $Q$  is the query involved in the cross-attention, and  $R_m$  is the representation of the  $m$  modality in the same dimension as  $Q$ .

The first stage is used to obtain the query embedding of the multi-view modalities combinations. First, each modality representation is mapped to the same dimension  $E_m \in \mathbb{R}^{S \times D}$  by MLP, where  $S$  represents the sequence length and  $D$  denotes the dimension determined by a hyperparameter. We initialize the Gaussian-distributed learnable query ( $Q_m$ ) for each modality during model initialization, where  $Q_m \in \mathbb{R}^{S_1 \times D}$ , with  $S_1$  set to 1. For each modality, the original modal representations are mapped to the key ( $K_m$ ) and value ( $V_m$ ). According to the equation, we use the cross-attention module to obtain the representations  $R_m$  for each modality. The Attention-guided Feature Gathering (AFG) module [27] concatenates the three modal representations and feeds them into an MLP to obtain attention weights. These weights are used to compute the weighted sums of unimodal, bimodal,

and trimodal representations, producing a multi-view query ( $Q_{multv}$ ) with seven combinations. For instance, in a bimodal scenario, we set the unused modality's weight to zero while maintaining the weights of the other two modalities. This results in a weighted sum that forms the bimodal query.

In the second stage, the multi-view query obtained in the previous stage is used as the query ( $Q_{multv}$ ). The original modal representations are mapped as key ( $K_m$ ) and value ( $V_m$ ). The modal representations  $R_m^{\text{seq}}$  are obtained by cross-attention, and the final representation  $r$  is obtained by attention weighted summing  $R_m^{\text{seq}}$  by AFG module, and regression is performed to obtain the sentiment valence.

### B. Missing Modality Imagination

1) *LLM-based Text Representation Simulation*: SLAM-ASR [22] uses WavLM as an audio encoder and LLM-based model Vicuna [24] as a decoder. Then the method trained the projector between them and generated ASR results from the audio encoder representations directly using the generative power of the LLM. Based on this, we work on unifying the representation of text modality presence and missing. Definitions are as follows:

$$E_t = \text{LLM}_f(\text{tokenizer}(T)) \quad (3)$$

$$\hat{E}_t = \text{LLM}_g(\text{cat}(\text{proj}(E_a), \text{tokenizer}(P))) \quad (4)$$

In the definitions,  $T$  denotes the text transcribed from video content. The symbol  $E_t$  refers to the real text representation obtained from the text modality. The function  $\text{LLM}_f(\cdot)$  is employed for the direct forward pass to obtain the hidden state. Additionally, the symbol  $\hat{E}_t$  denotes the simulated text representation generated by the LLM in scenarios where the text

modality is missing.  $E_a$  refers to the representation obtained from the audio modality and  $P$  refers to the input prompt for LLM. The notation  $\text{proj}(\cdot)$  is used to refer to the pre-trained projector network within the SLAM-ASR framework. As shown in the middle of the left side of Fig. 1,  $\text{LLM}_g(\cdot)$  refers to the hidden state derived from the generated results.

2) *Modality Imagination Autoencoder*: We use the pre-trained projector and Vicuna to obtain real text representations when text modality is present and generate simulated ones when text modality is missing. However, the gap between real and simulated representations can negatively impact performance if used directly for training. To address this, we use the Residual Autoencoder [28] to construct the missing modality imagination autoencoder, which activates only when the text modality is missing. The structure of MIA is depicted in the bottom left corner of Fig. 1. The definition is as follows:

$$H_t = f(\mathbf{W}^{(1)} \text{cat}(R_v, R_a, \hat{R}_t) + \mathbf{b}^{(1)}) \quad (5)$$

$$\hat{R}_t = \hat{R}_t + f(\mathbf{W}^{(2)} H_t + \mathbf{b}^{(2)}) \quad (6)$$

$R_v, R_a, \hat{R}_t \in \mathbb{R}^{S \times D}$ , where  $S$  is the sequence length and  $D$  is the unified dimension of modality representations.  $H_t \in \mathbb{R}^{S \times D'}$ , where  $D'$  is the dimension of the intermediate hidden state. We replace the original  $\hat{R}_t$  with the reconstructed  $\hat{R}_t$ . As shown in Fig. 1, since there are two MIA modules, the input is  $R_m$  in the MIA-1 module and  $R_m^{\text{seq}}$  in the MIA-2 module.

### C. Modality Missing Self-Distillation

We design a self-distillation framework to handle cases where all modalities are present and cases where the text modality is missing. The network has two data flows: one with full modalities input (visual, audio, text) and one with only visual and audio modalities input.

During training, we distill text modality representation knowledge from the complete multimodal flow to the scenario where the text modality is missing. To achieve this, we have designed the Modality Knowledge Distillation Loss ( $\mathcal{L}_{\text{MKD}}$ ) and Representation Similarity Loss ( $\mathcal{L}_{\text{RS}}$ ). To capture the continuous nature of sample orders in the sentiment valence regression task, we introduced the Regression Rank-N Contrast Loss ( $\mathcal{L}_{\text{RNC}}$ ). Additionally, the Mean Squared Error loss is used as the sentiment valence regression task loss ( $\mathcal{L}_{\text{Task}}$ ) for sentiment analysis. To ensure the model acquires these various capabilities, we use the weighted sum of these losses as the final loss function. The overall loss is defined as follows:

$$\mathcal{L} = \mathcal{L}_{\text{Task}} + \alpha \mathcal{L}_{\text{MKD1}} + \beta \mathcal{L}_{\text{MKD2}} + \gamma \mathcal{L}_{\text{RS}} + \delta \mathcal{L}_{\text{RNC}} \quad (7)$$

The  $\alpha$  to  $\delta$  here defines the weights of the various losses, with specific values displayed in our open-sourced code.

1) *Sentiment Valence Regression Loss*: This is a regression task for predicting sentiment valence. Thus, the most important loss is the MSE loss between the true and predicted valence, defined as follows:

$$\mathcal{L}_{\text{Task}} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (8)$$

Here,  $y_i$  is the valence label, and  $\hat{y}_i$  is the predicted valence.

2) *Modality Knowledge Distillation Loss*: The objective of  $\mathcal{L}_{\text{MKD}}$  is to align the simulated text modality representations, produced by the MIA, as similar as possible with the real text modality representations to enhance the model's robustness when text modality is missing. Specifically, we detach the real text modality representation and then compute the RMSE loss between it and the simulated text modality representation from MIA. Consequently, it should be highlighted that the gradient of  $\mathcal{L}_{\text{MKD}}$  will only back-propagate through text missing data flow. Considering the presence of two MIA modules, we calculate two MKD losses, which are defined as follows:

$$\mathcal{L}_{\text{MKD}} = \sqrt{\sum_{i=1}^N \frac{(R_{i(\text{detached})} - \hat{R}_i)^2}{N}} \quad (9)$$

During the process, we need to calculate  $\mathcal{L}_{\text{MKD1}}$  and  $\mathcal{L}_{\text{MKD2}}$ . First, we detach  $R_t$  and use it along with  $\hat{R}_t$  in the formula to calculate  $\mathcal{L}_{\text{MKD1}}$ . Then, we detach  $R_t^{\text{seq}}$  and use it along with  $\hat{R}_t^{\text{seq}}$  in the formula to calculate  $\mathcal{L}_{\text{MKD2}}$ .

3) *Representation Similarity Loss*: In order to make the representations entering the final regression head as similar as possible when the modality is complete and the text modality is missing, we design a representation similarity loss  $\mathcal{L}_{\text{RS}}$ .

$$\mathcal{L}_{\text{RS}} = \sqrt{\sum_{i=1}^N \frac{(r_i - \hat{r}_i)^2}{N}} \quad (10)$$

4) *Regression Rank-N Contrast Loss*: For the sentiment valence regression task, the RNC loss has two goals. First, it brings the representations of two data flows of the same data similar. Second, it allows the model to learn and align the representation distance spaces correctly for both the missing text modality and the complete modalities. The RNC loss is defined as follows:

$$\mathcal{L}_{\text{RNC}}^i = \sum_{j=1, j \neq i}^{2N} \log \frac{\exp(\text{sim}(r_i, r_j)/\tau)}{\sum_{r_k \in S_{i,j}} \exp(\text{sim}(r_i, r_k)/\tau)} \quad (11)$$

$$\mathcal{L}_{\text{RNC}} = \frac{1}{2N} \sum_{i=1}^{2N} \frac{-1}{2N-1} \mathcal{L}_{\text{RNC}}^i \quad (12)$$

Here,  $N$  is the batch size, and  $r$  refers to the final representation. The  $\text{sim}(\cdot)$  function computes the L2 distance.  $S_{i,j}$  refers to the set of final representations of all other data in the batch whose distance from the valence label of anchor  $i$  is greater than or equal to the distance between the labels of  $i$  and  $j$ .

## III. EXPERIMENTS

### A. Dataset and Metric

We test our model on the CMU-MOSEI [26] dataset, which includes 22,856 videos: 16,326 for training, 1,871 for validation, and 4,659 for testing. Each sample is labeled with a sentiment valence from -3 (strongly negative) to +3 (strongly positive). As is common in studies using this dataset, we primarily use Mean Absolute Error (MAE) and prediction Accuracy (ACC) as metrics to evaluate model performance.

## B. Implementation Details

In our experiments, we use MANet [29] as the visual encoder, WavLM-Large [30] as the audio encoder, and Vicuna-7B [24] as the text encoder. The dimension of the visual representation obtained from the visual encoder is 1024 frame-level features. And we use the output of the 20th hidden layer of WavLM-Large as the audio representation and the output of the last hidden layer to obtain the simulated text representation. When text is present, the sum of the last four hidden layers of Vicuna is used as the text representation. When the text modality is missing, the output of the last hidden layer of the audio encoder is fed into the pre-trained Projector and Vicuna to generate the hidden state, and the penultimate fourth layer of this hidden state is taken as the simulated text representation.

## C. Ablation Study

To evaluate the effectiveness of the missing modality imagination module, the text modality missing self-distillation, and the loss functions, we design a series of ablation experiments to illustrate the role of each component in the architecture. The results are shown in Table I, where  $LLM_g$  indicates that the LLM-based model generated text representations when the text modality is missing. When  $LLM_g$  is not used, the model does not use the text modality representations simulated by LLM, but only uses the audio and visual representations for modality imagination and inference.

TABLE I  
ABLATION STUDIES FOR EACH MODULE

$LLM_g$	MIA	$\mathcal{L}_{MKD}$	$\mathcal{L}_{RS}$	$\mathcal{L}_{RNC}$	w/o text		w ground truth	
					MAE↓	ACC↑	MAE↓	ACC↑
✗	✓	✓	✓	✓	0.584	82.47	0.522	86.57
✓	✗	✗	✓	✓	0.572	83.81	0.521	87.45
✓	✗	✓	✓	✓	0.560	83.73	0.515	86.95
✓	✓	✗	✓	✓	0.563	83.90	0.513	87.72
✓	✓	✓	✗	✓	0.552	83.84	0.510	87.53
✓	✓	✓	✓	✗	0.550	84.25	0.508	87.25
✓	✓	✓	✓	✓	<b>0.550</b>	<b>84.28</b>	<b>0.506</b>	87.64

Using the MAE metric for early stopping, we found that self-distillation significantly improves network performance, with or without the text modality. This demonstrates that our network ensures strong performance even with missing text.

We visualize the similarity matrix between the complete modality representation and the representation with the text modality missing for the same batch of data. The left panel displays the L2 similarity matrix between the complete modality case and the missing text case, arranged by label values. The right panel shows the similarity matrix after removing MIA and distillation losses. Obviously, our model effectively aligns representations under modal integrity and absence, learning the corresponding representation distance based on label distance.

## D. Overall Comparison

We compare our method with other high-performing approaches that have achieved excellent performance or studied modality missing. The approaches listed in Table II explore robustness against missing modalities using various methods:

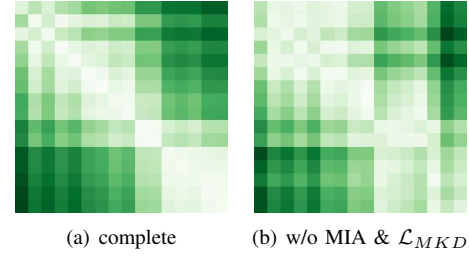


Fig. 2. Feature similarity visualization sorted by labels.

TABLE II  
METHODS PERFORMANCE COMPARISON TABLE.  $\Delta MAE$  AND  $\Delta ACC$  REPRESENT THE ABSOLUTE DIFFERENCE IN MODEL PERFORMANCE BETWEEN COMPLETE MODALITIES AND WHEN THE TEXT MODALITY IS MISSING. A SMALLER ABSOLUTE VALUE INDICATES BETTER ROBUSTNESS TO MODALITY MISSING.

Model	w text		w/o text		performance gap	
	MAE↓	ACC↑	MAE↓	ACC↑	$\Delta MAE↓$	$\Delta ACC↓$
UniMSE [31]	0.523	87.5	-	-	-	-
AMB [32]	0.536	85.8	-	-	-	-
ALMT [13]	0.526	-	-	-	-	-
MInD [33]	0.529	-	0.841	-	0.312	-
GMC [34]	-	80.0	-	66.8	-	13.2
GCN [35]	-	86.4	-	66.5	-	19.9
MMIN [36]	-	85.5	-	70.1	-	15.4
GMD [37]	-	87.1	-	72.6	-	14.5
Ours	<b>0.506</b>	<b>87.6</b>	<b>0.550</b>	<b>84.2</b>	<b>0.044</b>	<b>3.4</b>

contrastive learning (GMC), graph neural networks (GCN), modality dependence reduction via gradient-guided decoupling (GMD), missing modality reconstruction (MMIN), and modality information separation (MInD), et al. We referenced the experimental results of these models obtained from articles on the GMD method. The results are shown in Table II.

First, compared to studies conducted under complete modalities such as UniMSE [31], AMB [32], and ALMT [13], our model achieves significantly better MAE and ACC performance. Second, when the text modality is missing, our model's MAE increases by only 0.044, and the accuracy decreases by just 3.4%, outperforming other studies on modality-missing robustness. Additionally, the performance decay of our model is an order of magnitude less than that of other models when the text modality is missing, demonstrating its superior design.

## IV. CONCLUSIONS

This work addresses the challenge of high cost and heavy reliance on the text modality. Specifically, we designed the Double-Flow Self-Distillation Framework with Unified Modality Cross-Attention as the main network structure, combined with Modality Imagination Autoencoder to simulate missing text modality. Additionally, we developed a series of loss functions to enhance the model's performance in multimodal sentiment analysis, ensuring robustness both when all modalities are present and when the text modality is missing.

## V. ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China under Grant 62171427.

## REFERENCES

- [1] Z. Lian, H. Sun, L. Sun, K. Chen, M. Xu, K. Wang, K. Xu, Y. He, Y. Li, J. Zhao *et al.*, “Mer 2023: Multi-label learning, modality robustness, and semi-supervised learning,” in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 9610–9614.
- [2] Z. Lian, L. Sun, Y. Ren, H. Gu, H. Sun, L. Chen, B. Liu, and J. Tao, “Merbench: A unified evaluation benchmark for multimodal emotion recognition,” *arXiv preprint arXiv:2401.03429*, 2024.
- [3] S. Zhao, G. Jia, J. Yang, G. Ding, and K. Keutzer, “Emotion recognition from multiple modalities: Fundamentals and methodologies,” *IEEE Signal Processing Magazine*, vol. 38, no. 6, pp. 59–73, 2021.
- [4] U. Asiya and V. Kiran, “A novel multimodal speech emotion recognition system,” in *2022 Third International Conference on Intelligent Computing Instrumentation and Control Technologies (ICICICT)*. IEEE, 2022, pp. 327–332.
- [5] W. Yu, H. Xu, Z. Yuan, and J. Wu, “Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 12, 2021, pp. 10790–10797.
- [6] Z. Ma, Z. Zheng, J. Ye, J. Li, Z. Gao, S. Zhang, and X. Chen, “emotion2vec: Self-supervised pre-training for speech emotion representation,” *arXiv preprint arXiv:2312.15185*, 2023.
- [7] W. Chen, X. Xing, P. Chen, and X. Xu, “Vesper: A compact and effective pretrained model for speech emotion recognition,” *IEEE Transactions on Affective Computing*, 2024.
- [8] X. Wang and C. Zong, “Distributed representations of emotion categories in emotion space,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 2364–2375.
- [9] R. Fan, H. Liu, Y. Li, P. Guo, G. Wang, and T. Wang, “Atta-net: Attention aggregation network for audio-visual emotion recognition,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 8030–8034.
- [10] Q. Huang, P. Cai, T. Nie, and J. Zeng, “Clip-msa: Incorporating inter-modal dynamics and common knowledge to multimodal sentiment analysis with clip,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 8145–8149.
- [11] Y.-H. H. Tsai, S. Bai, P. P. Liang, J. Z. Kolter, L.-P. Morency, and R. Salakhutdinov, “Multimodal transformer for unaligned multimodal language sequences,” in *Proceedings of the conference. Association for computational linguistics. Meeting*, vol. 2019. NIH Public Access, 2019, p. 6558.
- [12] Y. H. Tsai, P. P. Liang, A. Zadeh, L. Morency, and R. Salakhutdinov, “Learning factorized multimodal representations,” in *ICLR*, 2019.
- [13] H. Zhang, Y. Wang, G. Yin, K. Liu, Y. Liu, and T. Yu, “Learning language-guided adaptive hyper-modality representation for multimodal sentiment analysis,” in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, Dec. 2023, pp. 756–767.
- [14] D. Zong, C. Ding, B. Li, J. Li, K. Zheng, and Q. Zhou, “Acformer: An aligned and compact transformer for multimodal sentiment analysis,” in *Proceedings of the 31st ACM International Conference on Multimedia*, ser. MM ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 833–842.
- [15] M. Li, D. Yang, X. Zhao, S. Wang, Y. Wang, K. Yang, M. Sun, D. Kou, Z. Qian, and L. Zhang, “Correlation-decoupled knowledge distillation for multimodal sentiment analysis with incomplete modalities,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 12458–12468.
- [16] H. Zuo, R. Liu, J. Zhao, G. Gao, and H. Li, “Exploiting modality-invariant feature for robust multimodal emotion recognition with missing modalities,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [17] X. Zhao, S. Poria, X. Li, Y. Chen, and B. Tang, “Toward robust multimodal learning using multimodal foundational models,” *arXiv preprint arXiv:2401.13697*, 2024.
- [18] Y. Dai, H. Chen, J. Du, R. Wang, S. Chen, H. Wang, and C.-H. Lee, “A study of dropout-induced modality bias on robustness to missing video frames for audio-visual speech recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 27445–27455.
- [19] J. Jang, Y. Wang, and C. Kim, “Towards robust multimodal prompting with missing modalities,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 8070–8074.
- [20] Y. Li, P. Bell, and C. Lai, “Speech emotion recognition with asr transcripts: A comprehensive study on word error rate and fusion techniques,” *arXiv preprint arXiv:2406.08353*, 2024.
- [21] S. Shon, P. Brusco, J. Pan, K. J. Han, and S. Watanabe, “Leveraging pre-trained language model for speech sentiment analysis,” in *Interspeech 2021*, 2021, pp. 3420–3424.
- [22] Z. Ma, G. Yang, Y. Yang, Z. Gao, J. Wang, Z. Du, F. Yu, Q. Chen, S. Zheng, S. Zhang *et al.*, “An embarrassingly simple approach for llm with strong asr capacity,” *arXiv preprint arXiv:2402.08846*, 2024.
- [23] X. Geng, T. Xu, K. Wei, B. Mu, H. Xue, H. Wang, Y. Li, P. Guo, Y. Dai, L. Li *et al.*, “Unveiling the potential of llm-based asr on chinese open-source datasets,” *arXiv preprint arXiv:2405.02132*, 2024.
- [24] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez, I. Stoica, and E. P. Xing, “Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality,” March 2023.
- [25] K. Zha, P. Cao, J. Son, Y. Yang, and D. Katabi, “Rank-n-contrast: learning continuous representations for regression,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [26] A. B. Zadeh, P. P. Liang, S. Poria, E. Cambria, and L.-P. Morency, “Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018, pp. 2236–2246.
- [27] H. Wang, Y. Xi, H. Chen, J. Du, Y. Song, Q. Wang, H. Zhou, C. Wang, J. Ma, P. Hu *et al.*, “Hierarchical audio-visual information fusion with multi-label joint decoding for mer 2023,” in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 9531–9535.
- [28] L. Tran, X. Liu, J. Zhou, and R. Jin, “Missing modalities imputation via cascaded residual autoencoder,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1405–1414.
- [29] Z. Zhao, Q. Liu, and S. Wang, “Learning deep global multi-scale and local attention features for facial expression recognition in the wild,” *IEEE Transactions on Image Processing*, vol. 30, pp. 6544–6556, 2021.
- [30] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [31] G. Hu, T.-E. Lin, Y. Zhao, G. Lu, Y. Wu, and Y. Li, “UniMSE: Towards unified multimodal sentiment analysis and emotion recognition,” in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Y. Goldberg, Z. Kozareva, and Y. Zhang, Eds. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics, Dec. 2022, pp. 7837–7851.
- [32] O. S. Chlapanis, G. Paraskevopoulos, and A. Potamianos, “Adapted multimodal bert with layer-wise fusion for sentiment analysis,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [33] W. Dai, X. Li, P. Hu, Z. Wang, J. Qi, J. Peng, and Y. Zhou, “Mind: Improving multimodal sentiment analysis via multimodal information disentanglement,” *arXiv preprint arXiv:2401.11818*, 2024.
- [34] P. Poklukar, M. Vasco, H. Yin, F. S. Melo, A. Paiva, and D. Kragic, “Geometric multimodal contrastive representation learning,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 17782–17800.
- [35] Z. Lian, L. Chen, L. Sun, B. Liu, and J. Tao, “Gcnet: Graph completion network for incomplete multimodal learning in conversation,” *IEEE Transactions on pattern analysis and machine intelligence*, vol. 45, no. 7, pp. 8419–8432, 2023.
- [36] J. Zhao, R. Li, and Q. Jin, “Missing modality imagination network for emotion recognition with uncertain missing modalities,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 2608–2618.
- [37] H. Wang, S. Luo, G. Hu, and J. Zhang, “Gradient-guided modality decoupling for missing-modality robustness,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 14, 2024, pp. 15483–15491.