

Article

Not peer-reviewed version

Multi-level Conflict-Aware Network for Multi-modal Sentiment Analysis

Yubo Gao , [Haotian Wu](#) , [Lei Zhang](#) *

Posted Date: 14 February 2025

doi: [10.20944/preprints202502.1056.v1](https://doi.org/10.20944/preprints202502.1056.v1)

Keywords: Multimodal sentiment analysis; Multi-level alignment; Multi-level conflict modeling



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Article

Multi-Level Conflict-Aware Network for Multi-Modal Sentiment Analysis

Yubo Gao ^{1,†}, Haotian Wu ^{2,†} and Lei Zhang ^{2,*}

¹ North University of China

² Beijing Jiaotong University

* Correspondence: zhlei@bjtu.edu.cn

† These authors contributed equally to this work.

Abstract: Multimodal Sentiment Analysis (MSA) aims to recognize human emotions by exploiting textual, acoustic, and visual modalities, and thus how to make full use of the interactions between different modalities is a central challenge of MSA. Interaction contains alignment and conflict aspects. Current works mainly emphasize alignment and the inherent differences between unimodal modalities, neglecting the fact that there are also potential conflicts between bimodal combinations. Additionally, multi-task learning-based conflict modeling methods often rely on the unstable generated labels. To address these challenges, we propose a novel multi-level conflict-aware network (MCAN) for multimodal sentiment analysis, which progressively segregates alignment and conflict constituents from unimodal and bimodal representations, and further exploits the conflict constituents with the conflict modeling branch. In the conflict modeling branch, we conduct discrepancy constraints at both the representation and predicted output levels, avoiding dependence on the generated labels. Experimental results on the CMU-MOSI and CMU-MOSEI datasets demonstrate the effectiveness of the proposed MCAN.

Keywords: multimodal sentiment analysis; multi-level alignment; multi-level conflict modeling

1. Introduction

In recent years, multimodal sentiment analysis (MSA) has attracted increasingly widespread attention [1–4]. Because of the heterogeneity among multimodal data, how to effectively fuse the representations of different modalities and ensure the semantic integrity of modalities is an important research topic in the community of MSA [5]. Some of the earlier works focus on the interaction between different modalities on low-level features, which results in limited fusion performance [1,6,7]. Inspired by the attention mechanism's [8] high-level relationship modeling capabilities, increasing MSA methods introduced attention when fusing unimodal representations. For example, Multimodal transformer (MulT) [2] employs the cross-modal attention mechanism to capture multimodal sequence interactions across different time steps. Some other works, such as Text Enhanced Transformer Fusion Network (TETFN) [9], Fine-grained Tri-modal Interaction Model (FGTI) [4], multimodal 3D stereoscopic attention [10], etc. have also witnessed the success of the attention-based methods in MSA application.

These methods fuse cross-modal features well but ignore the inherent information and potential conflicts of individual modalities, making the fused information somewhat incomplete. Some studies have noted this problem, either mapping unimodal representations to modality-invariant and modality-specific spaces and modeling them separately subsequently for fusion [3,11,12], or leveraging the multi-task learning (MTL) framework to model inter-modal differences in a supervised learning mode through unimodal label generation [13,14] or manual annotation [15].

However, these approaches still suffer from some shortcomings. First, there is still a potential conflict between emotional information contained by different bimodal combinations. Considering only inter-unimodal differences is not sufficient. For example, the combination of a smiling expression

and a positive word is positive, whereas audio represents sarcasm. In this case, the combination of textual and visual modalities and the combination of textual and acoustic modalities would conflict with the emotional polarity. Secondly, for those methods based on MTL, manual annotation of unimodal labels is costly, whereas label generation methods [13,14] rely on the quality of unimodal and cross-modal representations, and binary partitioning of the representation center may suffer from insufficient granularity.

To address these challenges, we propose a multi-level conflict-aware network (MCAN) that models consistency and discrepancy from different levels. Specifically, the MCAN is divided into the main branch and the conflict modeling branch. Wherein, the main branch progressively models the relationship between unimodal and bimodal representations utilizing Micro Multi-step Interaction Network (Micro-MSIN) and Macro Multi-step Intersection Network (Macro-MSIN) and segregates the inter-unimodal and inter-bimodal conflict components hierarchically, then feeds them to the conflict modeling branch. The conflict modeling branch models inter-unimodal and inter-bimodal conflicts through micro conflict-aware cross-attention (Micro-CACA) and macro conflict-aware cross-attention (Macro-CACA), respectively. To avoid introducing unstable representation-based generated labels, the conflict modeling branch directly encourages the unimodal and bimodal representations to generate inconsistent predictions to fully capture the conflict constituents, which will be joint-trained with the main branch. MCAN significantly outperforms the baselines on CMU-MOSI and CMU-MOSEI datasets. Extensive ablation experiments validate the effectiveness of the core component and the influence of the important hyperparameter of MCAN.

2. Methodology

The framework of the proposed multi-level conflict-aware network (MCAN) is shown in Figure 1. MCAN first conducts feature extraction for the three input modalities. For language modality, we feed the input text into BERT to obtain the language feature F_t . While LSTM is adopted to capture the intra-modality interaction F_v and F_a for visual and audio modalities.

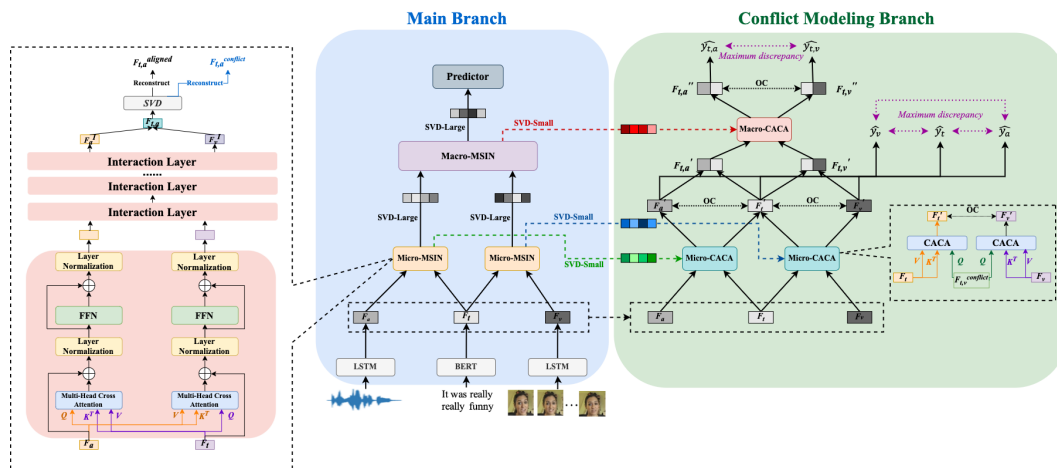


Figure 1. The overall framework of MCAN, MSIN and CACA.

2.1. Main Branch

The function of the main branch is to progressively fuse and align cross-modal representations of different granularities and to segregate conflict constituents. The two core components of the main branch are Transformer-style modules: Micro-MSIN and Macro-MSIN. Micro-MSIN receives F_t and F_v , F_t and F_a as inputs, and obtains the outputs $F_{t,a}$ and $F_{t,v}$. Then, inspired by [16,17], we conduct Singular Value Decomposition (SVD) of $F_{t,a}$ and $F_{t,v}$, and reconstruct the *top-k* singular values and the corresponding eigenvectors into alignment constituents ($F_{t,a}^{aligned}$ and $F_{t,v}^{aligned}$), which are fed to the

Macro-MSIN. The remaining singular values and their corresponding eigenvectors are reconstructed into conflict constituents ($F_{t,a}^{conflict}$, and $F_{t,v}^{conflict}$) to be delivered to the conflict modeling branch.

Macro-MSIN receives $F_{t,a}^{aligned}$ and $F_{t,v}^{aligned}$ as inputs and obtains the fused representation F_c , the aligned constituent $F_c^{aligned}$, and the conflicting constituent $F_c^{conflict}$ through a similar computational process to that of Micro-MSIN. The purpose of Macro-MSIN is to fully fuse and align the bimodal representations and separate out the conflict constituents between the bimodal representations. The cascade of Micro-MSIN and Macro-MSIN can make the modeling of MSA modal relationships more adequate and complete.

2.1.1. Micro Multi-Step Interaction Network

The Micro-MSIN modules receive the F_t and F_a , F_t and F_v as inputs. Following previous work [14,18–21], we treat the textual modality as the main contributing modality and thus do not set Micro-MSIN between F_a and F_v . It consists of multiple layers of Cross-Transformers. Taking the audio-text pairs as an example, the outputs of $(i-1)$ -th layer are $F_t^{(i-1)} \in \mathbb{R}^{n_t \times d}$ and $F_a^{(i-1)} \in \mathbb{R}^{n_a \times d}$, which will be fed to i -th Cross-Transformer layer. For textual modality, $F_t^{(i-1)}$ is transformed into Query to interact with audio modal input features, which are transformed into Key and Value. The computation for the multi-head cross-modal attention of textual modality is given as follows:

$$\text{head}_j^t = \text{SoftMax} \left(\frac{F_t^{(i-1)} W_Q \left(F_a^{(i-1)} W_K \right)^\top}{\sqrt{d_k}} \right) F_a^{(i-1)} W_V \quad (1)$$

$$\text{MultiHead}_t = \text{Concat} \left(\text{head}_1^t, \dots, \text{head}_e^t \right) W_O \quad (2)$$

where $W_Q, W_K, W_V \in \mathbb{R}^{d \times d_k}$, $W_O \in \mathbb{R}^{ed_k \times d}$, e is the number of attention heads. For audio modal, $F_a^{(i-1)}$ will be transformed into a Query and $F_t^{(i-1)}$ will be transformed into Key and Value, then conduct attention computation. Then, the output of cross-modal attention is processed by residual connection, layer normalization and feed-forward neural network (FFN), which is similar to naïve Transformer, and yield output of the i -th interaction layer $F_g^{(i)}$, $g \in t, a$. Assuming that the Micro-MSIN has a total of I layers, the output of the last layer is noted as $F_{t,a}$.

$$F_{t,a} = \text{Concatenate}(F_t^I, F_a^I) \quad (3)$$

To retain the alignment constituents and separate the conflict constituents to the greatest extent possible, we perform SVD, $F_{t,a} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top \in \mathbb{R}^{m \times n}$, $\mathbf{\Sigma} \in \mathbb{R}^{h \times h}$. In this case, the largest k singular values and the corresponding eigenvectors are considered to be the parts with significant alignment denoted as $F_{t,a}^{aligned}$, while the remaining singular values and the corresponding eigenvectors are regarded as the parts with insignificant alignment, i.e., conflicting, and are denoted as $F_{t,a}^{conflict}$.

$$\begin{aligned} F_{t,a}^{aligned} &= \mathbf{U}_{m \times k} \mathbf{\Sigma}_{k \times k} \mathbf{V}_{k \times n}^\top \\ F_{t,a}^{conflict} &= \mathbf{U}_{m \times (h-k)} \mathbf{\Sigma}_{(h-k) \times (h-k)} \mathbf{V}_{(h-k) \times n}^\top \end{aligned} \quad (4)$$

For text-visual pairs, the similar computation process is conducted, which yields $F_{t,v}^{aligned}$ and $F_{t,v}^{conflict}$ as outputs.

2.1.2. Macro Multi-Step Interaction Network

Macro-MSIN serves to model the alignment constituents and conflict constituents between bimodal representations. Macro-MSIN receives $F_{t,a}^{aligned}$ and $F_{t,v}^{aligned}$ as inputs, and its outputs are shown in the following calculations:

$$Z_c^{aligned}, Z_c^{conflict} = \text{Macro-MSIN}(F_{t,a}^{aligned}, F_{t,v}^{aligned}) \quad (5)$$

Micro-MSIN is more fine-grained compared to Macro-MSIN, and they are cascaded to progressively align cross-modal representations at different levels and effectively disentangle conflict knowledge.

2.2. Conflict Modeling Branch

The conflict modeling branch was designed to receive conflict constituents at different levels from the main branch, and model task conflict in terms of both representations and predicted outputs. It mainly consists of Micro Conflict-aware Cross-Attention (Micro-CACA) and Macro Conflict-aware Cross Attention (Macro-CACA), which are employed for further modeling of conflicts between unimodal representations and bimodal representations, respectively.

2.2.1. Micro Conflict-Aware Cross-Attention

The role of Micro-CACA is to adaptively fuse conflict constituents into unimodal representations. To illustrate with the case of text-visual pairs, the conflict constituent $F_{t,v}^{conflict}$ from the main branch will be transformed into Query. The output of the textual modality obtained after Micro-CACA processing is F'_t

$$F'_t = \text{SoftMax} \left(\frac{F_{t,v}^{conflict} W_Q^c (F_t W_K^t)^\top}{\sqrt{d_c}} \right) F_t W_V^t \quad (6)$$

Similarly, we can obtain Micro-CACA outputs F'_v and F'_a for visual and acoustic modalities. In particular, the two Micro-CACAs will generate two textual modal representations, which we average as the final outputs.

To further emphasize the discrepancy between unimodal representations, we impose orthogonal constraints on F'_t, F'_v and F'_a :

$$\mathcal{L}_{micro}^{oc} = \sum_{p \in \{l,v,a\}} \sum_{q \neq p} \|F_p'^\top F_q'\|_F^2 \quad (7)$$

Furthermore, we set individual FFN prediction heads for F'_t, F'_v and F'_a and encourage them to generate distinct predictions as much as possible to further emphasize the conflicting aspects between unimodal representations at the level of the prediction outputs.

$$\mathcal{L}_{micro}^{diff} = \sum_{p \in \{l,v,a\}} \sum_{q \neq p} |\hat{y}'_p - \hat{y}'_q|^2 \quad (8)$$

2.2.2. Macro Conflict-Aware Cross-Attention

The process of Macro-CACA is similar to that of Micro-CACA. Macro-CACA receives the separated conflict constituents of the main branch Macro-MSIN and transforms them into the Query of cross attention to capture and adaptively fuse inter-bimodal (between $F'_{t,a}$ and $F'_{t,v}$) conflicts. Similarly, the discrepancy constraints at the representation level and the predicted output level of Macro-CACA are represented as follows:

$$\mathcal{L}_{macro}^{oc} = \|F_{t,v}''^\top F_{t,a}''\|_F^2, \mathcal{L}_{macro}^{diff} = |\hat{y}''_{t,v} - \hat{y}''_{t,a}|^2 \quad (9)$$

where $F_{t,v}''$ and $F_{t,a}''$ are features extracted by Macro-CACA, $\hat{y}''_{t,v}$ and $\hat{y}''_{t,a}$ are predicted outputs of $F_{t,v}''$ and $F_{t,a}''$. The final loss function is represented as follows:

$$\mathcal{L} = \mathcal{L}_{main} + \alpha(\mathcal{L}_{micro}^{oc} + \mathcal{L}_{macro}^{oc}) + \beta(\mathcal{L}_{micro}^{diff} + \mathcal{L}_{macro}^{diff}) \quad (10)$$

where \mathcal{L}_{main} is mean squared error loss, α and β are trade-off parameters to control the intensity of conflict modeling.

3. Experiment

3.1. Datasets, Metrics and Implementation Details

We evaluate MCAN on CMU-MOSI [22] and CMU-MOSEI [23] datasets, which are the most widely used benchmark for MSA. Five different metrics are employed to evaluate the performance of MCAN and baselines: binary accuracy (Acc2), 7-class accuracy (Acc7), F1 Score (F1), Pearson correlation (Corr), and mean absolute error (MAE). For the Experimental setting, α and β are set to 1e-2 and 1e-3, respectively. Adam is adopted as the optimizer with an initial learning rate 5e-5 for BERT and 1e-4 for other parameters. Additionally, We select the *top* – 44 singular values and the corresponding eigenvectors for generating the alignment constituents

3.2. Comparison with Baselines

To validate the effectiveness of our proposed method, the baselines we chose cover classical MSA methods, and recent competitive approaches: TFN [1], LMF [6], MARN [7], RAVEN [24], MulT [2], MISA [3], Self-MM [13], GFML [14], MMIN [4], MSAN [10].

The results of the comparative analysis, as illustrated in Table 1, demonstrate that our model achieves significant improvement compared to baselines across different datasets. Fusion-based methods such as TFN, and LMF, despite their simplicity, have limited performance due to the difficulty of capturing high-level feature interactions. Compared to these fusion-based methods, attention-based methods such as MARN, RAVEN, and MulT demonstrate improved performance. Benefiting from the excellent high-level relationship capture capabilities of the attention mechanism, MMIN, and MSAN design novel attention modules to fine-grained align the representations of different modalities and achieve performance improvements. Self-MM and GFML focus on the intrinsic differences between modalities by introducing generated labels to model unimodal differences under the MTL framework. In contrast to the above methods, our approach balances alignment and conflict of modal representations at different levels and avoids the introduction of unstable generated labels by encouraging conflicting modeling branches to yield distinct predictions. As a result, the proposed MCAN further improves the performance of MSA.

Table 1. The experiment results on CMU-MOSI and CMU-MOSEI across various evaluation metrics.

Model	CMU-MOSI					CMU-MOSEI				
	Acc2	Acc7	F1	Corr	MAE	Acc2	Acc7	F1	Corr	MAE
TFN	76.8	32.5	76.3	0.601	0.998	78.5	43.7	78.0	0.665	0.709
LMF	77.4	33.9	76.5	0.638	0.922	78.8	42.9	79.1	0.644	0.682
MARN	78.1	34.7	77.0	0.655	0.908	79.3	44.8	79.7	0.673	0.672
RAVEN	79.8	36.2	79.3	0.699	0.886	80.5	45.7	80.0	0.678	0.631
MulT	81.3	38.4	81.4	0.734	0.802	82.9	47.7	82.8	0.744	0.586
MISA	81.7	40.6	81.3	0.720	0.793	83.3	49.8	83.2	0.767	0.572
Self-MM	82.5	40.9	82.4	0.769	0.725	84.1	49.8	84.4	0.786	0.555
GFML	83.9	41.9	83.8	0.804	0.694	85.1	50.1	84.8	0.795	0.541
MMIN	84.2	42.6	84.1	0.805	0.671	85.3	50.0	85.3	0.791	0.542
MSAN	83.6	41.5	83.7	0.794	0.712	84.6	49.5	84.2	0.768	0.551
MCAN (Ours)	84.5	43.1	84.8	0.811	0.675	85.8	51.6	85.9	0.798	0.527

3.3. Ablation Study

The effectiveness of core components and each loss in our method is verified by ablation experiments on the CMU-MOSI dataset, and the results are shown in Table 2. We individually removed \mathcal{L}_{diff} and \mathcal{L}_{oc} (Sum of corresponding terms for Micro-CACA and Macro-CACA) to assess the efficacy of these discrepancy constraints constraints. The experimental results reveal that the omission of either \mathcal{L}_{diff} or \mathcal{L}_{oc} results in a noticeable deterioration in model performance. Specifically, \mathcal{L}_{diff} and \mathcal{L}_{oc} function to regularize the feature and prediction aspects, respectively. Furthermore, the experiments verify the effect of the Conflict Modeling Branch (denoted as CMB in Table 2). The design of CMB improves the conflict-capturing ability of our model. Lastly, we confirmed that the choice of the

truncation position of singular values in SVD is critical to the outcomes. Different truncation positions will affect the amount of information assigned to the alignment and conflict constituents.

Table 2. Ablation study of MCAN on CMU-MOSI. “w/o” means without the specific components.

Ablation	Acc2	Acc7	F1	Corr	MAE
Effect of discrepancy constraints					
w/o \mathcal{L}_{diff}	82.1	42.3	82.0	0.763	0.814
w/o \mathcal{L}_{oc}	81.9	42.2	82.0	0.759	0.816
Effect of CMB					
w/o CMB	82.3	42.5	82.2	0.774	0.711
Effect of truncation positions					
Top-8	79.9	36.5	80.2	0.700	0.821
Top-16	83.8	42.5	83.6	0.796	0.701
Top-24	82.5	40.5	82.6	0.745	0.771
Top-36	84.3	42.7	84.3	0.807	0.698
Top-52	83.4	41.7	83.3	0.776	0.720
Top-64	83.0	41.1	83.0	0.762	0.742

4. Conlusions

In this paper, we develop a novel MCAN for MSA. To balance the discrepancies between unimodal and bimodal representations while fusing and aligning cross-modal representations, MCAN is divided into a main branch and a conflict modeling branch, which are jointly trained in a multi-task learning manner. The former progressively extracts different levels of cross-modal alignment and segregates the conflict constituents through the cascade of Micro-MSIN and Macro-MSIN, while the latter receives these conflict constituents and further models the conflicts. The experimental results show that MCAN outperforms the current state-of-the-art methods. In future work, we will endeavor to further analyze the modal conflict problem at the optimization level (e.g., gradient) and improve the proposed method.

Acknowledgments: The work was supported by the National Natural Science Foundation of China (No.72271017).

References

1. Zadeh, A.; Chen, M.; Poria, S.; Cambria, E.; Morency, L.P. Tensor Fusion Network for Multimodal Sentiment Analysis. In Proceedings of the Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, 2017, pp. 1103–1114.

2. Tsai, Y.H.H.; Bai, S.; Liang, P.P.; Kolter, J.Z.; Morency, L.P.; Salakhutdinov, R. Multimodal transformer for unaligned multimodal language sequences. In Proceedings of the Proceedings of the conference. Association for computational linguistics. Meeting. NIH Public Access, 2019, Vol. 2019, p. 6558.

3. Hazarika, D.; Zimmermann, R.; Poria, S. Misa: Modality-invariant and-specific representations for multi-modal sentiment analysis. In Proceedings of the Proceedings of the 28th ACM international conference on multimedia, 2020, pp. 1122–1131.

4. Fang, L.; Liu, G.; Zhang, R. Multi-Grained Multimodal Interaction Network for Sentiment Analysis. In Proceedings of the ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2024, pp. 7730–7734.

5. Wang, D.; Liu, S.; Wang, Q.; Tian, Y.; He, L.; Gao, X. Cross-modal enhancement network for multimodal sentiment analysis. *IEEE Transactions on Multimedia* **2022**, *25*, 4909–4921.

6. Liu, Z.; Shen, Y.; Lakshminarasimhan, V.B.; Liang, P.P.; Zadeh, A.; Morency, L.P. Efficient low-rank multi-modal fusion with modality-specific factors. *arXiv preprint arXiv:1806.00064* **2018**.

7. Zadeh, A.; Liang, P.P.; Poria, S.; Vij, P.; Cambria, E.; Morency, L.P. Multi-attention recurrent network for human communication comprehension. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2018, Vol. 32.

8. Vaswani, A. Attention is all you need. *Advances in Neural Information Processing Systems* **2017**.

9. Wang, D.; Guo, X.; Tian, Y.; Liu, J.; He, L.; Luo, X. TETFN: A text enhanced transformer fusion network for multimodal sentiment analysis. *Pattern Recognition* **2023**, *136*, 109259.

10. Huang, J.; Pu, Y.; Zhou, D.; Shi, H.; Zhao, Z.; Xu, D.; Cao, J. Multimodal Sentiment Analysis Based on 3D Stereoscopic Attention. In Proceedings of the ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2024, pp. 11151–11155.
11. Yang, D.; Huang, S.; Kuang, H.; Du, Y.; Zhang, L. Disentangled representation learning for multimodal emotion recognition. In Proceedings of the Proceedings of the 30th ACM International Conference on Multimedia, 2022, pp. 1642–1651.
12. Li, Y.; Wang, Y.; Cui, Z. Decoupled multimodal distilling for emotion recognition. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 6631–6640.
13. Yu, W.; Xu, H.; Yuan, Z.; Wu, J. Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis. In Proceedings of the Proceedings of the AAAI conference on artificial intelligence, 2021, Vol. 35, pp. 10790–10797.
14. Sun, X.; Ren, X.; Xie, X. A Novel Multimodal Sentiment Analysis Model Based on Gated Fusion and Multi-Task Learning. In Proceedings of the ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2024, pp. 8336–8340.
15. Yu, W.; Xu, H.; Meng, F.; Zhu, Y.; Ma, Y.; Wu, J.; Zou, J.; Yang, K. Ch-sims: A chinese multimodal sentiment analysis dataset with fine-grained annotation of modality. In Proceedings of the Proceedings of the 58th annual meeting of the association for computational linguistics, 2020, pp. 3718–3727.
16. Chen, X.; Wang, S.; Long, M.; Wang, J. Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In Proceedings of the International conference on machine learning. PMLR, 2019, pp. 1081–1090.
17. Wu, A.; Zhao, S.; Deng, C.; Liu, W. Generalized and discriminative few-shot object detection via SVD-dictionary enhancement. *Advances in Neural Information Processing Systems* **2021**, *34*, 6353–6364.
18. Yang, H.; Zhao, Y.; Wu, Y.; Wang, S.; Zheng, T.; Zhang, H.; Che, W.; Qin, B. Large Language Models Meet Text-Centric Multimodal Sentiment Analysis: A Survey. *arXiv preprint arXiv:2406.08068* **2024**.
19. Lin, R.; Hu, H. Multi-task momentum distillation for multimodal sentiment analysis. *IEEE Transactions on Affective Computing* **2023**.
20. Zhang, H.; Wang, Y.; Yin, G.; Liu, K.; Liu, Y.; Yu, T. Learning Language-guided Adaptive Hyper-modality Representation for Multimodal Sentiment Analysis. In Proceedings of the Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 2023, pp. 756–767.
21. Huang, J.; Ji, Y.; Yang, Y.; Shen, H.T. Cross-modality Representation Interactive Learning for Multimodal Sentiment Analysis. In Proceedings of the Proceedings of the 31st ACM International Conference on Multimedia, 2023, pp. 426–434.
22. Zadeh, A.; Zellers, R.; Pincus, E.; Morency, L.P. Multimodal Sentiment Intensity Analysis in Videos: Facial Gestures and Verbal Messages. *IEEE Intelligent Systems* **2016**, *31*, 82–88.
23. Zadeh, A.B.; Liang, P.P.; Poria, S.; Cambria, E.; Morency, L.P. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In Proceedings of the Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2018, pp. 2236–2246.
24. Wang, Y.; Shen, Y.; Liu, Z.; Liang, P.P.; Zadeh, A.; Morency, L.P. Words can shift: Dynamically adjusting word representations using nonverbal behaviors. In Proceedings of the Proceedings of the AAAI conference on artificial intelligence, 2019, Vol. 33, pp. 7216–7223.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.