

PSA-MF: Personality-Sentiment Aligned Multi-Level Fusion for Multimodal Sentiment Analysis

Heng Xie¹, Kang Zhu², Zhengqi Wen^{3*}, Jianhua Tao^{3,4*}, Xuefei Liu^{5*}, Ruibo Fu⁶, Changsheng Li^{1*}

¹School of Computer Science & Technology, Beijing Institute of Technology, Beijing, China

²School of Computer Science, Wuhan University, Wuhan, China

³Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing, China

⁴Department of Automation, Tsinghua University, Beijing, China

⁵College of Computer and Information Engineering, Tianjin Normal University, Tianjin, China

⁶Institute of Automation, Chinese Academy of Sciences, Beijing, China

Abstract

Multimodal sentiment analysis (MSA) is a research field that recognizes human sentiments by combining textual, visual, and audio modalities. The main challenge lies in integrating sentiment-related information from different modalities, which typically arises during the unimodal feature extraction phase and the multimodal feature fusion phase. Existing methods extract only shallow information from unimodal features during the extraction phase, neglecting sentimental differences across different personalities. During the fusion phase, they directly merge the feature information from each modality without considering differences at the feature level. This ultimately affects the model's recognition performance. To address this problem, we propose a personality-sentiment aligned multi-level fusion framework. We introduce personality traits during the feature extraction phase and propose a novel personality-sentiment alignment method to obtain personalized sentiment embeddings from the textual modality for the first time. In the fusion phase, we introduce a novel multi-level fusion method. This method gradually integrates sentimental information from textual, visual, and audio modalities through multimodal pre-fusion and a multi-level enhanced fusion strategy. Our method has been evaluated through multiple experiments on two commonly used datasets, achieving state-of-the-art results.

Introduction

With the rapid development of the internet and social media, user-generated video data has increased significantly. To enhance user experience, research on multimodal sentiment analysis (MSA) has attracted growing attention and expanded to applications in human-computer interaction (Cambria et al. 2017), risk prediction (Zhu et al. 2023), psychotherapy (Yang, Jiang, and Sahli 2018), and other fields (Das and Singh 2023; Fan et al. 2024). Compared to unimodal sentiment analysis, MSA integrates data from multiple modalities, such as text, visual, and audio, into a unified sentiment representation. This modal complementarity is more conducive to accurate and stable sentiment recognition (Gandhi et al. 2023; Han, Chen, and Poria 2021). How-

ever, it also presents two key challenges: how to extract effective sentimental representations from each modality, and how to effectively combine sentiment information from multiple modalities.

Existing methods approach the problem from the perspectives of unimodal feature extraction and multimodal fusion. Initially, various unimodal features are extracted using pre-trained models or other unimodal feature extraction tools (Caschera, Ferri, and Grifoni 2016; Xie et al. 2023; Yang et al. 2022). On this basis, multimodal semantic-level fusion is further performed. From the feature extraction perspective, methods include unimodal feature encoding (Sun, Mai, and Hu 2022), modality feature decomposition (Yang et al. 2023), and contrastive learning (Lin et al. 2022). From the fusion perspective, methods include tensor-based fusion (Liu et al. 2018), attention mechanism-based fusion (Wei et al. 2020), and pre-trained model-based fusion (Rahman et al. 2020). However, current methods have issues: features extracted in the encoding phase do not incorporate personalized sentimental information and fail to account for differences in sentimental expression across personalities (Zhang, Peng, and Winkler 2019). These personality differences significantly affect sentimental expression and perception in real-world scenarios (Zhu et al. 2024). Additionally, due to the heterogeneity of different modal data, extracting unimodal information for basic fusion cannot fully understand complex sentimental states (Han et al. 2021). Traditional fusion strategies struggle to achieve deep interactions between modalities.

To address these problems, we propose a personality-sentiment aligned multi-level fusion (PSA-MF) method for MSA. From the feature extraction perspective, we integrate personality information into sentiment features by combining a pre-trained personality model with fine-tuned BERT to extract personalized sentiment features from text. We then achieve alignment between sentiment and personality through contrastive learning and a personalized sentimental constraint loss. From the perspective of fusion, we design a multi-level fusion method. First, the deep layers of BERT are used as the multimodal pre-fusion layer, combining shallow text embeddings with visual and audio features for preliminary alignment between modalities. Second,

*Corresponding author

the output from the multimodal pre-fusion layer serves as a query for cross-modal interaction with visual and audio modalities. This process guides the personalized generation of visual and audio features while reducing information bias from modality distribution differences. Finally, through serial and parallel fusion, semantic-level fusion of personalized sentiments is conducted from multiple angles, enhancing cross-modal transmission of personalized sentimental cues while maintaining the complementary characteristics of each modality. This refined fusion architecture overcomes modality noise interference in early fusion and avoids the shortcomings of late fusion that neglect cross-modal correlations, achieving a balance between global consistency and local specificity in sentimental semantics.

Our contributions are as follows.

- We are the first to combine personality pre-trained models with a phase-wise fine-tuned BERT architecture to extract personalized sentimental features from the textual modality. This addresses the limitation of traditional methods that ignore personality factors.
- We propose an innovative personality-sentiment alignment strategy that uses contrastive learning to achieve this alignment at the semantic level. Furthermore, we construct a constraint function that dynamically adjusts the alignment strength while confining the alignment process within the appropriate sentiment space.
- We propose a multi-level fusion architecture that includes pre-fusion, cross-modal interaction, and enhanced fusion. The pre-fusion features direct audio-visual modalities toward personalized feature generation in the cross-modal module. Finally, by combining serial fusion and parallel fusion in a collaborative mechanism, we effectively balance global sentimental consistency with local modality complementarity.

Related Work

Personality-Sentiment Interplay

In the current field of affective computing, personality and sentiment analysis are considered two core elements. Eladhari et al. (Eladhari 2014) proposed a personality-sentiment mapping model that maps different sentiments to corresponding personality traits. They developed a mental module applied in gaming, which provides a foundation for understanding the dynamic relationship between sentiments and personality. To conduct in-depth research on personality recognition, Ponce-López et al. (Ponce-López et al. 2016) created a Big Five personality dataset by screening 10,000 YouTube videos through fitting a Bradley-Terry-Luce model with maximum likelihood. This dataset has facilitated further research on the Big Five personality traits and sentiment analysis. Since most studies on sentiment and personality analysis are conducted separately, Zhang et al. (Zhang, Peng, and Winkler 2019) built upon Ponce-López’s work to combine personality and sentiment for analysis. They constructed an end-to-end network structure and demonstrated the interrelationship between sentiment and personality traits through a multi-task learning strategy. Mohammadi et al. (Mohammadi and Vuilleumier 2020) used a

component-based framework to analyze sentimental experiences and explored the relationship between personality traits and sentiments. The study found that the component model could effectively predict individual differences in discrete sentiments. It also emphasized the important role of personality in the generation of sentiments. With the rapid iteration of generative technologies, dialogue systems based on large language models have also matured. Ma et al. (Ma et al. 2020) further explored dialogue systems that are more humanlike, incorporating sentimental awareness and personality awareness as key features of the system.

The above research demonstrates that it is crucial to consider the influence of personality traits in sentiment analysis. They provide a theoretical foundation and experimental support for understanding the complex relationship between sentiments and personality traits. Based on this, we are applying a personality pre-trained model in a sentiment analysis framework and design a personality-sentiment alignment module. This method effectively addresses the limitation of traditional approaches that ignore personality factors and restrict sentimental recognition performance. It also solves the problem that existing models cannot conduct personalized sentiment research on datasets with only sentiment labels.

Multimodal Fusion Paradigms

Existing multimodal works typically utilize unimodal encoders to extract features and employ multimodal encoders for fusion, combining encoders with varying parameter counts for tasks such as classification or retrieval. In early multimodal work, visual semantic embedding (VSE) (Chen et al. 2021) proposed that the visual encoder is heavier compared to the text embedding, and multimodal interactions are relatively simple. Building on this, CLIP (Radford et al. 2021) elevated the text encoder to the same scale as the visual encoder, and the multimodal aspect used contrastive learning to achieve the alignment of visual and textual semantic information, making the fusion part lightweight. The work of Lu and Chen et al. (Lu et al. 2019; Chen et al. 2020) focused on modal interaction, achieving significant performance improvements by replacing traditional fusion methods such as concatenation and dot product with transformer blocks. To improve inference speed, Kim et al. (Kim, Son, and Kim 2021) drew inspiration from ViT (Dosovitskiy et al. 2020), replacing the visual encoder with patch embedding and using lightweight text and visual encoders for feature extraction, while employing heavier transformer blocks for multimodal fusion. To balance model performance and speed, Li et al. (Li et al. 2021) use transformer-based pre-trained models as unimodal encoders and utilize the deeper layers of the text encoder as multimodal encoders.

In multimodal research, sentiment analysis is an emerging direction that typically combines information from text, visual, and audio modalities for sentiment classification. Early research adopted early fusion (Morency, Mihalcea, and Doshi 2011) and late fusion (Nojavanasghari et al. 2016) methods. Current research utilizes pre-trained models to extract features from the textual modality, while visual and audio features are usually pre-extracted. Therefore, researchers are more focused on effectively fusing sentiment informa-

tion from different modalities. This can be specifically divided into two aspects: decoupled feature fusion and direct fusion.

Decoupling relies on sophisticated constraint conditions and requires the design of multiple loss functions, making the model relatively complex. For example, Zeng et al. (Zeng et al. 2024) proposed a DTN that employs multiple encoders and decoders to map features from various modalities into homogeneous and heterogeneous spaces. Constraints are applied to the heterogeneous features output by the decoder to distinguish between homogeneous and heterogeneous features. Finally, homogeneous features are fused for sentiment classification. Direct fusion, on the other hand, focuses more on the multimodal fusion framework. For example, Yang et al. (Yang, Xu, and Gao 2020) first utilized BERT and COVAREP (Degottex et al. 2014) to extract feature information from text and audio modalities, respectively. To achieve word-level modality alignment, they padded the audio sequence with zero vectors and finally employed masked multimodal attention for fusion.

In current MSA, pre-extracted features are commonly used for audio and video, while text features are extracted through pre-trained models. This results in text representation spaces having significantly higher information density than visual and audio modalities. Simple same-level feature fusion makes it difficult to coordinate information density and semantic depth between different modalities. Therefore, we constructed a multi-level fusion structure that gradually alleviates the semantic gap between modalities through progressive feature interaction.

Methodology

In this section, we provide a detailed description of our method, with the model architecture shown in Figure 1.

Unimodal Feature Extraction

In the feature extraction phase, we utilize pre-trained models to extract sentiment and personality features from the textual modality, and apply LSTM models to extract visual and audio features.

Textual Sentiment and Personality Feature Extraction. We use two pre-trained models, the first N layers of a fine-tuned BERT_s¹ and Personality BERT_p² to extract sentiment and personality embeddings from text, respectively. Given the text input X_t , the representation for both sentimental and personality features can be expressed as:

$$\text{CLS}_i = \text{BERT}_i(X_t; \theta_{\text{BERT}_i}), \quad i \in \{s, p\} \quad (1)$$

CLS_i represents the output embeddings of the [CLS] token from the BERT model, which captures different types of features based on the value of i . θ_{BERT_i} denotes the parameters of the respective BERT models. Here, $i \in \{s, p\}$ where s indicates the sentiment and p indicates the personality.

Visual and Audio Feature Extraction. We use LSTM to process visual and audio features. Given the visual input X_v and audio input X_a , the LSTM encoding outputs are:

$$h_m = \text{LSTM}(X_m; \theta_{\text{LSTM}_m}), \quad m \in \{v, a\} \quad (2)$$

where h_m represents the temporal feature embeddings extracted from different modalities, and θ_{LSTM_m} denotes the parameters of the LSTM models, with $m \in \{v, a\}$ indicating visual and audio modalities, respectively.

Personality-Sentiment Alignment

To align sentimental and personality features, we introduce contrastive learning. We first map these features to a common lower-dimensional space using linear transformations.

$$T_i = W_i \cdot \text{CLS}_i, \quad i \in \{s, p\} \quad (3)$$

where W_i is a linear learning mapping matrix and T_i represents the textual mapped embeddings, with $i \in \{s, p\}$ indicating sentiment and personality, respectively.

Specifically, the contrastive learning loss \mathcal{L}_{cl} follows the CLIP (Radford et al. 2021) formulation:

$$\mathcal{L}_{cl} = -\log \frac{\exp(\text{sim}(T_s^i, T_p^i)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(T_s^i, T_p^j)/\tau)} \quad (4)$$

where the total number of samples is N with $i, j \in [1, N]$ denoting sample indices, $\text{sim}(\cdot, \cdot)$ denotes the computation of cosine similarity, and τ is a temperature parameter that controls the smoothness of the feature distribution. This loss function encourages matched sentiment-personality pairs (positive pairs) to be closer in the feature space while pushing unmatched pairs (negative pairs) apart. Combined with cosine similarity, we define the compound contrastive learning loss:

$$\mathcal{L}_{ccl} = \text{sim}(T_s^i, T_p^i) \cdot \mathcal{L}_{cl} \quad (5)$$

where i denotes the i -th sample.

To further enhance the alignment of personalized sentiment features, we design a personalized sentimental constraint loss. We calculate the regression loss of personalized sentiment features using true sentimental labels y_i :

$$\mathcal{L}_{ps} = (1 - \text{sim}(T_s^i, T_p^i)) \cdot \|W_y \cdot T_s^i - y_i\|_1 \quad (6)$$

where W_y is a linear mapping and $\|\cdot\|_1$ denotes the L1 norm. This step optimizes the model’s error in personalized sentimental modeling tasks while balancing the integration of sentiment features with personality.

The overall personality-sentiment alignment loss is a combination of the two losses:

$$\mathcal{L}_{Align} = \mathcal{L}_{ccl} + \mathcal{L}_{ps} \quad (7)$$

In summary, contrastive learning aligns sentiment with personality, and the personalized sentimental constraint loss further refines this alignment by constraining the difference between personalized features and true labels, effectively enhancing the model’s sentiment classification performance.

¹<https://huggingface.co/google-bert/bert-base-uncased>

²<https://huggingface.co/Minej/bert-base-personality>

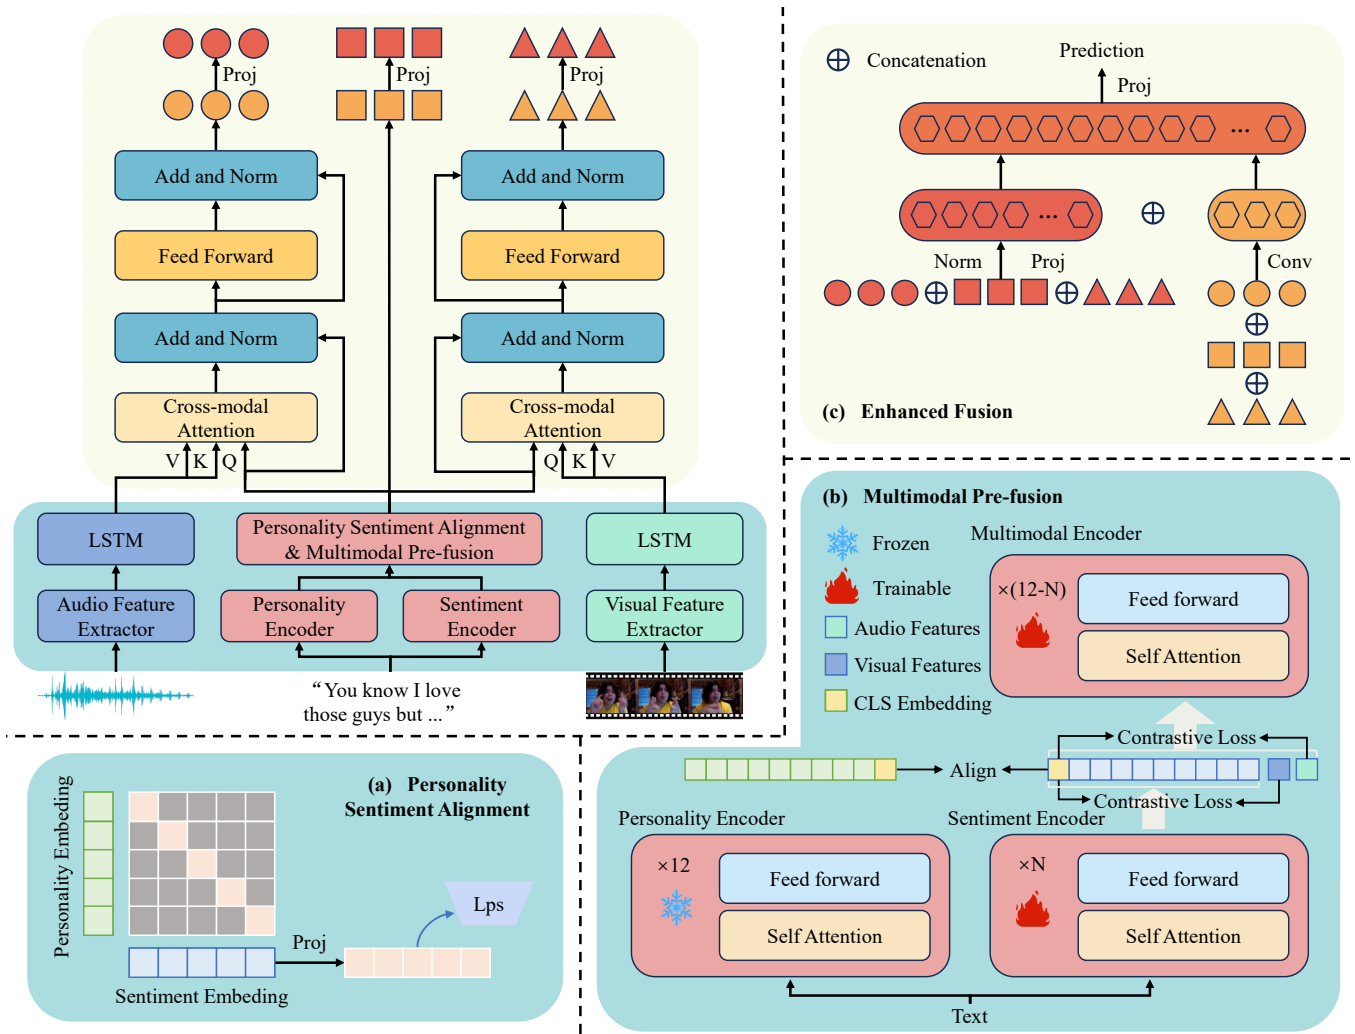


Figure 1: The proposed personality-sentiment aligned multi-level fusion model (PSA-MF): In the top left corner of the diagram is the main framework of the model, which includes feature extraction and cross-attention modal interaction. (a) shows the personality-sentiment alignment module, which includes personality-sentiment contrastive learning and personalized sentimental constraints. (b) displays the multimodal pre-fusion module of the model, which utilizes the deep layers of BERT as a multimodal encoder for initial alignment across three modalities. (c) depicts the enhanced fusion module, which performs serial fusion and parallel fusion of the multi-level features from the upper-layer cross-attention, culminating in the final prediction.

Multi-Level Modality Fusion

In the fusion phase, we propose a multi-level fusion strategy. This involves a step-by-step integration of semantics and modalities to enable the model to integrate and understand personalized emotional information more comprehensively.

Multimodal Pre-fusion Module. The feature information of each modality has been extracted through the above modules, including the personality information features from the textual modality. We begin the multimodal pre-fusion of the features above. First, we perform cross-modal contrastive learning between the textual modality and the other two modalities, with the loss function designed as follows:

$$\mathcal{L}_{clm} = \sum_{m \in \{v, a\}} -\log \frac{\exp(\text{sim}(\text{CLS}_s^i, X_m^i)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(\text{CLS}_s^i, X_m^j)/\tau)} \quad (8)$$

where CLS_s^i represents the [CLS] token of the text for the i -th sample, while X_v^i and X_a^i correspond to the visual and audio features of the same sample, respectively. \mathcal{L}_{clm} is the total contrastive learning loss for the text and both modalities v and a .

Then, we use the last (12-N) layers of the fine-tuned BERT as a multimodal encoder. The features from the three modalities are concatenated and input for pre-fusion. Through bidirectional encoding, contextual information is

effectively captured, particularly enhancing the semantic association capture in cross-modal space and the progressive transfer of personality representation. This produces the multimodal pre-fusion embedding.

$$\text{CLS}_m = \text{BERT}_m([\text{CLS}_s, X_v, X_a]; \theta_{\text{BERT}_m}) \quad (9)$$

where, $[\text{CLS}_s, X_v, X_a]$ represents the concatenated feature vectors of the three modalities, and θ_{BERT_m} denotes the parameters of the multimodal encoder, resulting in the multimodal pre-fusion embedding CLS_m .

Cross-modal Interaction Module. The cross-modal personality representation output from the pre-fusion stage interacts with the visual and audio modalities. In this process, the pre-fusion features serve as query vectors for attention, guiding personalized weight allocation for the other two modalities, while further eliminating stepwise shifts between heterogeneous modalities, achieving personality-driven modality-specific reconstruction. The feature representation that fuses visual and audio is:

$$V_t = \text{Att}_v(M_s, h_v, h_v) \quad (10)$$

$$A_t = \text{Att}_a(M_s, h_a, h_a) \quad (11)$$

where M_s represents the pre-fusion features after linear transformation of CLS_m , and h_v and h_a are the temporal features of the visual and audio modalities, respectively. Additionally, Att_v and Att_a are multi-head attention mechanisms. The fusion features V_t , A_t , and M_s undergo linear transformations to obtain the cross-modal sentimental enhancement representations V'_t , A'_t , and M'_s .

Enhanced Fusion Module. We perform enhanced fusion on the features obtained from the upper-level cross-modal interaction. This is aimed at strengthening the propagation of personality and sentimental information across modalities and enhancing the complementarity of information within each modality. Specifically, we first construct a dual-stream network with serial fusion and parallel fusion. Then, we concatenate the fusion results and pass them through a subnet for the final prediction.

The serial fusion layer uses a linear layer to merge the feature dimensions of the cross-modal sentimental enhancement representations V'_t , A'_t , and M'_s :

$$F_s = W_p \cdot [V'_t, A'_t, M'_s] \quad (12)$$

where W_p represents the learnable weight matrix, and $[V'_t, A'_t, M'_s]$ denotes the concatenated feature vectors. The LayerNorm operation ensures stable training by normalizing the features across the batch dimension. The fully connected layer further facilitates the transfer of personalized sentiments between different modalities while allowing for complementary learning among the various modalities.

The parallel fusion layer uses convolution operations to merge the stacked fusion features V_t , A_t , and M_s :

$$F_p = \text{Conv}([V_t, A_t, M_s]) \quad (13)$$

where Conv represents a convolutional operation with kernel size 3, which transforms the concatenated features

$[V_t, A_t, M_s]$ into a compressed representation, effectively capturing local spatial information across modalities.

The enhanced fusion features can be obtained by connecting two fusion paths. These enhanced fusion features are then fed into the final subnet to obtain the prediction result:

$$\hat{y} = \text{Sub}([F_s, F_p]) \quad (14)$$

where \hat{y} represents the final prediction result of the model, and $[F_s, F_p]$ denotes the connected enhanced fusion features. The Sub is a fully connected layer network structure with two layers of non-linear activation. The final sentiment prediction is made by enhancing and fusing personalized sentiment features from multiple perspectives.

The total loss function includes the loss of personality-sentiment alignment, the loss of cross-modal contrastive learning, and the loss L1 between predicted results and ground truth. The total loss function is expressed as:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{Align}} + \mathcal{L}_{\text{clm}} + \|\hat{y}_i - y_i\|_1 \quad (15)$$

Experiments

Datasets

To evaluate the performance of this method, experiments are conducted using two commonly used datasets: MOSI(Zadeh et al. 2016) and MOSEI(Zadeh et al. 2018). The MOSI dataset comprises 93 videos from YouTube, divided into 2,199 video segments. The MOSEI dataset includes 23,453 video clips from 1,000 speakers and 250 topics. Both datasets have sentiment scores ranging from -3 (strongly negative) to 3 (strongly positive).

Experimental Setup

Our method is implemented with PyTorch and trained on an NVIDIA RTX 3090 GPU. Initial learning rates for MOSI and MOSEI datasets are 10^{-4} , with batch sizes of 128 and 256, respectively. For consistency with existing work, video embeddings derive from 35 facial action units extracted using Facet (Littlewort et al. 2011), while audio embeddings come from features extracted using COVAREP (Degottex et al. 2014). We evaluated performance on regression and classification tasks using mean absolute error (MAE), Pearson correlation coefficient (Corr), binary accuracy (Acc2), seven-class accuracy (Acc7), and F1 score (F1).

Baselines

We compared our method with several state-of-the-art approaches in MSA, including both classic and recent advanced methods:

Tensor Fusion Methods: TFN (Zadeh et al. 2017) and LMF (Liu et al. 2018) achieve multimodal tensor-level fusion through specialized tensor fusion layers.

Cross-modal Attention and Encoding: MuLT (Tsai et al. 2019) and MISA (Hazarika, Zimmermann, and Poria 2020) utilize cross-modal attention mechanisms and specific modality feature encoding methods, respectively.

Contrastive Learning: MVCL (Liu et al. 2023) and Hy-Con (Mai et al. 2022) employ various contrastive learning

Method	MOSI					MOSEI				
	MAE↓	Corr↑	Acc7↑	Acc2↑	F1↑	MAE↓	Corr↑	Acc7↑	Acc2↑	F1↑
TFN	0.901	0.698	34.9	-/80.8	-/80.7	0.593	0.700	50.2	-/82.5	-/82.1
LMF	0.914	0.689	33.2	79.8/-	79.5/-	0.680	0.590	48.0	80.4/-	78.2/-
MuLT	0.871	0.698	-	-/83.0	-/82.8	0.580	0.703	-	-/82.5	-/82.3
MISA	0.783	0.761	42.3	81.8/83.4	81.7/83.6	0.555	0.756	-	83.6/85.5	83.8/85.3
MVCL	0.769	0.783	-	83.7	84.2	0.573	0.741	-	85.0	85.0
HyCon	0.713	0.790	46.6	85.2	85.1	0.601	0.776	52.8	85.4	85.1
PriSA	0.714	0.792	47.3	83.38/85.52	83.24/85.45	0.523	0.772	54.65	82.84/85.93	83.18/85.87
FGTI	0.702	0.791	48.0	85.8	85.8	0.536	0.771	53.4	86.0	86.0
ULMD	0.700	0.799	47.8	85.82	85.71	0.531	0.770	53.81	85.95	85.91
PSA-MF (Ours)	0.686	0.807	46.5	83.67/86.43	83.30/86.19	0.521	0.774	55.0	83.76/86.30	84.21/86.28

Table 1: Comparing the results of various baseline methods on the CMU-MOSI and CMU-MOSEI datasets.

techniques to learn modality differences and extract effective sentimental representations.

Priority-based Fusion: PriSA (Ma, Zhang, and Sun 2023) utilizes priority-based fusion to guide inter-modal correlation computation and employs distance-aware contrastive learning to handle mixed-modal correlations.

Multi-granularity Fusion: FGTI (Zhi et al. 2024) and MMIN (Fang, Liu, and Zhang 2024) focus on reducing modality redundancy and enhancing modality-invariant representations through multi-granularity fusion approaches.

Feature Decoupling Method: ULMD (Zhu et al. 2025) achieves feature decoupling by introducing modality separators and a unimodal label generation module, dividing representations into invariant and specific components.

Experimental Results

Table 1 presents the comprehensive experimental results of our proposed method compared to the baseline approaches on the MOSI and MOSEI datasets. Our method consistently outperforms existing state-of-the-art models across almost all metrics, demonstrating its effectiveness in MSA.

Compared to traditional tensor fusion methods like TFN and LMF, our approach shows significant improvements. On the MOSI dataset, our method achieves a 4.8% increase in Acc2 and a 4.6% increase in F1 compared to TFN. This substantial improvement can be attributed to our incorporation of personality information at the feature extraction level, which enables the model to capture individual differences in sentimental expression more accurately.

Compared to methods that utilize cross-modal attention mechanisms and specific modal encoding (MuLT and MISA), our approach still shows advantages. On the MOSEI dataset, our method improved Acc2 by 3.2% and F1 by 3% compared to MISA. This enhancement stems from the multi-level fusion mechanism of our model. During the multimodal pre-fusion stage and the cross-modal interaction stage, the model achieves a dynamic transfer of personalized sentimental features across modalities, initially bridging the modal differences. In the enhanced fusion stage, the model further strengthens these capabilities.

Our method also outperforms advanced techniques that utilize contrastive learning (MVCL and HyCon). This can be attributed to our pioneering introduction of personality-

sentiment contrastive learning in addition to cross-modal contrastive learning. Compared to their methods, we simultaneously consider both multimodal semantic differences and personalized sentimental differences.

Compared to state-of-the-art methods, our approach achieves 0.91% improvement in Acc2 and 0.74% increase in F1 over PriSA on the MOSI dataset, and 0.63% improvement in Acc2 and 0.39% increase in F1 over FGTI. Unlike PriSA, which uses priority-based fusion for inter-modal correlation and distance-aware contrastive learning, our personality-sentiment alignment and multi-level fusion better capture both fine-grained and high-level sentimental cues. While PriSA handles mixed-modal correlations, our method uses personality features for more nuanced, personalized sentiment analysis. FGTI focuses on multi-granularity fusion and reduces modal gaps through similarity metrics and self-supervised learning. Although FGTI enhances modal specificity through constraints, our method uses personality information for more differentiated sentiment understanding across modalities.

Regarding decoupling methods, our model yields higher performance, improving MOSI Acc2 and F1 by 0.61% and 0.48% over ULMD. Our approach is simpler, aligning personality and sentiment via constraint functions without complex decoupling networks. By adapting BERT into a multimodal encoder and employing multi-level fusion, we effectively mitigate data heterogeneity while streamlining model training and convergence.

In summary, our model shows superior performance across datasets, demonstrating the efficacy of personality-sentiment alignment. By progressively integrating multimodal data and learning personality-sentiment dependencies, our model better captures subtle expressions and improves overall sentiment analysis accuracy.

Ablation Study

To validate each component in our model, we conducted a comprehensive ablation study on the MOSI dataset, which mainly includes module ablation and loss function ablation. Results are shown in Table 2.

Module Ablation Study. We investigated removing key components: the personality feature extraction (w/o-PF), the BERT-based multimodal pre-fusion (w/o-BF), and the en-

Model	MAE↓	Corr↑	Acc-7↑	Acc-2↑	F1↑
w/o-PF	0.711	0.795	44.6	84.60	84.47
w/o-BF	0.735	0.778	43.2	84.76	84.44
w/o-EF	0.806	0.788	37.3	85.21	84.96
w/o-(\mathcal{L}_{ps})	0.724	0.784	44.5	83.99	83.92
w/o-(\mathcal{L}_{clm})	0.754	0.784	40.8	85.06	84.92
PSA-MF(Ours)	0.686	0.807	46.5	86.43	86.19

Table 2: Module and loss function ablation study on the CMU-MOSI.

hanced fusion module (w/o-MF).

Removing the personality feature extraction caused the most significant decline, with Acc2 dropping by 1.83% and F1 by 1.72%. This highlights the critical role of personality information in obtaining personalized sentimental features. By incorporating personality information, the model recognizes subtle differences in sentimental expression among individuals, improving sentiment recognition accuracy.

Removing the BERT-based multimodal pre-fusion decreased performance, with Acc2 dropping by 1.67% and F1 by 1.75%. This indicates its crucial role in alleviating modal heterogeneity and facilitating the dynamic transfer of personality information across modalities in early fusion.

Removing the multi-level fusion reduces Acc2 and F1 by 1.22% and F1 by 1.23%, respectively. This demonstrates enhanced fusion strengthens cross-modal transfer of personalized sentimental information while maintaining complementarity across modalities. It enables capturing fine-grained and high-level sentimental cues, improving sentiment classification capability.

Loss Function Ablation Study. We also investigated the impact of removing loss functions: the personalized sentiment constraint loss (w/o-(\mathcal{L}_{ps})) and the multimodal contrastive learning loss (w/o-(\mathcal{L}_{clm})).

Removing the personalized sentimental constraint loss caused the most significant decline among all ablation experiments, with Acc2 dropping by 2.44% and F1 by 2.27%. This loss is vital for balancing training, as it guides the model to align personality traits with appropriate sentiments while preventing overfitting to personality at the cost of label accuracy. By constraining deviations between representations and labels, \mathcal{L}_{ps} captures complex personality-sentiment dependencies. Removing this constraint would cause an imbalance between personality matching and sentiment classification, leading to performance degradation.

Removing the multimodal contrastive learning loss degrades Acc2 and F1 by 1.39% and F1 by 1.27%. This suggests contrastive learning during multimodal pre-fusion helps align different modalities. By reducing unimodal noise and enhancing complementary sentimental features, the model establishes semantic relationships between modalities in later cross-modal interactions.

Layer Analysis

This study investigates the effects of the personality-sentiment alignment strategy at different network layers through a hierarchical experiment. Figure 2 shows the changes in Acc2 and F1 metrics for each layer’s alignment




Figure 2: Acc2 and F1 for personality-sentiment alignment applied at different layers on the CMU-MOSI.

strategy on the test set. The model architecture consists of 13 layers, with alignment strategies as follows:

Layers 1-11: Personality embeddings are aligned with text embeddings outputted by the text encoder. The alignment effect in the shallow layers is relatively weak, as the shallow features of the text encoder have not captured sufficient sentimental semantic information.

In particular, in pure text processing layers (Layers 1-11), as network depth increases, model performance shows a stable upward trend, peaking at Layer 11. This confirms that deeper text encoders can extract richer sentimental semantic features, enhancing personality-sentiment alignment.

Layer 12: Personality embeddings align with multimodal features obtained from concatenated visual and audio features. Performance in this layer significantly declines, likely due to modality noise in unfused visual and audio features, making cross-modal semantic matching difficult.

Layer 13: Personality embeddings align with fused features from multimodal encoding. Performance in this layer improves compared to Layer 12 but remains below Layer 11. Analysis attributes this to personality embeddings serving as pure text representations, creating semantic space discrepancy with multimodal joint features.

Conclusions

This paper presents PSA-MF, an MSA method that incorporates personality information and employs multi-level fusion for personalized sentimental feature extraction. It extracts personality traits from text and maps personality embeddings into the sentimental representation space through a personality-sentiment alignment module. The multi-level fusion strategy enhances personality information flow across modalities and facilitates progressive fusion of personalized sentimental information. Comparative experiments and ablation studies validate the method’s superiority. Future work could extend personalized sentiment analysis to other modalities and explore new fusion strategies.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No.U2441242).

References

- Cambria, E.; Das, D.; Bandyopadhyay, S.; and Feraco, A. 2017. Affective computing and sentiment analysis. *A practical guide to sentiment analysis*, 1–10.
- Caschera, M. C.; Ferri, F.; and Grifoni, P. 2016. Sentiment analysis from textual to multimodal features in digital environments. In *proceedings of the 8th international conference on Management of Digital EcoSystems*, 137–144.
- Chen, J.; Hu, H.; Wu, H.; Jiang, Y.; and Wang, C. 2021. Learning the best pooling strategy for visual semantic embedding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15789–15798.
- Chen, Y.-C.; Li, L.; Yu, L.; El Kholy, A.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2020. Uniter: Universal image-text representation learning. In *European conference on computer vision*, 104–120. Springer.
- Das, R.; and Singh, T. D. 2023. Multimodal sentiment analysis: a survey of methods, trends, and challenges. *ACM Computing Surveys*, 55(13s): 1–38.
- Degottex, G.; Kane, J.; Drugman, T.; Raitio, T.; and Scherer, S. 2014. COVAREP—A collaborative voice analysis repository for speech technologies. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 960–964. IEEE.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Eladhari, M. P. 2014. The mind module—Using an affectand personality computational model as a game-play element. *IEEE Transactions on Affective Computing*, 5(1): 3–16.
- Fan, C.; Xie, H.; Tao, J.; Li, Y.; Pei, G.; Li, T.; and Lv, Z. 2024. ICaps-ResLSTM: Improved capsule network and residual LSTM for EEG emotion recognition. *Biomedical Signal Processing and Control*, 87: 105422.
- Fang, L.; Liu, G.; and Zhang, R. 2024. Multi-Grained Multimodal Interaction Network for Sentiment Analysis. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7730–7734. IEEE.
- Gandhi, A.; Adhvaryu, K.; Poria, S.; Cambria, E.; and Husain, A. 2023. Multimodal sentiment analysis: A systematic review of history, datasets, multimodal fusion methods, applications, challenges and future directions. *Information Fusion*, 91: 424–444.
- Han, W.; Chen, H.; Gelbukh, A.; Zadeh, A.; Morency, L.-p.; and Poria, S. 2021. Bi-bimodal modality fusion for correlation-controlled multimodal sentiment analysis. In *Proceedings of the 2021 international conference on multimodal interaction*, 6–15.
- Han, W.; Chen, H.; and Poria, S. 2021. Improving multimodal fusion with hierarchical mutual information maximization for multimodal sentiment analysis. *arXiv preprint arXiv:2109.00412*.
- Hazarika, D.; Zimmermann, R.; and Poria, S. 2020. Misa: Modality-invariant and-specific representations for multimodal sentiment analysis. In *Proceedings of the 28th ACM international conference on multimedia*, 1122–1131.
- Kim, W.; Son, B.; and Kim, I. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In *International conference on machine learning*, 5583–5594. PMLR.
- Li, J.; Selvaraju, R.; Gotmare, A.; Joty, S.; Xiong, C.; and Hoi, S. C. H. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34: 9694–9705.
- Lin, Z.; Liang, B.; Long, Y.; Dang, Y.; Yang, M.; Zhang, M.; and Xu, R. 2022. Modeling intra-and inter-modal relations: Hierarchical graph contrastive learning for multimodal sentiment analysis. In *Proceedings of the 29th international conference on computational linguistics*, volume 29, 7124–7135. Association for Computational Linguistics.
- Littlewort, G.; Whitehill, J.; Wu, T.; Fasel, I.; Frank, M.; Movellan, J.; and Bartlett, M. 2011. The computer expression recognition toolbox (CERT). In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, 298–305. IEEE.
- Liu, P.; Zheng, X.; Li, H.; Liu, J.; Ren, Y.; Zhu, H.; and Sun, L. 2023. Improving the modality representation with multi-view contrastive learning for multimodal sentiment analysis. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Liu, Z.; Shen, Y.; Lakshminarasimhan, V. B.; Liang, P. P.; Zadeh, A.; and Morency, L.-P. 2018. Efficient low-rank multimodal fusion with modality-specific factors. *arXiv preprint arXiv:1806.00064*.
- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.
- Ma, F.; Zhang, Y.; and Sun, X. 2023. Multimodal Sentiment Analysis with Preferential Fusion and Distance-aware Contrastive Learning. In *2023 IEEE International Conference on Multimedia and Expo (ICME)*, 1367–1372. IEEE.
- Ma, Y.; Nguyen, K. L.; Xing, F. Z.; and Cambria, E. 2020. A survey on empathetic dialogue systems. *Information Fusion*, 64: 50–70.
- Mai, S.; Zeng, Y.; Zheng, S.; and Hu, H. 2022. Hybrid contrastive learning of tri-modal representation for multimodal sentiment analysis. *IEEE Transactions on Affective Computing*, 14(3): 2276–2289.
- Mohammadi, G.; and Vuilleumier, P. 2020. A multi-componential approach to emotion recognition and the effect of personality. *IEEE Transactions on Affective Computing*, 13(3): 1127–1139.
- Morency, L.-P.; Mihalcea, R.; and Doshi, P. 2011. Towards multimodal sentiment analysis: Harvesting opinions from the web. In *Proceedings of the 13th international conference on multimodal interfaces*, 169–176.

- Nojavanasghari, B.; Gopinath, D.; Koushik, J.; Baltrušaitis, T.; and Morency, L.-P. 2016. Deep multimodal fusion for persuasiveness prediction. In *Proceedings of the 18th ACM international conference on multimodal interaction*, 284–288.
- Ponce-López, V.; Chen, B.; Oliu, M.; Corneanu, C.; Clapés, A.; Guyon, I.; Baró, X.; Escalante, H. J.; and Escalera, S. 2016. Chalearn lap 2016: First round challenge on first impressions-dataset and results. In *Computer Vision—ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8-10 and 15-16, 2016, Proceedings, Part III 14*, 400–418. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rahman, W.; Hasan, M. K.; Lee, S.; Zadeh, A.; Mao, C.; Morency, L.-P.; and Hoque, E. 2020. Integrating multimodal information in large pretrained transformers. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2020, 2359. NIH Public Access.
- Sun, Y.; Mai, S.; and Hu, H. 2022. Learning to learn better unimodal representations via adaptive multimodal meta-learning. *IEEE Transactions on Affective Computing*, 14(3): 2209–2223.
- Tsai, Y.-H. H.; Bai, S.; Liang, P. P.; Kolter, J. Z.; Morency, L.-P.; and Salakhutdinov, R. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for computational linguistics. Meeting*, volume 2019, 6558. NIH Public Access.
- Wei, X.; Zhang, T.; Li, Y.; Zhang, Y.; and Wu, F. 2020. Multi-modality cross attention network for image and sentence matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10941–10950.
- Xie, H.; Cui, J.; Cao, Y.; Chen, J.; Tao, J.; Fan, C.; Liu, X.; Wen, Z.; Lu, H.; Yang, Y.; et al. 2023. Multimodal Cross-Lingual Features and Weight Fusion for Cross-Cultural Humor Detection. In *Proceedings of the 4th on Multimodal Sentiment Analysis Challenge and Workshop: Mimicked Emotions, Humour and Personalisation*, 51–57.
- Yang, B.; Wu, L.; Zhu, J.; Shao, B.; Lin, X.; and Liu, T.-Y. 2022. Multimodal sentiment analysis with two-phase multi-task learning. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30: 2015–2024.
- Yang, J.; Yu, Y.; Niu, D.; Guo, W.; and Xu, Y. 2023. Confede: Contrastive feature decomposition for multimodal sentiment analysis. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7617–7630.
- Yang, K.; Xu, H.; and Gao, K. 2020. Cm-bert: Cross-modal bert for text-audio sentiment analysis. In *Proceedings of the 28th ACM international conference on multimedia*, 521–528.
- Yang, L.; Jiang, D.; and Sahli, H. 2018. Integrating deep and shallow models for multi-modal depression analysis—hybrid architectures. *IEEE Transactions on Affective Computing*, 12(1): 239–253.
- Zadeh, A.; Chen, M.; Poria, S.; Cambria, E.; and Morency, L.-P. 2017. Tensor fusion network for multimodal sentiment analysis. *arXiv preprint arXiv:1707.07250*.
- Zadeh, A.; Zellers, R.; Pincus, E.; and Morency, L.-P. 2016. Multimodal sentiment intensity analysis in videos: Facial gestures and verbal messages. *IEEE Intelligent Systems*, 31(6): 82–88.
- Zadeh, A. B.; Liang, P. P.; Poria, S.; Cambria, E.; and Morency, L.-P. 2018. 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)*, 2236–2246.
- Zeng, Y.; Yan, W.; Mai, S.; and Hu, H. 2024. Disentanglement translation network for multimodal sentiment analysis. *Information Fusion*, 102: 102031.
- Zhang, L.; Peng, S.; and Winkler, S. 2019. PersEmon: A deep network for joint analysis of apparent personality, emotion and their relationship. *IEEE Transactions on Affective Computing*, 13(1): 298–305.
- Zhi, Y.; Li, J.; Wang, H.; Chen, J.; and Cao, T. 2024. A Fine-Grained Tri-Modal Interaction Model for Multimodal Sentiment Analysis. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5715–5719. IEEE.
- Zhu, K.; Liu, X.; Xie, H.; Cai, C.; Fu, R.; Li, G.; Wen, Z.; Tao, J.; Fan, C.; Lv, Z.; et al. 2024. Transferring Personality Knowledge to Multimodal Sentiment Analysis. In *2024 IEEE 14th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, 431–435. IEEE.
- Zhu, L.; Zhao, H.; Zhu, Z.; Zhang, C.; and Kong, X. 2025. Multimodal sentiment analysis with unimodal label generation and modality decomposition. *Information Fusion*, 116: 102787.
- Zhu, L.; Zhu, Z.; Zhang, C.; Xu, Y.; and Kong, X. 2023. Multimodal sentiment analysis based on fusion methods: A survey. *Information Fusion*, 95: 306–325.