Multi-Modality Collaborative Learning for Sentiment Analysis

Shanmin Wang[12], Chengguang Liu[3], Qingshan Liu[12], *Senior Member, IEEE*

***Abstract*—Multimodal sentiment analysis (MSA) identifies individuals' sentiment states in videos by integrating visual, audio, and text modalities. Despite progress in existing methods, the inherent modality heterogeneity limits the effective capture of interactive sentiment features across modalities. In this paper, by introducing a Multi-Modality Collaborative Learning (MMCL) framework, we facilitate cross-modal interactions and capture enhanced and complementary features from modality-common and modality-specific representations, respectively. Specifically, we design a parameter-free decoupling module and separate uni-modality into modality-common and modality-specific components through semantics assessment of cross-modal elements. For modality-specific representations, inspired by the act-reward mechanism in reinforcement learning, we design policy models to adaptively mine complementary sentiment features under the guidance of a joint reward. For modality-common representations, intra-modal attention is employed to highlight crucial components, playing enhanced roles among modalities. Experimental results, including superiority evaluations on four databases, effectiveness verification of each module, and assessment of complementary features, demonstrate that MMCL successfully learns collaborative features across modalities and significantly improves performance. The code can be available at https://github.com/smwanghhh/MMCL.**

***Index Terms*—Multimodal sentiment analysis, Multi-modality collaboration, Enhanced features, Complementary features.**

I. Introduction

Multimodal sentiment analysis (MSA) aims to infer individuals' sentiment states from signals they emit, and has wide applications in human-computer interactions, intelligent healthcare, intelligent driving, etc [1]–[4]. Typically, individuals convey their inner states through a combination of facial expressions, spoken language, and acoustic behaviors, which can be captured from video clips [5]–[8]. Since these signals are driven by the same sentiment states, they definitely exhibit collaborative characteristics [9]. Facilitating interactions between these signals and capturing collaborative features can enhance and compensate sentiment expressions, making it one of the keys for MSA [10]–[12].

This work was supported by the National key research and development program 2021ZD0112200, the National Natural Science Foundation of China (NSFC) under Grant 61825601, and the Natural Science Foundation of Jiangsu Province (NSF-JS) under Grant BK20192004B.

[1] Nanjing University of Posts and Telecommunications. [2] Nanjing University of Aeronautics and Astronautics. [3] Nanjing University of Information Science and Technology.

Shanmin Wang, lecture in Nanjing University of Posts and Telecommunications and PhD graduate in Nanjing University of Aeronautics and Astronautics. E-mail: smwang1994@163.com.

Chengguang Liu, PhD candidate in the School of Computer Science in Nanjing University of Information Science and Technology. E-mail: lcg70kg@163.com.

Qingshan Liu, corresponding author, professor in Nanjing University of Posts and Telecommunications and PhD supervisor of Nanjing University of Aeronautics and Astronautics. E-mail: qsliu@njupt.edu.cn.

Modality fusion [13], [14] and representation disentanglement [15], [16] are popular MSA methods. Modality fusion leverages tensor-based [17], attention-based [18], or transfer-based [19] techniques to integrate representations from all modalities. Nevertheless, representations from various modalities often differ significantly in constituent elements, elements' distributions, noise levels, task relevance, etc, which are termed modality heterogeneity [20]. The heterogeneity of modalities hinders methods that directly integrate semantically unaligned representations from extracting interactive sentiment features across modalities seriously [16], [21], [22].

Representation disentanglement methods suggest that cross-modal interactive features include modality-common and modality-specific components [15], [16]. Modality-common features have shared semantics among modalities, such as rhythm, which can be expressed simultaneously through lip movements and sound waves. Modality-specific features are exclusively inherent to a particular modality, such as visual color properties. Based on this insight, these methods propose a new multimodal data processing paradigm that prioritizes disentanglement before fusion. Specifically, techniques such as subspace learning [15], [23], adversarial learning [16], [21], [24], and transfer learning [25]–[27] are employed to separate unimodal representations into modality-common and modality-specific components, with the intention of providing cross-modal enhancement and complementarity, respectively [21], [25], [28]. Subsequently, these disentangled components from all modalities are combined into a joint multimodal representation for prediction. However, two critical issues remain. First of all, modality-specific representations often fail to fully serve complementary roles due to their inclusion of task-irrelevant semantics, resulting in fragile multimodal representations. Besides, these disentanglement approaches rely heavily on the synergistic use of complex models paired with rigorous loss functions, leading to limited adaptability.

To alleviate the above issues, in this paper, we propose a Multi-Modality Collaborative Learning (MMCL) framework, capturing interactive sentiment features that enhance and complement across modalities. MMCL begins with a parameter-free decoupling module. Due to the inherent temporal asynchrony nature between modalities, the decoupling module assesses semantic correlations between temporal elements and divides unimodal representation into common and specific components, avoiding complex learning prerequisites. For decoupled specific representations, MMCL learns from the act-reward mechanism in reinforcement learning and captures complementary features adaptively. Specifically, it assigns a policy model for each specific representation for feature learning, and rewards learned features across modalities unitedly. The reward is incorporated into a centralized critic model

to evaluate and adjust policies jointly. For decoupled common features, MMCL highlights their important elements as cross-modal enhanced properties. For disentangled common features, MMCL highlights their important elements as cross-modal enhanced properties. Eventually, MMCL integrates enhanced and complementary features from all modalities into a joint representation and serves multiple downstream tasks. Our contributions can be concluded as follows.

- The MMCL framework captures collaborative properties that are enhanced and complementary across modalities to predict sentiment states. It captures collaborative properties through progressive representation decoupling and elaborate processing of decoupled features.
- We propose a parameter-free decoupling module that obtains common and specific representations by assessing the semantic correlation between cross-modal temporal elements, avoiding complex structural design and parameter learning.
- For learning collaborative cross-modal features, MMCL enhances vital common features and mines complementary properties from specific features. In complementary feature mining, the act-reward mechanism in reinforcement learning is utilized for adaptive feature learning, supported by a centralized critic model to coordinate feature-learning policies in multiple specific representations.

II. Related Work

Multimodal sentiment analysis (MSA) aims to determine individuals' sentiment states or emotional categories [29], [30]. In the early stages, researchers shift from unimodal to multimodal sentiment analysis [13], [14], [31], [32], leading to modality-fusion techniques becoming the focus.

Multimodal Fusion Strategies. Early and late fusion are two paradigms to integrate multiple modalities. Early-fusion approaches merge multiple representations into a joint one before making decisions. Late-fusion methods independently predict from uni-modalities and combine predicted vectors to make final decisions [33], [34]. Compared to late fusion, early fusion is more widely used. For example, TFN combined bi- or tri-modal data through tensor operations [32]. MFN and MARN integrated multiple signals over time via gated units and attention mechanisms [13], [31], respectively. MCTN performs cyclic translations among modalities and treats intermediate results as the joint multimodal representation [19]. MulT exploited multiple transformers to combine pairwise modalities by establishing long-time dependencies [14]. However, the gaps between heterogeneous signals pose significant challenges for these models in effectively capturing features for cross-modal interactions [16], [21], [22].

Disentangled Representation Learning. To efficiently facilitate cross-model interactions, recent studies have divided unimodal representations into modality-common and -specific components before modality fusion, aiming to provide enhanced and complementary sentimental clues, respectively [15], [25], [28]. To learn common and specific features, a variety of disentanglement techniques have emerged [35], [36]. Subspace learning, adversarial learning, and transfer learning are leading techniques for disentanglement. Subspace learning approaches employ distinct encoders paired with strict constraints to map raw inputs into common and specific subspaces [23]. For instance, MISA integrates similarity loss, orthogonal loss, and reconstruction loss to guide the learning of features in each branch [15]. Additionally, some research leverages adversarial learning based on encoders to further increase discrepancies between disentangled common and specific features. [16], [21], [24]. Several studies further conclude from previous research that the text modality provides higher-order semantic information compared to acoustic and visual modalities [37]. Based on this premise, they use the text modality as an anchor and divide the visual and acoustic data into common and specific features. For example, encoders, seq2seq, and transformer models are utilized in CRNet [27], TCSP [25], and CJTF [26], respectively, to transfer acoustic and visual data to the text modality.

The aforementioned models impose stringent constraints on common feature learning to ensure a high level of similarity. The residuals between inputs and common features are treated as specific features. Inheriting the semantic similarity requirement for common features from the above works and considering the temporal asynchronicity across modalities, we design a parameter-free decoupling module. It assesses the semantic similarities of each paired cross-modal element along the time dimension, which avoids complex parameter learning and structural design, getting better adaptability.

After disentanglement, common and specific from all modalities are combined through concatenation [23], [25]–[27], distribution alignment [16], or popular cross-modal attention [14], [15], [18], [24], [38]–[40]. The premise behind the above operations is that disentangled common and specific features perfectly play the role of cross-modal enhancement and complementarity, respectively. Nevertheless, according to the disentanglement scheme, specific features also convey task-irrelevant semantics. In this paper, inspired by adaptive policy learning under the rewarding mechanism in reinforcement learning [41]–[43], we design a specialized policy model for each modality-specific representation to mine complementary features.

III. Multi-Modality Collaborative Learning (MMCL)

A. Model Overview

In this section, we elaborate on the proposed MMCL method. The overall structure is illustrated in Figure 1. For the three input modalities $\{X^v, X^a, X^t\}$, we represent and map them into the same subspace as $\{Z^v, Z^a, Z^t \in \mathbb{R}^{L\times d^k}\}$, where L is the sequence length, and d^k is representations' dimension. Our objective is to coordinate multi-modalities to make accurate predictions. The collaboration between modalities contains enhanced and complementary features, which can be captured from common and specific representations, respectively. To this end, we first employ the Common-Specific Feature Decoupling (CSD) module to separate unimodal representations. Subsequently, crucial elements in decoupled common features are emphasized to promote mutual enhancement

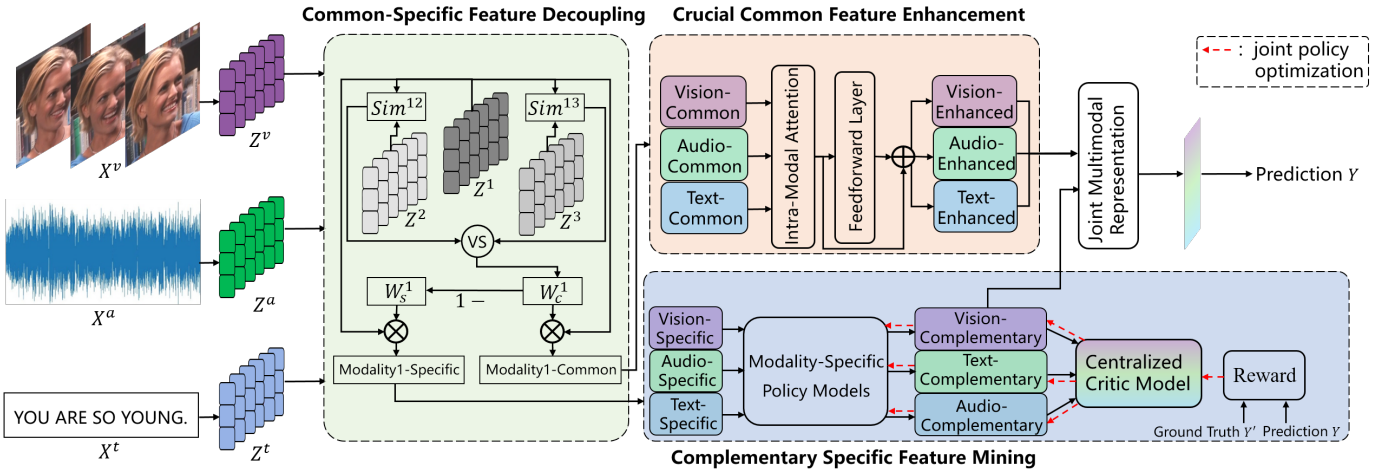

Fig. 1. The overall structure of the proposed MMCL. For multimodal sequences $\{X^v, X^a, X^t\}$, we initially map their features into the same subspace and obtain modality-common and modality-specific representations. Subsequently, MMCL captures enhanced and complementary features from modality-common and modality-specific representations, respectively. Ultimately, the enhanced and complementary representations are integrated into the joint multimodal representations for prediction. The decoupling module applies to three modalities simultaneously, obtaining common and specific representations for each modality.

among modalities. Regarding decoupled specific features, we exploit an act-reward mechanism to mine complementary components adaptively. Eventually, we integrate enhanced and complementary features from all modalities into a joint representation and apply it to predict sentiment states.

B. Common-Specific Representation Decoupling

Previous decoupling models, which rely on complex structural design and parameter learning, exhibit poor adaptability when incorporated with further collaborative feature learning. Building on the core principles of these decoupling models and accounting for temporal asynchrony across modalities, MMCL introduces a parameter-free approach based on the semantic similarities between cross-modal elements. For each unimodal representation, the CSD module decouples common features by calculating the semantic similarity between each of its temporal elements and those from other modalities. Features with low semantic similarity are then identified as modality-specific. Take the visual representation as an example ('1' of the CSD module in Figure 1 signifies 'v', '2' and '3' denote 'a' and 't', respectively.), its similarities with the text and audio modalities are calculated as follows.

$$\begin{aligned} Sim^{va} &= \frac{Z^v {Z^a}^T}{\|Z^v\| \, \|Z^a\|}; \\ Sim^{vt} &= \frac{Z^v {Z^t}^T}{\|Z^v\| \, \|Z^t\|}; \end{aligned} \tag{1}$$

where Sim^{va}, $Sim^{vt} \in \mathbb{R}^{L \times L}$. Elements in Sim^{va} and Sim^{vt} represent the cosine similarities between corresponding cross-modal elements. Elements exhibiting high similarity suggest a strong likelihood of being common across both modalities. To extend the similarities from two to three or more modalities, the CSD module incorporates a function F_{vs} to integrate similarity matrices.

$$(W_c^v)_{ij} = F_{vs}(Sim_{ij}^{va}, Sim_{ij}^{vt}); \forall i \in [0, L-1], \forall j \in [0, L-1]; \tag{2}$$

where F_{vs} is the comparison function that takes the smaller similarity score and can be expressed as follows.

$$F_{vs}(a, b) = \begin{cases} a, & if \quad a \leq b; \\ b, & otherwise; \end{cases} \tag{3}$$

In the CSD module, representations that show significant relevance across all modalities are treated as common through F_{vs}. Otherwise, they are marked as specific. In the experimental section, we also examine other comparison functions to explore broader common and specific features. Based on the similarity matrix W_c^v, the common features Z_c^v from the visual modality can be obtained as follows.

$$Z_c^v = W_c^v \cdot Z^v \tag{4}$$

The matrix reflecting the discrepancy between cross-modal elements is calculated as $W_s^v = 1 - W_c^v$. Based on the difference matrix W_s^v, the specific features can be obtained.

$$Z_s^v = W_s^v \cdot Z^v \tag{5}$$

During common and specific feature calculations, the sum of elements per row in W_c^v and W_s^v is limited to 1. The CSD module is also performed on text and audio modalities to obtain the decoupled representations, respectively.

C. Crucial Common Feature Enhancement

The common representations are shared across all modalities. For example, rhythm can be simultaneously depicted through visual cues like lip movements and acoustic signals such as sound waves. Common representations may express repeated content between modalities, impacting the efficiency of information processing. To mitigate this issue, we highlight

crucial common features as cross-modal enhanced properties as follows.

$$\tilde{Z}_c^v = F_{mh}^v(Z_c^v; \theta_{mh}^v) \tag{6}$$

As shown in Figure 1, the crucial common feature enhancement F_{mh}^v contains three operations. Intra-modal attention, i.e., self-attention, is first performed on common representations, followed by a feed-forward layer. Finally, a residual connection is established between the outputs of the former two layers. Similar functions F_{mh}^a and F_{mh}^t are also applied to acoustic and text common features to get $\tilde{Z}_c^a$ and $\tilde{Z}_c^t$, respectively.

D. Complementary Specific Feature Mining

Individuals' sentiments are time-varying, and the semantics conveyed by specific features from various modalities change at different moments. Dynamically adjusting the importance of these features enables effective utilization of the complementary strengths across modalities. For instance, audio features may dominate during moments of subdued tone, while visual features often provide richer cues during intense segments. Unfortunately, current modality disentanglement methods directly integrate specific features into a joint multimodal representation, presuming equal importance for all modalities at all time [16], [25]–[27]. This assumption risks obscuring critical emotional signals and limits the effective exploitation of modalities' complementary nature. To better utilize complementary features, we propose collaboratively learning the temporal importance of specific features but lack efficient supervised signals. In reinforcement learning, adaptive policy adjustments are facilitated by feedback from professional rewards without prior knowledge of optimal policies [43]. Inspired by this act-reward mechanism in reinforcement learning, we design modality-specific policy models to explore the temporal importance of features. By incorporating shared rewards, our method jointly optimizes policies to ensure each learned feature contributes valuable information during fusion.

As shown in Figure 1, we assign an independent policy model to each specific representation Z_s^i, taking actions to mine features $\tilde{Z}_s^i$ from it. To ensure that the mined features are complementary across modalities, we facilitate collaborative interaction among multiple specific representations Z_s^i using a centralized critic model. This centralized critic model gathers observed representations and actions from all modalities to evaluate the policy models from a global perspective. Based on this global evaluation, the policy models are optimized jointly, enabling the adaptive capture of complementary features.

Formally, each policy model takes observed specific representation Z_s^i as input and outputs action A^i according to current policy π^i.

$$A^i = F^{\pi^i}(Z_s^i; \theta^{\pi^i}) \tag{7}$$

where θ^{π^i} is the parameter of the policy model F^{π^i}. $A^i \in \mathbb{R}^L$ identifies complementary clues within Z_s^i. In practice, the policy model is implemented as a fully connected layer. After all policy models execute Eq. 7, the actions and observed representations are sent to the centralized critic model to obtain a global evaluation Q.

$$Q = F^{critic}(Z_s^v, A^v, Z_s^a, A^a, Z_s^t, A^t; \theta^{critic}) \tag{8}$$

The centralized critic model evaluates the rationality of policies by referring to the observed specific representations. It also takes into account the long-term impact of actions on the task, resulting in Q being the cumulative reward. Specifically, the critic model is implemented as an 8-head transformer that takes the concatenation of representations and actions as input. The critic model is optimized with the Temporal-Difference (TD) Error algorithm [43], [44]: $L^{critic} = Q - Q'$. Q' represents the ground truth of the cumulative reward and is calculated as $Q' = R + \gamma Q''$, Q'' is the cumulative reward in the next stage. R denotes the immediate reward, and γ is the discount factor.

Joint Policy Adjustment. Benefitting from the centralized critic model, each policy model strives to take action A^i in order to jointly maximize the cumulative reward. Thus, the objective function for policy models is $L^{policy} = -Q$. Take the vision modality as an example, its policy π^v is adjusted as follows.

$$\nabla_{\theta^{\pi^v}} L^{policy} = -\nabla_{\theta^{A^v}} Q \cdot \nabla_{\theta^{\pi^v}} A^v \tag{9}$$

where Q and A^v are obtained from Eq. 8 and Eq. 7, respectively. Clearly, observed specific representations and actions from audio and text modalities have explicitly contributed to adjusting the visual policy through Q. Both audio and text modalities update their policies in similar ways, suggesting that these modalities provide references for learning complementary features within each modality. The joint policy adjustment mechanism allows learned features to play compensating roles across modalities. Notably, the centralized critic model will be removed after training, leaving only the policy models for inference.

Reward. At each stage, we get complementary features by combining observed representations and actions: $\tilde{Z}_s^i = Z_s^i \times A^i$. Complementary features are further integrated with enhanced ones to form the joint representation: $Z = [\tilde{Z}_s^v, \tilde{Z}_c^v] * [\tilde{Z}_s^a, \tilde{Z}_c^a] * [\tilde{Z}_s^t, \tilde{Z}_c^t]$, where $[\cdot, \cdot]$ denotes concatenation, and $*$ signifies the modality fusion operation. The multimodal representation Z is then applied to downstream tasks to feed back a joint reward R for optimizing policy and critic models. R varies with tasks. For regression tasks, $R = -|Y' - Y|$, where Y' and Y are true and predicted sentiment states, respectively. For classification tasks, $R = \frac{e^{Y_i}}{\sum_j^C e^{Y_j}}$, where Y_i is the predicted probability for the true class, and C is the total number of categories. The calculation of rewards relies on both policy-critic and prediction modules. So, except for the loss from policy and critic modules, MMCL also involves a prediction loss L^p. The objective function of MMCL can be expressed as follows.

$$L = \alpha_1 L^p + \alpha_2 (L^{policy} + L^{critic}) \tag{10}$$

where α_1 and α_2 are weights for the prediction and policy-critic modules, respectively.

TABLE I
ABLATION STUDY ON THE MOSI AND IEMOCAP BENCHMARKS. CSD, CCE, AND CSM DENOTE COMMON-SPECIFIC FEATURE DECOUPLING, CRUCIAL COMMON FEATURE ENHANCEMENT, AND SPECIFIC FEATURE MINING MODULES, RESPECTIVELY.

Settings	IEMOCAP		MOSI				
	Acc (↑)	F1 (↑)	Acc7 (↑)	Acc2 (↑)	F1 (↑)	MAE (↓)	Corr (↑)
MMCL	**84.9**	**84.5**	**50.4**	**87.3**	**87.3**	**0.672**	**0.817**
Importance of Modules							
w/o CSD	81.9	81.2	47.0	85.7	85.6	0.714	0.804
w/o CCE	83.7	83.1	47.6	86.9	86.8	0.684	0.814
w/o CSM	83.1	82.4	47.2	85.7	85.6	0.701	0.806
Different Comparison Functions							
$F_{vs} = major$	83.7	83.2	49.2	86.3	86.2	0.706	0.802
$F_{vs} = mean$	83.2	82.6	48.2	85.5	85.4	0.714	0.796
Role of Representations							
Enhanced Features	83.2	82.1	47.0	86.6	86.5	0.695	0.805
Complementary Features	83.8	83.3	48.9	86.3	86.2	0.695	0.803
Importance of Modalities							
Text	81.4	80.8	47.0	85.2	85.1	0.734	0.795
Vision	76.5	72.8	20.3	57.6	55.4	1.440	0.088
Audio	80.3	78.8	21.3	55.4	54.6	1.470	0.062
Text + Vision	81.9	81.2	47.3	86.4	86.4	0.697	0.811
Text + Audio	83.3	82.7	47.7	85.5	85.5	0.704	0.806
Vision + Audio	81.1	79.7	21.6	57.0	55.7	1.430	0.116

IV. EXPERIMENT

A. Benchmark and Evaluation Metrics

We conduct extensive experiments on two Multimodal Sentiment Analysis (MSA), a Multimodal Emotion Recognition (MER), and a Multimodal Depression Assessment (MDA) benchmarks.

CMU-MOSI [17] and CMU-MOSEI [45] are widely used for MSA, collected from online sharing websites. The MOSI database contains 1281, 229, and 685 utterance-level sequences for training, validation, and testing, respectively. Each sentence is annotated with the sentiment scores ranging from -3(strongly negative) to +3(strongly positive). The MOSEI dataset comprises 16,265 utterances for training, 1,869 for validation, and 4,643 for testing.

IEMOCAP [46] is a laboratory-collected MER database. It has about 10K utterances labeled with nine emotions: angry, happy, sad, neutral, surprised, fearful, excited, frustrated, and other. Following the popular work [14], we take the first four emotions for experiments.

CMDC [47] is a popular MDA database that includes 78 subjects, each responding to 12 questions. 45 subjects are both audio and video recorded. The subjects' depression levels are evaluated using PHQ-9 scores [48], which range from 0 to 27 in total. We perform five-fold cross-validation on the CMDC database..

For the MSA task, we evaluate MMCL with (1) MAE: mean absolute error; (2) Corr: correlation between predictions and ground truth; (3) Acc2: binary accuracy, samples to be positive if its sentiment value is greater than 0, and vice versa; (4) F1 score; (5) Acc7: 7-class accuracy, we round up the predicted sentiment value for each sample as the class. For the MER task, we report the accuracy and F1 score for each category. For the MDA task, we focus on the regression metrics, including MAE, root mean squared error (RMSE), and Pearson correlation coefficients. Individuals scoring below 9 are classified as normal, while those scoring 9 or above are considered depressed. We also provide classification metrics such as precision, recall, and F-measure.

B. Implemented Details

Following popular works [14], [47], [49]–[52], the pre-trained BERT model [53] is used for extracting 768-dimensional text features for MOSI and MOSEI databases. For the IEMOCAP database, we utilize 768-dimensional Glove features to represent text data. For audio data, COVAREP is used to extract 12-Mel frequency cepstral coefficients, pitch tracking, spectral envelope, and so on, resulting in 74-dimensional representations. Visual data is processed with Facet to capture action units, facial landmarks, head pose, etc., forming 47-dimensional, 35-dimensional, and 35-dimensional features for the MOSI, MOSEI, and IEMOCAP databases, respectively. For the CMDC database, we use the provided 768-dimensional BERT, 768-dimensional TimesFormer [54], and 128-dimensional Viggish features [55] for three modalities. We map three modalities into a 256-dimensional subspace. Without loss of generality, we merge three modalities through weighted summation and concatenation, respectively. The discounted factor γ is set as 0.5. For MSA and MDA

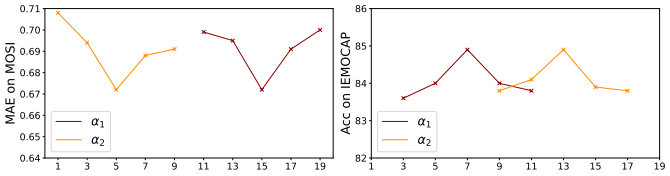

Fig. 2. Ablation studies on hyper-parameters. α_1 and α_2 are weights for the prediction and policy-critic modules, respectively.

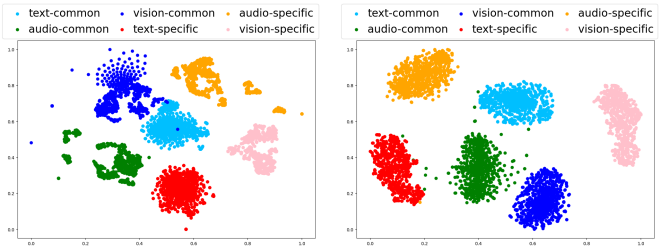

(a) MISA decoupled representations. (b) MMCL decoupled representations.

Fig. 3. Visualization of decoupled representations on the IEMOCAP benchmark.

tasks, weights are set as $\alpha_1 = 15$ and $\alpha_2 = 5$. For the MER task, weights are set as $\alpha_1 = 7$ and $\alpha_2 = 13$. During model training, the batch is set to 64, 128, and 128 for MSA, MER, and MDA tasks, respectively. The model is trained for 200 epochs on 2080Ti, using the Adam optimizer.

C. Ablation Study

We conduct comprehensive experiments on the MOSI and IEMOCAP benchmarks to provide a thorough analysis of MMCL. These experiments include the ablation study on each module, analysis of various comparison functions F_{vs}, examination of the role of enhanced and complementary features, as well as assessment of the importance of each modality. Besides, we also present the performance under different hyper-parameters and visualized representation comparisons.

Ablation studies on Hyper-parameters. MMCL contains objective functions for the prediction and policy-critic modules, weighted by α_1 and α_2, respectively. To create a more rigorous experimental setup, we conduct extensive experiments across a wide range of values for two weights. Figure 2 shows the weights and their corresponding performance. MMCL achieves the optimal MAE for the MOSI database when $\alpha_1 = 15$, $\alpha_2 = 5$. Weights for the MOSEI and CMDC databases align with those used for MOSI. For the IEMOCAP database, MMCL reaches the best accuracy with $\alpha_1 = 7$, $\alpha_2 = 13$.

Importance of Modules. In Table I, we analyze the importance of each module by removing them one at a time. First, we eliminate the decoupling module (CSD) and proceed with enhanced and complementary feature learning on the raw input. Without the decoupling module, the subsequent operations fail to perform effectively, resulting in a notable drop in performance compared to MMCL. Next, we remove the enhanced feature learning module (CCE) and directly integrate the decoupled common features into the joint representation, which leads to a slight decrease in performance. Finally, we

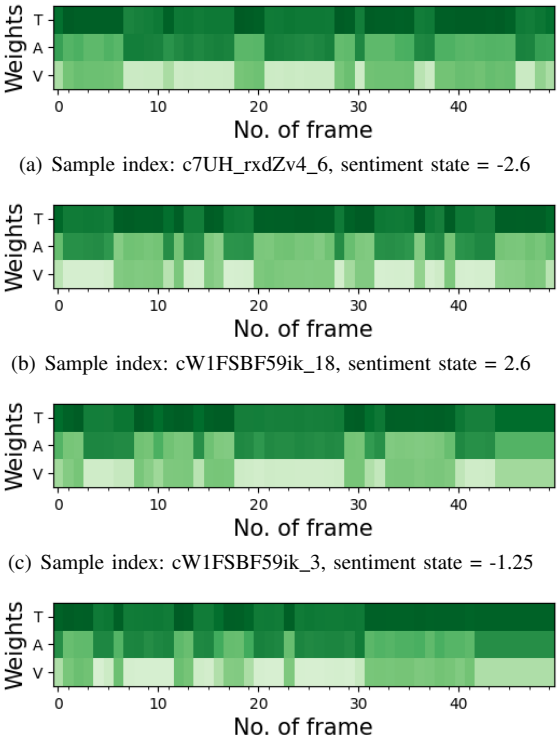

(a) Sample index: c7UH_rxdZv4_6, sentiment state = -2.6

(b) Sample index: cW1FSBF59ik_18, sentiment state = 2.6

(c) Sample index: cW1FSBF59ik_3, sentiment state = -1.25

(d) Sample index: tIrG4oNLFzE_6, sentiment state = 1.25

Fig. 4. Temporal weights of four samples' specific features on the MOSI database. "T", "A", and "V" denote text, audio, and vision modalities, respectively.

eliminate the complementary feature mining module (CSM) and treat the decoupled specific features as complementary. This operation significantly impacts performance by introducing task-irrelevant modality information and noise.

Different Comparison Functions F_{vs}. We also investigate various comparison functions F_{vs} during representation decoupling, which include taking the minor, major, or average similarity score, marked as 'minor', 'major', and 'mean', respectively. MMCL in Table I corresponds to the setting of $F_{vs} = minor$. With $F_{vs} = minor$, only elements that are similar across all modalities are considered common features. Instead, $F_{vs} = major$ identifies features as modality-common if they are similar to at least one modality. $F_{vs} = mean$ calculates the average similarity score. Among these functions, $F_{vs} = mean$ performs the worst because an element's lack of relevance to one modality can be offset by high similarity to another, introducing noise into the common features. A similar issue arises with $F_{vs} = major$. $F_{vs} = minor$ imposes stricter criteria for identifying common features and achieves the best performance among the three functions.

Visualization of Representations. Decoupling representation into common and specific components is an important premise for MMCL. Therefore, we employ the T-SNE technique [56] to visualize the decoupled representations in Figure 3. For comparison, we also visualize the representations obtained by MISA, a popular decoupled model. By inheriting MISA's core principles and avoiding complex parameter learning, features decoupled by MMCL even exhibit better separability.

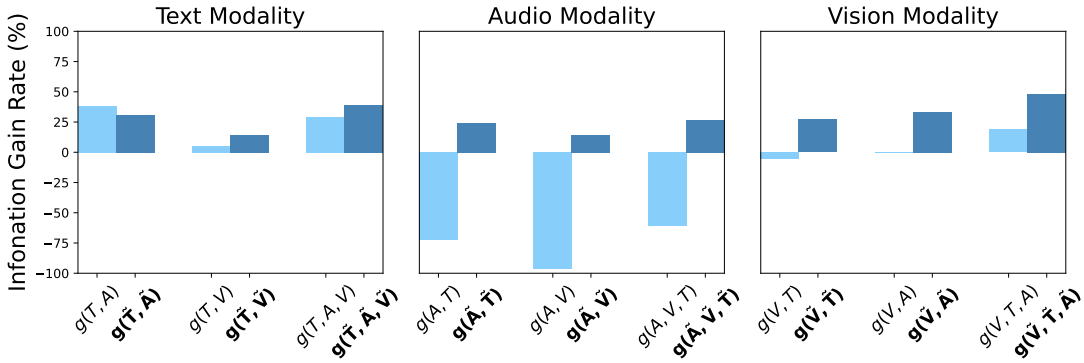

Fig. 5. Information gain rate comparisons between specific and complementary representations. For brevity, features from the vision, audio, and text modalities are denoted as "V", "A", and "T", respectively. Complementary features are bolded for distinction.

Role of Representations. We also analyze the roles of enhanced and complementary features across two tasks. Table I presents the results of independently using enhanced and complementary features for predictions, respectively. On the MOSI database, both types of representations contribute similarly to overall performance. Yet, in the 7-class classification task, complementary features significantly outperform enhanced ones. On the IEMOCAP database, complementary features also demonstrate higher recognition rates compared to enhanced features.

Importance of Modalities. We further explore the performance of bi-modal MMCL. To enable effective comparison, we first present the unimodal performance in Table I. For instance, the text modality shows varying degrees of improvement when the MMCL framework is applied to promote its interaction with visual and acoustic modalities, respectively. Similar improvements are observed in the visual and acoustic modalities as well. This comparative analysis highlights MMCL's versatility across different modalities. Additionally, by examining the results of uni-, bi-, and tri-modal MMCL, we conclude that (1) the text modality plays a crucial role, and (2) each modality is essential, with their integration achieving the best performance.

Temporal Weights of Specific Features. As shown in Figure 4, we select four samples with various sentiment states from the MOSI database and visualize their temporal weights for modality-specific features. On a macro level, the text modality consistently holds larger weights than the other two modalities during interactions. This can be attributed to two factors: first, text contains intuitive emotional descriptors; and second, unlike audio and visual signals, which rely on extracted features, raw spoken words are publicly available and can be directly used in end-to-end training of the MSA model without privacy restrictions. The modality weights also align with the modality importance in Table I. More importantly, the temporal weights learned by the MMCL model show that the importance of each modality-specific feature varies over time, reflecting the dynamic interactions in joint multimodal representations. Furthermore, these interactive weights differ across samples, being fully determined by the input videos. These findings suggest that the proposed MMCL model can adaptively adjust modality importance in real time during emotional expression.

TABLE II
COMPARISON WITH SOTA METHODS ON THE MOSI AND MOSEI BENCHMARKS. 'A' AND 'C' DENOTE WEIGHTED SUMMATION AND CONCATENATION MODALITY FUSIONS, RESPECTIVELY.

	Acc7 (↑)	Acc2 (↑)	F1 (↑)	MAE (↓)	Corr (↓)
		MOSI			
TFN [32]	44.7	82.6	82.6	0.761	0.789
LMF [57]	45.1	84.0	84.0	0.742	0.785
MulT [14]	41.5	83.7	83.7	0.767	0.799
MAG [49]	42.9	83.5	83.5	0.790	0.769
MISA [15]	42.3	83.4	83.6	0.783	0.761
TFR-Net [58]	42.6	84.0	83.9	0.787	0.788
HyCon [59]	46.6	85.2	85.1	0.713	0.790
MCL [51]	49.2	86.1	86.1	0.713	0.793
MHE [60]	41.5	83.6	83.5	0.801	0.722
CMHFM [36]	37.0	81.0	81.3	0.912	0.677
CRNet [27]	47.4	86.4	86.4	0.712	0.797
TMBL [23]	36.3	83.8	84.3	0.867	0.762
DTN [16]	48.1	86.2	86.2	0.714	0.807
MMCL(c)	49.9	86.7	86.7	0.681	0.811
MMCL(a)	**50.4**	**87.3**	**87.3**	**0.672**	**0.817**
		MOSEI			
TFN [32]	51.8	84.5	84.5	0.622	0.781
LMF [57]	51.2	84.2	84.3	0.612	0.779
MulT [14]	50.7	84.7	84.6	0.625	0.775
MAG [49]	51.9	85.0	85.0	0.602	0.778
MISA [15]	52.2	85.5	85.3	0.555	0.756
TFR-Net [58]	51.7	85.2	85.1	0.606	0.781
HyCon [59]	52.8	85.4	85.6	0.601	0.776
MCL [51]	53.3	86.2	86.2	0.581	0.791
MHE [60]	52.5	84.2	84.0	0.577	0.712
CMHFM [36]	52.6	84.0	83.6	0.558	0.731
CRNet [27]	53.8	86.2	86.1	0.541	0.771
TMBL [23]	52.4	85.8	85.9	0.545	0.766
DTN [16]	52.5	86.3	86.3	0.579	0.788
MMCL(c)	54.5	**86.6**	**86.5**	**0.528**	0.794
MMCL(a)	**54.7**	86.5	**86.5**	**0.528**	**0.799**

Evaluation on Complementary Features. We calculate the information gain rate g [61] for specific and complementary features to investigate cross-modal complementarity. First,

TABLE III
COMPARISON WITH SOTA METHODS ON THE IEMOCAP BENCHMARK. "A" AND "C" DENOTE WEIGHTED SUMMATION AND CONCATENATION MODALITY FUSIONS, RESPECTIVELY.

	Happy		Sad		Angry		Neutral		Average	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
MFN [31]	86.5	84.0	83.5	82.1	85.0	83.7	69.6	69.2	81.2	79.8
LMF [57]	86.9	82.3	85.4	84.7	87.1	86.8	71.6	71.4	82.8	81.3
MulT [14]	87.4	84.1	84.2	83.1	88.0	87.5	69.9	68.4	82.4	80.8
MISA [15]	86.1	80.8	82.3	79.1	84.1	83.8	69.0	67.7	80.4	77.8
HyCon [59]	88.0	85.5	86.2	85.9	89.4	89.2	70.4	70.5	83.5	82.8
MCL [51]	**88.8**	86.8	86.6	86.6	90.3	90.3	71.6	71.4	84.3	83.8
UniMF [62]	83.4	85.3	82.9	84.0	84.0	83.2	69.5	70.0	80.0	80.6
TLRF [63]	84.2	86.1	84.9	85.0	87.9	87.7	72.9	73.1	82.5	83.0
TMBL [23]	86.4	85.1	82.8	83.0	86.0	86.3	69.9	69.5	81.3	81.0
MMCL(c)	**88.8**	**87.8**	**87.5**	**87.6**	89.3	89.3	73.1	73.2	84.7	**84.5**
MMCL(a)	88.4	87.3	87.3	87.0	**90.6**	**90.4**	**73.3**	**73.3**	**84.9**	**84.5**

we extract specific Z_s and complementary features $\tilde{Z}_s$ from MMCL trained on the IEMOCAP benchmark. Specialized classifiers are then trained to get predicted vectors from Z_s and $\tilde{Z}_s$. Take Z_s^v as an example. Its predicted result is obtained: $Y^v = F^v(Z_s^v; \theta^v)$. The entropy is further calculated: $H(Z_s^v) = -\sum_{j=1}^{C} Y_j^v \log_2 Y_j^v$, which represents the predicted uncertainty using Z_s^v. After adding the audio-specific features as a condition, another predicted vector Y^{va} and the conditional entropy $H(Z_s^v \,|Z_s^a)$ are calculated. Eventually, we compute the information gain rate of Z_s^a relative to Z_s^v: $g(Z_s^v, Z_s^a) = \frac{H(Z_s^v) - H(Z_s^v|Z_s^a)}{H(Z_s^v)}$. Information gain rate for other modalities and features can be obtained similarly. In Figure 5, we compare g between specific and complementary features. For specific features, the audio modality predominantly reduces predictive uncertainty. However, when other modalities are added to audio, g becomes negative, suggesting that these modalities either lack complementary information or inject noise. In contrast, complementary features exhibit positive g for each other, demonstrating strong compensation effects among the modalities, with the visual modality benefiting the most. These results confirm that we have effectively captured mutually compensating features as anticipated.

D. Comparison With State-of-the-Art Methods

We compare MMCL with some state-of-the-art models, including TFN [32], LMF [57], MulT [14], MAG [49], MISA [15], TFR-Net [58], HyCon [59], MCL [51], MHE [60], CMHFM [36], CRNet [27], TMBL [23], UniMF [62], TLRF [63], DTN [16]. For fair comparisons, we compare BERT-based models for MSA and MDA, Glove-based models for MER, respectively.

Multimodal Sentiment Analysis. Comparative results on the MOSI and MOSEI benchmarks are reported in Table II. Compared to the recent DTN and popular MISA, which disentangles representations via adversarial learning and subspace learning, respectively, MMCL obtains common and specific features through the simple decoupling process and further captures collaborative properties. Compared to MCL, which retains modality-specific information via contrastive learning, MMCL effectively differentiates between specific and complementary features. Finally, both fusion modes of MMCL demonstrate superior performance on the MOSI and MOSEI benchmarks.

TABLE IV
COMPARISON WITH SOTA METHODS ON THE CMDC BENCHMARK. "A" AND "C" DENOTE WEIGHTED SUMMATION AND CONCATENATION MODALITY FUSIONS, RESPECTIVELY.

	MAE (↓)	RMSE (↓)	Pearson (↑)	Precision (↑)	Recall (↑)	F1 (↑)
Bi-LSTM [64]	4.55	5.67	0.68	0.87	0.89	0.88
MulT [14]	4.32	5.61	0.72	0.97	0.85	0.91
MISA [15]	2.47	3.43	0.88	**1.00**	**0.93**	**0.96**
TMBL [23]	3.09	3.98	0.90	**1.00**	**0.93**	**0.96**
MMCL(c)	2.31	3.43	**0.92**	**1.00**	**0.93**	**0.96**
MMCL(a)	**2.27**	**2.93**	**0.92**	**1.00**	**0.93**	**0.96**

Multimodal Emotion Recognition. Results on IEMOCAP are presented in Table III. Among the compared models, MCL gets the best average accuracy and F1 score. MMCL significantly outperforms MCL with arbitrary fusion mode. In addition, compared with previous methods, MMCL has the greatest improvement in recognizing neutral expressions.

Multimodal Depression Assessment. Comparative results on CMDC are shown in Table IV. CMDC is a new multimodal database for depression assessment. We reimplement several popular multimodal models, including Bi-LSTM and MulT, and representation decoupling models, including MISA and TMBL, to compare with MMCL. Among the compared methods, MMCL consistently achieves superior results, particularly on regression metrics.

V. CONCLUSION

In this paper, we capture enhanced and complementary collaborative features among modalities to analyze senti-

ment states. For collaborative sentimental feature learning, the proposed MMCL model first decouples unimodal representations into common and specific components, and then learns enhanced and complementary properties from these decoupled features. Our main contributions include highlighting the distinction between collaborative attributes and decoupled representations, as well as implementing adaptive complementary feature mining using the rewarding mechanism in reinforcement learning. Experimental evaluations of the information gain rate demonstrate that MMCL successfully mines complementary properties from specific representations.

REFERENCES

[1] S. Wang, H. Shuai, C. Liu, and Q. Liu, "Bias-based soft label learning for facial expression recognition," *IEEE Transactions on Affective Computing*, 2022.

[2] Q. Liu *et al.*, "Phase space reconstruction driven spatio-temporal feature learning for dynamic facial expression recognition," *IEEE Transactions on Affective Computing*, 2020.

[3] P. Song, D. Guo, X. Yang, S. Tang, E. Yang, and M. Wang, "Emotion-prior awareness network for emotional video captioning," in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 589–600.

[4] Z. Zhao and Q. Liu, "Former-dfer: Dynamic facial expression recognition transformer," in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 1553–1561.

[5] R. Chen, W. Zhou, Y. Li, and H. Zhou, "Video-based cross-modal auxiliary network for multimodal sentiment analysis," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 12, pp. 8703–8716, 2022.

[6] Y. Ou, Z. Chen, and F. Wu, "Multimodal local-global attention network for affective video content analysis," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 31, no. 5, pp. 1901–1914, 2020.

[7] Y. Dai, Y. Li, D. Chen, J. Li, and G. Lu, "Multimodal decoupled distillation graph neural network for emotion recognition in conversation," *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.

[8] H. Zhou, S. Huang, F. Zhang, and C. Xu, "Ceprompt: Cross-modal emotion-aware prompting for facial expression recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.

[9] C. E. Izard, "Differential emotions theory," *Human emotions*, pp. 43–66, 1977.

[10] L. He, Z. Wang, L. Wang, and F. Li, "Multimodal mutual attention-based sentiment analysis framework adapted to complicated contexts," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 33, no. 12, pp. 7131–7143, 2023.

[11] M. Hou, Z. Zhang, C. Liu, and G. Lu, "Semantic alignment network for multi-modal emotion recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 33, no. 9, pp. 5318–5329, 2023.

[12] D. Yang, M. Li, L. Qu, K. Yang, P. Zhai, S. Wang, and L. Zhang, "Asynchronous multimodal video sequence fusion via learning modality-exclusive and-agnostic representations," *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.

[13] A. Zadeh, P. P. Liang, S. Poria, P. Vij, E. Cambria, and L.-P. Morency, "Multi-attention recurrent network for human communication comprehension," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.

[14] 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.

[15] D. Hazarika, R. Zimmermann, and S. Poria, "Misa: Modality-invariant and-specific representations for multimodal sentiment analysis," in *Proceedings of the 28th ACM international conference on multimedia*, 2020, pp. 1122–1131.

[16] Y. Zeng, W. Yan, S. Mai, and H. Hu, "Disentanglement translation network for multimodal sentiment analysis," *Information Fusion*, vol. 102, p. 102031, 2024.

[17] A. Zadeh, R. Zellers, E. Pincus, and L.-P. Morency, "Multimodal sentiment intensity analysis in videos: Facial gestures and verbal messages," *IEEE Intelligent Systems*, vol. 31, no. 6, pp. 82–88, 2016.

[18] M. S. Akhtar, D. S. Chauhan, D. Ghosal, S. Poria, A. Ekbal, and P. Bhattacharyya, "Multi-task learning for multi-modal emotion recognition and sentiment analysis," in *Proceedings of NAACL-HLT*, 2019, pp. 370–379.

[19] H. Pham, P. P. Liang, T. Manzini, L.-P. Morency, and B. Póczos, "Found in translation: Learning robust joint representations by cyclic translations between modalities," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, 2019, pp. 6892–6899.

[20] T. Baltrušaitis, C. Ahuja, and L.-P. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 2, pp. 423–443, 2018.

[21] S. Mai, H. Hu, and S. Xing, "Modality to modality translation: An adversarial representation learning and graph fusion network for multimodal fusion," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 01, 2020, pp. 164–172.

[22] Y. Zeng, Z. Li, Z. Chen, and H. Ma, "A feature-based restoration dynamic interaction network for multimodal sentiment analysis," *Engineering Applications of Artificial Intelligence*, vol. 127, p. 107335, 2024.

[23] J. Huang, J. Zhou, Z. Tang, J. Lin, and C. Y.-C. Chen, "Tmbl: Transformer-based multimodal binding learning model for multimodal sentiment analysis," *Knowledge-Based Systems*, vol. 285, p. 111346, 2024.

[24] D. Yang, S. Huang, H. Kuang, Y. Du, and L. Zhang, "Disentangled representation learning for multimodal emotion recognition," in *Proceedings of the 30th ACM International Conference on Multimedia*, 2022, pp. 1642–1651.

[25] Y. Wu, Z. Lin, Y. Zhao, B. Qin, and L.-N. Zhu, "A text-centered shared-private framework via cross-modal prediction for multimodal sentiment analysis," in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2021, pp. 4730–4738.

[26] Q. Lu, X. Sun, Z. Gao, Y. Long, J. Feng, and H. Zhang, "Coordinated-joint translation fusion framework with sentiment-interactive graph convolutional networks for multimodal sentiment analysis," *Information Processing & Management*, vol. 61, no. 1, p. 103538, 2024.

[27] H. Shi, Y. Pu, Z. Zhao, J. Huang, D. Zhou, D. Xu, and J. Cao, "Co-space representation interaction network for multimodal sentiment analysis," *Knowledge-Based Systems*, vol. 283, p. 111149, 2024.

[28] D. Chen, W. Su, P. Wu, and B. Hua, "Joint multimodal sentiment analysis based on information relevance," *Information Processing & Management*, vol. 60, no. 2, p. 103193, 2023.

[29] F. Qian, J. Han, Y. Guan, W. Song, and Y. He, "Capturing high-level semantic correlations via graph for multimodal sentiment analysis," *IEEE Signal Processing Letters*, 2024.

[30] Z. Xie, Y. Yang, J. Wang, X. Liu, and X. Li, "Trustworthy multimodal fusion for sentiment analysis in ordinal sentiment space," *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.

[31] A. Zadeh, P. P. Liang, N. Mazumder, S. Poria, E. Cambria, and L.-P. Morency, "Memory fusion network for multi-view sequential learning," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.

[32] A. Zadeh, M. Chen, S. Poria, E. Cambria, and L.-P. Morency, "Tensor fusion network for multimodal sentiment analysis," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 1103–1114.

[33] V. Lopes, A. Gaspar, L. A. Alexandre, and J. Cordeiro, "An automl-based approach to multimodal image sentiment analysis," in *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021, pp. 1–9.

[34] O. Kampman, E. J. Barezi, D. Bertero, and P. Fung, "Investigating audio, visual, and text fusion methods for end-to-end automatic personality prediction," *arXiv preprint arXiv:1805.00705*, 2018.

[35] A. Ando, R. Masumura, A. Takashima, S. Suzuki, N. Makishima, K. Suzuki, T. Moriya, T. Ashihara, and H. Sato, "On the use of modality-specific large-scale pre-trained encoders for multimodal sentiment analysis," in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 739–746.

[36] L. Wang, J. Peng, C. Zheng, T. Zhao *et al.*, "A cross modal hierarchical fusion multimodal sentiment analysis method based on multi-task learning," *Information Processing & Management*, vol. 61, no. 3, p. 103675, 2024.

[37] J. Huang, Y. Ji, Y. Yang, and H. T. Shen, "Cross-modality representation interactive learning for multimodal sentiment analysis," in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 426–434.

[38] A. Shenoy, A. Sardana, and N. Graphics, "Multilogue-net: A context aware rnn for multi-modal emotion detection and sentiment analysis in conversation," *ACL 2020*, p. 19, 2020.

[39] S. Rahmani, S. Hosseini, R. Zall, M. R. Kangavari, S. Kamran, and W. Hua, "Transfer-based adaptive tree for multimodal sentiment analysis based on user latent aspects," *Knowledge-Based Systems*, vol. 261, p. 110219, 2023.

[40] H. Sun, J. Liu, Y.-W. Chen, and L. Lin, "Modality-invariant temporal representation learning for multimodal sentiment classification," *Information Fusion*, vol. 91, pp. 504–514, 2023.

[41] A. Tampuu, T. Matiisen, D. Kodelja, I. Kuzovkin, K. Korjus, J. Aru, J. Aru, and R. Vicente, "Multiagent cooperation and competition with deep reinforcement learning," *PloS one*, vol. 12, no. 4, p. e0172395, 2017.

[42] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," *Advances in neural information processing systems*, vol. 12, 1999.

[43] S. S. Mousavi, M. Schukat, and E. Howley, "Deep reinforcement learning: an overview," in *Proceedings of SAI Intelligent Systems Conference (IntelliSys) 2016: Volume 2*. Springer, 2018, pp. 426–440.

[44] R. S. Sutton, "Learning to predict by the methods of temporal differences," *Machine learning*, vol. 3, pp. 9–44, 1988.

[45] 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.

[46] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, pp. 335–359, 2008.

[47] B. Zou, J. Han, Y. Wang, R. Liu, S. Zhao, L. Feng, X. Lyu, and H. Ma, "Semi-structural interview-based chinese multimodal depression corpus towards automatic preliminary screening of depressive disorders," *IEEE Transactions on Affective Computing*, 2022.

[48] K. Kroenke, R. L. Spitzer, and J. B. Williams, "The phq-9: validity of a brief depression severity measure," *Journal of general internal medicine*, vol. 16, no. 9, pp. 606–613, 2001.

[49] W. Rahman, M. K. Hasan, S. Lee, A. Zadeh, C. Mao, L.-P. Morency, and E. Hoque, "Integrating multimodal information in large pretrained transformers," in *Proceedings of the conference. Association for Computational Linguistics. Meeting*, vol. 2020. NIH Public Access, 2020, p. 2359.

[50] M. K. Hasan, S. Lee, W. Rahman, A. Zadeh, R. Mihalcea, L.-P. Morency, and E. Hoque, "Humor knowledge enriched transformer for understanding multimodal humor," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 14, 2021, pp. 12 972–12 980.

[51] S. Mai, Y. Sun, Y. Zeng, and H. Hu, "Excavating multimodal correlation for representation learning," *Information Fusion*, vol. 91, pp. 542–555, 2023.

[52] S. Mai, S. Xing, and H. Hu, "Analyzing multimodal sentiment via acoustic-and visual-lstm with channel-aware temporal convolution network," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1424–1437, 2021.

[53] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[54] G. Bertasius, H. Wang, and L. Torresani, "Is space-time attention all you need for video understanding?" in *ICML*, vol. 2, no. 3, 2021, p. 4.

[55] X. Li, "Xmnlp: A lightweight chinese natural language processing toolkit," *GitHub*, 2018.

[56] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of machine learning research*, vol. 9, no. 11, 2008.

[57] Z. Liu, Y. Shen, V. B. Lakshminarasimhan, P. P. Liang, A. Zadeh, and L.-P. Morency, "Efficient low-rank multimodal fusion with modality-specific factors," *arXiv preprint arXiv:1806.00064*, 2018.

[58] Z. Yuan, W. Li, H. Xu, and W. Yu, "Transformer-based feature reconstruction network for robust multimodal sentiment analysis," in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 4400–4407.

[59] S. Mai, Y. Zeng, S. Zheng, and H. Hu, "Hybrid contrastive learning of tri-modal representation for multimodal sentiment analysis," *IEEE Transactions on Affective Computing*, 2022.

[60] Y. Gao, Y. Fu, M. Sun, and F. Gao, "Multi-modal hierarchical empathetic framework for social robots with affective body control," *IEEE Transactions on Affective Computing*, 2024.

[61] J. R. Quinlan, "Induction of decision trees," *Machine learning*, vol. 1, pp. 81–106, 1986.

[62] R. Huan, G. Zhong, P. Chen, and R. Liang, "Unimf: A unified multimodal framework for multimodal sentiment analysis in missing modalities and unaligned multimodal sequences," *IEEE Transactions on Multimedia*, 2023.

[63] L. Zhao, Y. Yang, and T. Ning, "A three-stage multimodal emotion recognition network based on text low-rank fusion," *Multimedia Systems*, vol. 30, no. 3, p. 142, 2024.

[64] S. Zhang, D. Zheng, X. Hu, and M. Yang, "Bidirectional long short-term memory networks for relation classification," in *Proceedings of the 29th Pacific Asia conference on language, information and computation*, 2015, pp. 73–78.