

in its weakened position across all modalities. Hence, from the outset, it is crucial to enhance representations for each modality, regardless of whether they are weak or strong, as it can affect the imbalance in learning across modalities.

Moreover, current methodologies primarily focus on interactions between pairs of modalities [6, 22, 25, 39], resulting in complex computations and inadequate treatment across all modalities. These methods are commonly applied in tasks such as audio-visual learning [25, 39] and multimodal affective computing [44], often using datasets related to sarcasm detection, sentiment analysis, or humor detection. However, there is a lack of methods explicitly tailored for multimodal emotion recognition in conversation, especially for well-known multimodal datasets like IEMOCAP [3], MELD [27], and CMU-MOSEI [1]. Additionally, in recent prominent studies [17, 32], while overall performance for multimodal emotion recognition tasks has notably increased, a closer examination of the “importance of modality” reveals that pairwise modalities consistently fail to achieve satisfactory performance, creating a significant gap compared to leveraging all three modalities simultaneously. Therefore, it is crucial to simultaneously leverage learning from all modalities while also significantly enhancing the capabilities of weaker modalities to improve the overall learning performance of multimodal emotion recognition models in practical applications.

In this paper, we propose a novel framework named **Ada2I** that addresses imbalances in learning across audio, text, and visual modalities for multimodal ERC. It consists of two primary modules including **Adaptive Feature Weighting (AFW)** and **Adaptive Modality Weighting (AMW)** for *feature-level* and *modality-level* balancing respectively in the consideration of **Inter-** and **Intra-**modal interactions. Focusing on *feature-level balancing* using Adaptive Feature Weighting (AFW), we apply tensor contraction to infer feature-aware attention weights for each modality, which aims to produce a feature-level balanced representation for each conversation. As an important component of AFW, Attention Mapping Network controls the balancing via maximizing the alignment between unimodal features and their corresponding attention coefficients. For *modality-level balancing* using Adaptive Modality Weighting (AMW), we further exploit feature-level balanced representations from the preceding AFW module to generate modality-level balanced ones through modality-wise normalization of features and learning weights before being used to enhance the emotion recognition. Additionally, we utilize the concept of disparity ratio, although with modifications compared to the study by Peng et al. [25], called OGM-GE, as a value to supervise the training process and evaluate the model. Specifically, while OGM-GE [25] introduced gradient modulation for pairs of modalities, we refine it to handle all three modalities simultaneously—textual, visual, and audio—in the Multimodal emotion recognition in conversation task. This adjustment reduces model complexity and overall processing time, leading to enhanced efficiency.

To summarize, our contributions are as follows:

- We propose an end-to-end framework named **Ada2I** that addresses the issue of imbalance learning across modalities comprehensively for the multimodal ERC task. It not only considers modality-level imbalances but also leverages feature-level representations to contribute to the balancing step in the learning process.

- With two modules intricately designed yet inseparable, Adaptive Feature Weighting (AFW) is crafted to enhance the representation of each conversation at the feature level, while Adaptive Modality Weighting (AMW) is proposed to optimize the modality-level learning weights during training. Additionally, we redefine the disparity ratio, a simple yet effective measure, to assess the overall discrepancy of the model’s learning process when simultaneously handling multiple modalities, rather than just two as in the original approach from [25].
- Our empirical experiments illustrate the effectiveness and enhancements provided by our method in comparison to existing approaches across three prevalent multimodal ERC datasets: IEMOCAP [3], MELD [27], and CMU-MOSEI [1].

The paper is structured as follows: Section 2 presents related work while we describe the proposed framework Ada2I in Section 3. Experiment settings are explained in Section 4, and Section 5 covers the experimental evaluation and results. Finally, Section 6 summarizes findings.

2 RELATED WORK

2.1 Multimodal Emotion Recognition in Conversation

Multi-modal Emotion Recognition (ERC) has emerged as a focal point within the affective computing community, garnering significant attention in recent years. The integration of multimodal data provides a multidimensional perspective, enabling a more nuanced understanding of emotions. Moreover, researchers have increasingly turned to multimodal fusion techniques, combining text, audio, and visual cues to enhance multimodal ERC performance [9, 10, 16, 18, 24]. ICON [9] employs two Gated Recurrent Units (GRUs) to capture speaker information, supplemented by global GRUs to track changes in emotional states throughout conversations. Similarly, MMGCN [37] utilizes Graph Convolutional Networks (GCNs) to capture contextual information, effectively leveraging multimodal dependencies and speaker information. On the other hand, Multilogue-Net [30] introduces a solution utilizing a context-aware RNN and employing pairwise attention as a fusion mechanism. TBJE [4], adopts a transformer-based architecture with modular co-attention to jointly encode multiple modalities. Additionally, COGMEN [16] is a multimodal context-based graph neural network that integrates both local (speaker information) and global (contextual information) aspects of conversation. Moreover, CORECT [24] employs relational temporal Graph Neural Networks (GNNs) with cross-modality interaction support, effectively capturing conversation-level interactions and utterance-level temporal relations. GraphMFT [18] utilizes multiple enhanced graph attention networks to capture intra-modal contextual information and inter-modal complementary information. More recently, DF-ERC [17] emphasizes both feature disentanglement and fusion while taking into account both multimodalities and conversational contexts. Moreover, AdaGN [32] employs the Gumbel Softmax trick to adaptively select nodes and edges, enhancing intra- and cross-modal interactions. While these methods primarily focus on designing model structures, they overlook the challenges posed by modality imbalance during multimodal learning.

2.2 Imbalanced multimodal learning

Despite the suggestion by [13] that integrating multiple modalities could enhance the accuracy of latent space estimations, thereby improving the efficacy of multimodal models, our investigation within the multimodal ERC task reveals a phenomenon contradicting this notion. The problem of modality imbalance persists as a significant challenge in multimodal learning frameworks, particularly in tasks like multimodal Emotion Recognition in Conversations (ERC). Conventional methods often prioritize one modality over others, assuming that certain types of sensory data are more relevant for a given task. For example, textual cues may receive greater emphasis, while visual or audio cues alone might be prioritized [16, 24, 37].

Current methodologies for addressing imbalanced multimodal learning primarily focus on tasks such as audio-visual learning with a focus on optimizing pairwise modality learning [6, 25, 39], sentiment analysis, and sarcasm detection [44]. However, these approaches often have task-specific limitations and framework restrictions, limiting their broader applicability. For instance, Wang et al. [35] identified that different modalities overfit and generalize at different rates, leading to suboptimal solutions when jointly trained using a unified optimization strategy. Peng et al. [25] proposed OGM-ME method where the better-performing modality dominates the gradient update, suppressing the learning process of other modalities. MMCosine [39] employs normalization techniques on features and weights to promote balanced and improved fine-grained learning across multiple modalities. Notably, there is a lack of specific approaches tailored for multimodal Emotion Recognition in Conversations (ERC) apart from the work by Wang et al. [36]. Recently, Wang et al. [36] observed a phenomenon referred to as “diminishing modal marginal utility” and proposed fine-grained adaptive gradient modulation, which was applied to ERC, while I²MCL considers both data difficulty and modality balance for multimodal learning based on curriculum learning for affective computing, though not specifically for emotion recognition. To comprehensively address the challenge of modality imbalance in multimodal ERC, we propose an end-to-end model that ensures balance among text, audio, and visual modalities during training.

3 METHODOLOGY

In the context of a conversation C with N utterances $\{u_1, u_2, \dots, u_N\}$, the task of Emotion Recognition in Conversations (ERC) is to predict the emotion label for each utterance in the conversation from a predefined emotion category set \mathcal{E} . Each utterance is associated with M modalities, i.e. textual (t), audio (a), and visual (v) modalities, represented as:

$$u_i = \{u_i^t, u_i^a, u_i^v\}, i \in \{1, \dots, N\} \quad (1)$$

where $u_i \in \mathbb{R}^{M \times d}$, d signifies the dimension of modal features. For each modality m , we derive multimodal features $\{\mathbf{X}^m\}_{m \in \{t,a,v\}} \in \mathbb{R}^{d_m \times N}$ for the conversation C . Here, $\{d_m\}_{m \in \{t,a,v\}}$ is the feature dimension of each modality.

In this section, we outline our proposed model Ada2I, including its main sub-modules: Modality Encoder, Adaptive Feature Weighting and Adaptive Modality Weighting. We also refine the disparity ratio metric as part of our Training Optimization Strategy. The architecture of our model is shown in Figure 2.

3.1 Modality Encoder

Given a conversation C , a **Transformer** [33] network is utilized as the encoder to generate a unimodal representation $\mathbf{Z}^m \in \mathbb{R}^{N \times d_m}$ respecting to the modality m as:

$$\mathbf{Z}^m = \phi(\theta^{(m)}, \mathbf{X}^m), m \in \{t, a, v\} \quad (2)$$

where the function $\phi(\theta^{(m)})$ is the Transformer network with learnable parameter $\theta^{(m)}$.

3.2 Adaptive Feature Weighting (AFW)

3.2.1 Tensor-based Multimodal Interaction Representation. Motivated by the tensor-ring decomposition method introduced by [42], we extend the traditional attention mechanism by replacing the query (**Q**) and key (**K**) representations with tensor-ring decomposition-based counterparts. This modification results in query tensor-ring representation \mathcal{G}_Q and key tensor-ring representation \mathcal{G}_K , which facilitate the acquisition of more compact modality representations. Additionally, inspired by [31], we integrate a tensor-based multi-way interaction transformer architecture into our model. This enhancement allows the model to capture multi-way interactions among modalities, thereby enhancing its capability to discern intricate multimodal relationships.

We employ a tensor-ring-based generation function to retrieve the multi-interaction multimodal query tensor Q and key tensor \mathcal{K} from the input modality presentations \mathbf{Z}^m . Specifically, we compute Q and \mathcal{K} as follows:

$$\begin{cases} Q = \text{Tr}\{\mathcal{G}_Q^{(t)}, \mathcal{G}_Q^{(a)}, \mathcal{G}_Q^{(v)}\} \in \mathbb{R}^{d_t \times d_a \times d_v} \\ \mathcal{K} = \text{Tr}\{\mathcal{G}_K^{(t)}, \mathcal{G}_K^{(a)}, \mathcal{G}_K^{(v)}\} \in \mathbb{R}^{d_t \times d_a \times d_v} \end{cases} \quad (3)$$

Here, $\text{Tr}\{\cdot\}$ represents the tensor-ring decomposition function, which naturally provides the low-rank core tensor representations \mathcal{G}_Q^m and \mathcal{G}_K^m for each modality.

To perform multimodal attention in the tensor space, we need to compute the attention coefficient matrix, Θ , from the tensorized input. To achieve this, we can first compute the Tensor-ring Key representation and Tensor-ring Query representation of input data, $\mathcal{G}_Q^m \in \mathbb{R}^{d_m \times r_s \times r_w}$ and $\mathcal{G}_K^m \in \mathbb{R}^{d_m \times r_s \times r_w}$, where $m \in \{t, a, v\}$, the index $s, w \in \{1, 2, 3\}$, and $s \neq w$. The attention coefficient matrix Θ of modality m is formulated as follows:

$$\Theta^m = \text{softmax}\left(\frac{1}{\sqrt{d_k}} \mathcal{G}_Q^m \odot \mathcal{G}_K^m\right) \quad (4)$$

where \odot denotes the element-wise product, $\sqrt{d_k}$ is a scaling factor.

More specifically, the modality m core tensor \mathcal{G}_K and \mathcal{G}_Q are expressed as follows:

$$\begin{cases} \mathcal{G}_Q^m = \text{reshape}((\mathbf{Z}^m W_{Q_m}^{(1)}) \otimes_1 (\mathbf{Z}^m W_{Q_m}^{(2)})) \\ \mathcal{G}_K^m = \text{reshape}((\mathbf{Z}^m W_{K_m}^{(1)}) \otimes_1 (\mathbf{Z}^m W_{K_m}^{(2)})) \end{cases} \quad (5)$$

where $m \in \{t, a, v\}$, $W_{Q_m}^{(1)} \in \mathbb{R}^{d_m \times r_s}$, $W_{Q_m}^{(2)} \in \mathbb{R}^{d_m \times r_w}$, $W_{K_m}^{(1)} \in \mathbb{R}^{d_m \times r_s}$, $W_{K_m}^{(2)} \in \mathbb{R}^{d_m \times r_w}$ are the linear transformation matrix; \otimes_1 denotes the mode-1 Khatri-Rao product.

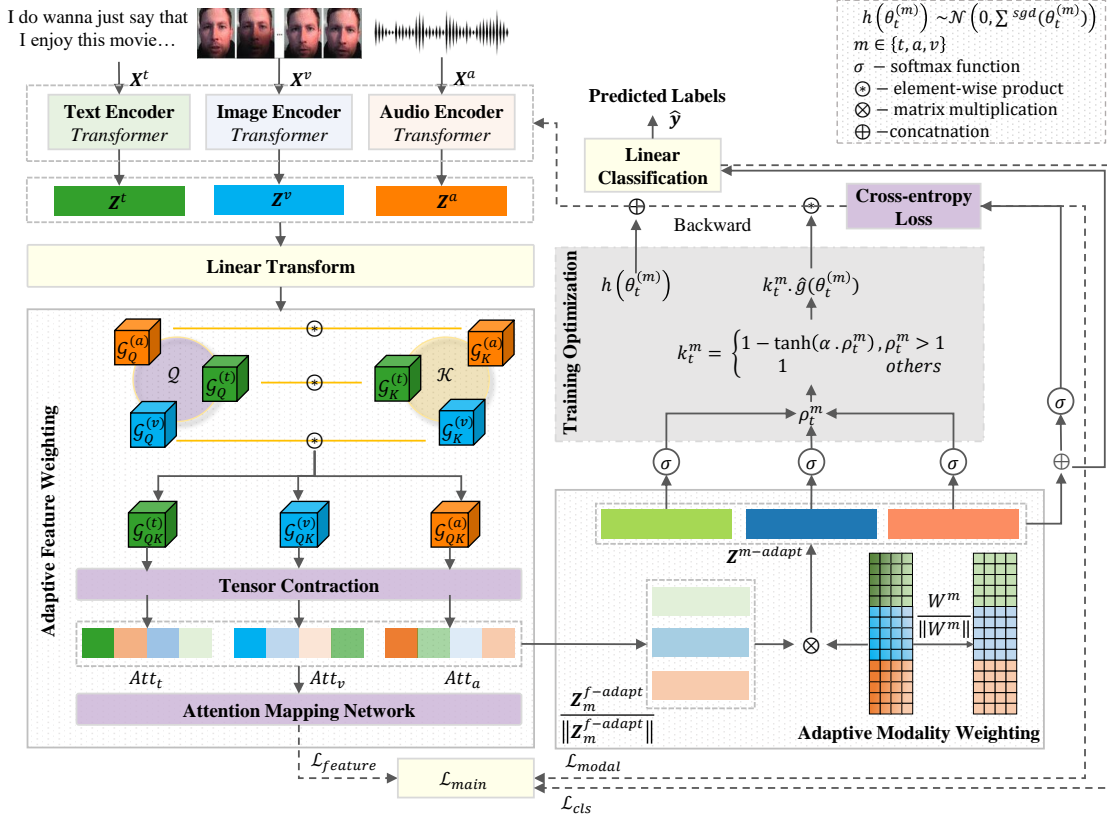


Figure 2: Illustration of Ada2I framework during the training phase

3.2.2 *Adaptive Feature Weighting (AFW)*. This module addresses the varying impact of each modality on inter-modality and intra-modality interactions using attention mechanism. First, we calculate the attention pooling matrices $\mathbf{A}^{(m)} \in \mathbb{R}^{r_s \times r_w}$ by averaging $\Theta^{(m)}$ across the modality dimension d_m , $m \in \{t, a, v\}$. Inspired by MMT [31], the *feature-aware* attention matrix $Att_m \in \mathbb{R}^{N \times d_m}$ for a given modality m is computed as follows:

$$Att_m = \text{Linear} \left(\Theta^m \times_3^1 \mathbf{A}^{(t)} \times_3^1 \mathbf{A}^{(a)} \times_3^1 \mathbf{A}^{(v)} \right) \quad (6)$$

where \times_3^1 is the *mode-1* tensor contraction. The *feature-aware* balanced representation $\mathbf{Z}_m^{f-adapt} \in \mathbb{R}^{N \times d_m}$ of the conversation C for a given modality m is computed as:

$$\mathbf{Z}_m^{f-adapt} = Att_m \mathbf{Z}^m + \beta \mathbf{Z}^m \quad (7)$$

where $\beta \in [0, 1]$ is a balancing parameter to regulate the contribution of the original unimodal feature vector \mathbf{Z}^m .

3.3 Adaptive Modality Weighting (AMW)

Our key focus is to achieve balanced contributions from each modality during the training. Similar to [39], we observe the imbalance problem in multimodal ERC through experiments analyzing the modality-wise weight in norm of each label during training. Apparently, the dominant unimodal encoder, e.g., text, tends to have its

weight in norm increase much faster than the weaker modalities, i.e., audio and visual, leading to divergent unimodal logit scores and distorting the joint fusion representation. Inspired by [34, 43], we propose to incorporate modality-wise L2 normalization to properly weight features, mitigating imbalances arising from differing data distributions and noise levels across modalities. This dynamic adjustment prevents any single modality from dominating the fusion process, thus enhancing overall performance.

Therefore, the modality-level balanced representation $\mathbf{Z}^{m-adapt}$ of the given conversation is calculated as follows:

$$\mathbf{Z}^{m-adapt} = \sum_m^{\{t, a, v\}} \frac{W^m \mathbf{Z}_m^{f-adapt}}{\|W^m\| \|\mathbf{Z}_m^{f-adapt}\|} + b \quad (8)$$

where $W^m \in \mathbb{R}^{d_m \times |\mathcal{E}|}$ symbolizes the output matrix of the model pertaining to modality m , and \mathcal{E} is the set of emotion classes.

For **emotion recognition**, we feed $\mathbf{Z}^{m-adapt}$ into the multilayer perceptron (MLP) with ReLU activation function to compute the output $\hat{y}_i \in \mathbb{R}^{N \times |\mathcal{E}|}$.

$$\hat{y}_i = \text{MLP}(\mathbf{Z}^{m-adapt}) \quad (9)$$

The output \hat{y}_i is utilized to predict emotion labels.

3.4 Learning

First, we investigate the standard cross-entropy loss for this downstream task, i.e., multimodal ERC as:

$$\mathcal{L}_{cls} = -\frac{1}{B} \sum_i^B y_i \log \hat{y}_i \quad (10)$$

where B is the batch size.

Second, in order to align between the original unimodal representation of modality m and its respective *feature-aware* attention weights as Eq (6), we employ Attention Mapping Network as follows:

$$\hat{Att}_m = \Phi_m(\mathbf{Z}_m, \psi^{(m)}), m \in \{t, a, v\} \quad (11)$$

where $\Phi_m(\cdot)$ is a feed-forward neural network with the parameter $\psi^{(m)}$, $\hat{Att}_m \in \mathbb{R}^{N \times d_m}$ is the *feature-aware* self-attention weights of the modality m . To enhance feature-level balance across all modalities, we introduce a L1-norm loss $\mathcal{L}_{feature}$ as:

$$\mathcal{L}_{feature} = \frac{1}{B} \sum_i^B \left(\sum_m^{\{t,a,v\}} |Att_m^i - \hat{Att}_m^i| \right) \quad (12)$$

Additionally, we also consider the modality-level balance loss \mathcal{L}_{modal} , which is computed as:

$$\mathcal{L}_{modal} = -\frac{1}{B} \sum_i^B \log \frac{e^{\mathbf{Z}_i^{m-adapt}}}{\sum_{j=1}^{|\mathcal{E}|} e^{\mathbf{Z}_j^{m-adapt}}} \quad (13)$$

where $\mathbf{Z}_j^{m-adapt}$ represents the output of the j -th class for the i -th sample. Finally, we combine the all loss functions into a joint objective function, which is used to optimize all trainable parameters in an end-to-end manner.

$$\mathcal{L}_{main} = \mathcal{L}_{modal} + \mathcal{L}_{feature} + \mathcal{L}_{cls} \quad (14)$$

Training Optimization Strategy

Recent studies have brought attention to the challenge of handling imbalanced optimization in joint learning models, particularly when dealing with multiple modalities. Peng et al. [25] introduce the OGM-GE method to address optimization imbalances encountered during the simultaneous training of dual-modal systems, i.e., visual and audio. However, directly applying the OGM-GE method to our framework is not practical as it only deals with two modalities. In contrast, our framework caters to more than two modalities across different domains, specifically tailored for the multimodal ERC task. Therefore, learnable parameter of encoder layer is optimized during training process as the following strategy:

$$\theta_{t+1}^{(m)} = \theta_t^{(m)} - \eta \cdot \hat{g}(\theta_t^{(m)}) \quad (15)$$

where $\hat{g}(\theta^m t) = \frac{1}{o} \sum x \in B_t \nabla_{\theta_t^m} \ell(x, \theta_t^{(i)})$ represents an unbiased estimation of the full gradient $\nabla_{\theta_t^m} \ell(x, \theta_t^{(i)})$ using a random mini-batch B_t chosen at the t -th step with size o . The term $\nabla_{\theta_t^m} \ell(x, \theta_t^{(i)})$ denotes the gradient with respect to B_t .

We adjust the balance of modalities through gradient parameter adjustments. For each output at step t , we compute the discrepancy ratio for each modality using the softmax of the cosine similarity

Algorithm 1 Ada2I Training Procedure

Input: The training set $\mathcal{D} = \{(x_t^t, x_t^a, x_t^v), y_i\}_{i=1}^N, m \in \{t, a, v\}$

Output: Prediction emotion label \hat{y}

for each training epoch **do**

for minibatch $\mathcal{B} = \{(x_t^t, x_t^a, x_t^v), y_i\}_{i=1}^N$ sampled from \mathcal{D} **do**

 #Refer to Subsection 3.1

 Encode unimodal feature \mathbf{X}^m to \mathbf{Z}^m as Eq (2)

 #Refer to Subsection 3.2

 Multimodal feature representation as Eq (3)

 Calculate coefficient matrix Θ^m as Eq (4)

 Calculate modality-aware attention Att_m as Eq (6)

 Compute fused feature $\mathbf{Z}_m^{f-adapt}$ with β using Eq (7)

 #Refer to Subsection 3.3

 Compute logit output $\mathbf{Z}^{m-adapt}$ with modality-wise L2 normalization as Eq (8)

 Produce prediction of multimodal data \hat{y}_i as Eq (9)

 #Refer to Subsection 3.4

 Use cross-entropy loss to calculate \mathcal{L}_{cls} as Eq (10)

 Use L_1 to calculate $\mathcal{L}_{feature}$ as Eq (12)

 Use cross-entropy to calculate \mathcal{L}_{modal} as Eq (13)

 Add $\mathcal{L}_{feature}$, \mathcal{L}_{modal} and \mathcal{L}_{cls} to compute \mathcal{L}_{main} as

Eq (14)

 Compute discrepancy ratio $\rho_t^m = \frac{s_t^m}{\min_{m \in \{t,a,v\}} (s_t^j)}$

 Compute modulation coefficient k_t^m

 Update using $\theta_{t+1}^{(i)} = \theta_t^{(i)} - \eta \cdot \hat{g}(\theta_t^{(i)}) \cdot k_t^i + \eta \cdot h(\theta_t^{(i)})$

end for

end for

between the output weights and the corresponding feature vectors:

$$s_t^m = \sum_{j=1}^L \sum_{k=1}^{\mathcal{E}} \mathbb{I}_{k=y_j} \text{softmax}(\cos(\langle W_k^m, \mathbf{Z}_k^m \rangle + \frac{b_k}{M})) jk \quad (16)$$

where $\mathbb{I}_{k=y_j}$ equals 1 if $k = y_j$ and 0 otherwise, and $\text{softmax}(\cdot)$ estimates the unimodal performance of the multimodal model, M denotes the count of modalities. Specifically, for the Multimodal ERC task under consideration, we delineate three modalities: text (t), audio (a), and visual (v). The discrepancy ratio is calculated as:

$$\rho_t^m = \frac{s_t^m}{\min_{m \in \{t,a,v\}} (s_t^j)} \quad (17)$$

The learnable parameters are updated according to:

$$\theta_{t+1}^{(m)} = \theta_t^{(m)} - \eta \cdot \hat{g}(\theta_t^{(m)}) \cdot k_t^m \quad (18)$$

where the modulation coefficient k_t^m is determined by $1 - \tanh(\alpha \cdot \rho_t^m)$ if $\rho_t^m > 1$, and 1 otherwise. Here, α is a hyperparameter controlling the degree of modulation. Additionally, to enhance the adaptability of the modulation process, Gaussian noise $h(\theta_t^{(i)})$ sampled from a distribution $\mathcal{N}(0, \Sigma^{sgd}(\theta_t^{(i)}))$ is introduced after parameter updates:

$$\theta_{t+1}^{(i)} = \theta_t^{(i)} - \eta \cdot \hat{g}(\theta_t^{(i)}) \cdot k_t^i + \eta \cdot h(\theta_t^{(i)}) \quad (19)$$

The training process of Ada2I is illustrated in Algorithm 1.

4 EXPERIMENTAL SETUP

4.1 Datasets

Datasets: We consider three benchmark datasets for multimodal ERC namely: IEMOCAP [3], MELD [27], and CMU-MOSEI [1]. The dataset statistics are illustrated in Table 1.

IEMOCAP. This dataset offers 12 hours of video recordings capturing dyadic conversations involving 10 speakers. Each video contains a single dialogue, segmented into utterances, resulting in a total of 7,433 utterances across 151 dialogues. Notably, each utterance is annotated with one of six emotion labels: happy, sad, neutral, angry, excited, or frustrated.

MELD. This dataset is based on the TV series Friends, includes 13,709 video clips featuring multi-party conversations, each labeled with one of Ekman’s six universal emotions: joy, sadness, fear, anger, surprise, and disgust.

CMU-MOSEI. This dataset is a prominent resource for sentiment and emotion analysis, comprises 3,228 YouTube videos divided into 23,453 segments, featuring contributions from 1,000 speakers covering 250 topics. It includes six emotion categories: happy, sad, angry, scared, disgusted, and surprised, with sentiment intensity ranging from -3 to 3.

Table 1: Data Statistics

Datasets	Dialogues			Utterances		
	train	valid	test	train	valid	test
IEMOCAP	120	31		5,810	1,623	
MELD	1,039	114	280	9,989	1,109	2,610
CMU-MOSEI	2,248	300	676	16,326	1,871	4,659

4.2 Baselines and Evaluation Metrics

Baselines: Ada2I is compared against several state-of-the-art (SOTA) baseline approaches for evaluating performance in multimodal ERC, particularly addressing modality imbalance problems. For the IEMOCAP and MELD datasets, we consider baseline models such as DialogueRNN [23], DialogueGCN [8], MMGCN [37], BiDIN [40], and MM-DFN [11]. We report the best results obtained from [36], which enhanced these models to address modality imbalance. Additionally, we consider other SOTA models for multimodal ERC that do not explicitly address modality imbalance, including COGMEN [16], CORECT [24], GraphMFT [18], DF-ERC [17], and AdaIGN [32]. For the CMU-MOSEI dataset, we evaluated various baseline models for sentiment classification tasks, which include both 2-class sentiment, featuring only positive and negative sentiment, and 7-class sentiment, ranging from highly negative (-3) to highly positive (+3). These baseline models include Multilouge-Net [30], TBJE [4], COGMEN [16], CORECT [24], OGM-GE [25], and I²MCL [44]. Notably, OGM-GE and I²MCL specifically address the issue of imbalanced modalities in multimodal ERC, whereas the other baseline models do not.

Evaluation Metrics: Similar to prior studies [23, 36, 37], we evaluate the effectiveness of emotion recognition using Accuracy (Acc) and Weighted F1 Score (WF1) as our primary evaluation metrics.

4.3 Experimental Settings

We derive multimodal features for each utterance from acoustic, lexical, and visual modalities using a combination of models and pre-trained models, as outlined in Table 2. We employ PyTorch¹ for training our architecture and Comet² for logging all experiments, leveraging its Bayesian optimizer for hyperparameter tuning. Additional parameters can be found in Table 2.

Table 2: Hyper-parameter settings

Parameter/Module	IEMOCAP	MELD	CMU-MOSEI
Text Feature Extraction	sBERT ³		
Audio Feature Extraction	Wave2vec-Large [28], OpenSmile [5]		
Visual Feature Extraction	MTCNN [41], MA-Net ⁴ , DenseNet [12]		
Text embedding dim. d_t	768	768	768
Audio embedding dim. d_a	512	300	512
Visual embedding dim. d_v	1024	342	1024
hidden dim	300	200	500
tensor rank	11	6	10
α	0.037	0.4	0.4
β	0.01	0.55	0.2
learning rate	1.7e-4	1.2e-4	1.9e-4
batch size	10	10	32
epoch	50	50	30

5 RESULTS AND DISCUSSION

5.1 Performance Comparison against Baselines

IEMOCAP and MELD dataset: As depicted in Table 3, our model Ada2I performs better than the previous SOTA baselines in the context of balanced modality consideration on all modality combinations on both datasets. Indeed, in the AV modality pair on the MELD dataset, traditionally deemed the weakest, we observe a substantial performance boost in Multimodal ERC. Specifically, there is a noteworthy enhancement of 10.77% on WF1 and 6.98% on Accuracy compared to the previous SOTA model. This progress effectively reduces the performance discrepancy compared to modality pairs where text plays a dominant role.

We also compare Ada2I with SOTA baseline models for multimodal ERC, particularly those focusing solely on multimodal fusion and architectural design without addressing modality imbalance. Figure 3b demonstrates that our proposed Ada2I significantly reduces the performance gap in WF1 between learning from all three modalities simultaneously (T+A+V) and pair-wise modality combinations on the MELD dataset. Most notably, with the weaker

¹<https://pytorch.org/>

²<https://comet.ml>


³<https://www.sbert.net/>

⁴<https://github.com/zengqunzhao/MA-Net>

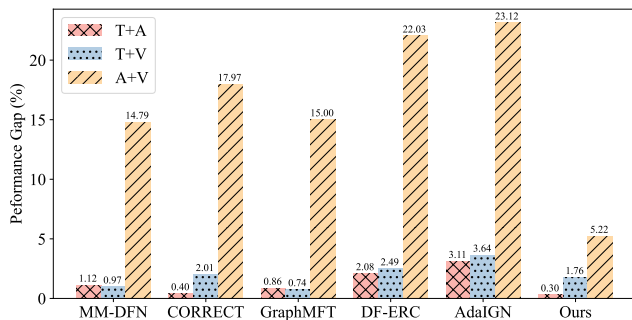
Table 3: Comparison results in the multimodal setting of Ada2I with the modal-balanced baseline model enhanced by FAGM [36]. The best performance is marked in bold, and the second best performance is indicated by underlining.

Methods	IEMOCAP								MELD							
	TAV		TA		TV		AV		TAV		TA		TV		AV	
	W-F1	Acc	W-F1	Acc	W-F1	Acc	W-F1	Acc	W-F1	Acc	W-F1	Acc	W-F1	Acc	W-F1	Acc
DialogueRNN†	61.31	61.61	61.90	61.98	60.19	59.95	48.31	50.71	56.42	58.05	56.46	58.01	55.67	57.39	40.46	45.39
DialogueGCN†	62.76	63.22	<u>64.36</u>	<u>64.39</u>	<u>61.25</u>	<u>62.23</u>	49.20	49.85	54.61	58.96	54.80	57.28	55.26	57.10	10.02	44.44
BiDDIN†	58.81	58.84	58.88	58.16	59.04	58.96	46.36	46.77	57.47	59.18	56.56	58.05	56.93	58.10	<u>44.39</u>	48.62
MM-DFN†	<u>64.92</u>	<u>64.57</u>	63.91	64.20	61.02	60.60	<u>54.48</u>	<u>55.03</u>	55.75	60.8	57.10	60.00	<u>57.73</u>	<u>60.65</u>	42.05	<u>48.66</u>
MMGCN†	64.53	64.51	63.25	63.40	61.02	61.06	54.14	54.90	<u>58.48</u>	<u>61.15</u>	<u>57.59</u>	<u>60.69</u>	57.14	59.46	43.49	48.43
Ada2I	68.97	68.76	66.91	67.28	65.48	65.43	55.16	55.64	60.38	63.03	60.08	62.64	58.62	61.95	55.16	55.64
Δ	$\uparrow 4.05$	$\uparrow 4.19$	$\uparrow 2.55$	$\uparrow 2.89$	$\uparrow 4.23$	$\uparrow 3.20$	$\uparrow 0.68$	$\uparrow 0.61$	$\uparrow 1.90$	$\uparrow 1.88$	$\uparrow 2.49$	$\uparrow 1.95$	$\uparrow 0.89$	$\uparrow 1.30$	$\uparrow 10.77$	$\uparrow 6.98$

modality pair (audio+visual) consistently lagging behind in performance compared to the full modality combination (i.e., with AdaIGN, this gap is 23.12%), Ada2I boosts the model and shortens the gap to only 5.22%. Similarly, with the text+audio (T+A) and text+visual (T+V) pairs, this gap is also substantially reduced, indicating that the model has learned in a more balanced manner, leveraging additional useful information from non-dominant modalities. The significant improvement is similarly observed on the IEMOCAP dataset in Figure 3a.



(a) on IEMOCAP dataset



(b) on MELD dataset

Figure 3: Performance gap visualizations between the multimodal setting (T+A+V) and pair-wise modality combinations are evaluated using the WF1 metric across the IEMOCAP and MELD datasets.**Table 4: Results on the CMU-MOSEI dataset. The best performance is highlighted in bold. For cells displaying "-", the results were not provided in the paper. † denotes results obtained from running the provided code in original paper.**

Methods	2-class				7-class			
	TAV	TA	TV	AV	TAV	TA	TV	AV
Multilouge-Net [30]	82.10	80.18	80.06	75.16	44.83	-	-	-
TBJE [4]	81.50	82.40	-	-	44.40	<u>45.50</u>	-	-
COGMEN† [16]	82.95	<u>85.00</u>	82.99	65.95	43.90	44.31	42.68	24.27
CORRECT† [24]	83.98	84.28	82.83	68.89	<u>46.31</u>	44.89	43.76	24.55
I ² MCL [44]	81.05	-	-	-	-	-	-	-
OGM-GE† [25]	<u>84.58</u>	84.03	<u>83.67</u>	71.53	45.43	43.68	<u>44.44</u>	<u>31.53</u>
Ada2I	85.25	85.08	85.21	<u>74.93</u>	47.71	47.35	47.37	34.64
Δ	$\uparrow 0.67$	$\uparrow 0.08$	$\uparrow 1.54$	$\downarrow 0.23$	$\uparrow 2.28$	$\uparrow 1.85$	$\uparrow 2.93$	$\uparrow 3.11$

CMU-MOSEI dataset: Table 4 shows that Ada2I outperforms all baseline models. Specifically, when compared to OGM-GE and I²MCL, two models proposed for addressing modality imbalance during training, Ada2I demonstrates superior performance across all modality combinations. When compared to other baseline models that do not consider modality balancing, Ada2I also demonstrates significant balancing capabilities, reducing the performance gap between modality pairs. For instance, in the CORRECT model, the gap between T+A+V and A+V is 15.09% for 2-class sentiment, and this figure increases to 21.76% for 7-class sentiment. However, with Ada2I, these gaps are significantly reduced to 10.32% and 13.07%, respectively, underscoring the effectiveness of Ada2I in addressing modality imbalances.

5.2 Ablation Study

5.2.1 Balancing Interpretation. We conduct ablation studies with the two main modules of the model, AMW and AFW, to assess their impact on the Ada2I model. Additionally, through the Discrepancy Ratio, we interpret the model's balancing by observing its changes. A smaller Discrepancy Ratio indicates a more balanced optimization process. Figure 4 shows that the discrepancy ratios ρ^t , ρ^v , and ρ^a significantly decrease when both AMW and AFW are combined within Ada2I, with all ratios approaching approximately 1 on the IEMOCAP dataset. In contrast, when one of the modules is ablated,

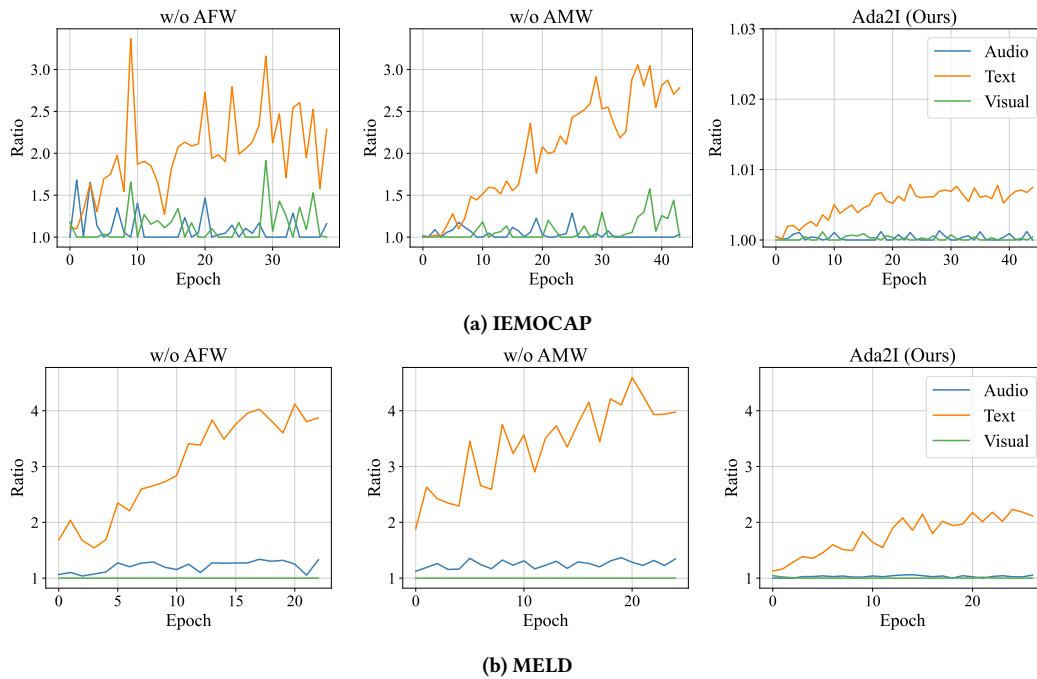


Figure 4: The change of the discrepancy ratio ρ^t, ρ^a, ρ^v on the IEMOCAP and MELD datasets during training, along with various ablation tests including without AMW and without AFW, are compared to the Ada2I model.

the ratios for audio (ρ^a) and visual (ρ^v) are approximately 1.5, while for text, it increases to around 3. Similarly, on the MELD dataset, our proposed model Ada2I has reduced this discrepancy ratio of text from over 4 (w/o AFW) to approximately half, reaching around 2, while for audio and visual, it brings them close to the 1 mark. In summary, the combined design of both modules AMW and AFW enhances balanced learning across modalities during training, highlighting the significance and inseparability of feature-level and modality-level balancing.

5.2.2 Effect of Weight Normalization. As mentioned earlier, the unimodal weights also directly influence the encoder updating process. The imbalanced weight components induce gradients and subsequently lead to the inconsistent convergence of unimodalities. Here, we provide a clearer visualization of these unimodal weights before imbalance processing (Only Encoder) and in the Ada2I model in Figure 5 for the IEMOCAP dataset. It is evident that with Only Encoder, the text encoder (dominant modality) weight in norm grows much faster than audio and visual. After balancing, our model exhibits a more balanced optimization process.

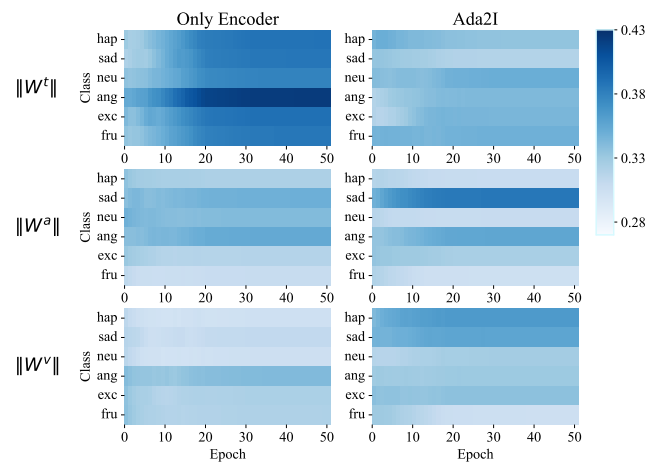


Figure 5: The observation of modality-wise weights for each label normalized for the IEMOCAP dataset

6 CONCLUSION

In this work, we present **Ada2I**, a framework designed to address modality imbalances and optimize learning in multimodal ERC. We identify and analyze existing issues in current ERC models that overlook the imbalance problem. From there, we propose a solution comprising integral modules: Adaptive Feature Weighting (AFW) and Adaptive Modality Weighting (AMW). The former enhances intra-modal representations for feature-level balancing, while the latter optimizes inter-modal learning weights with the balancing at modality level. Furthermore, we introduce a refined disparity ratio to optimize training, offering a straightforward yet effective measure to evaluate the model's overall discrepancy when handling multiple modalities simultaneously. Extensive experiments on the IEMOCAP, MELD, and CMU-MOSEI datasets validate its effectiveness, showcasing SOTA performance. In the future, we anticipate enhancing the efficiency of the framework and maximizing the utilization of emotional cues.

REFERENCES

- [1] AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 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)*, Iryna Gurevych and Yusuke Miyao (Eds.). Association for Computational Linguistics, Melbourne, Australia, 2236–2246. <https://doi.org/10.18653/v1/P18-1208>
- [2] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence* 41, 2 (2018), 423–443.
- [3] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. IEMOCAP: Interactive emotional dyadic motion capture database. *Language resources and evaluation* 42, 4 (2008), 335–359.
- [4] Jean-Benoit Delbrouck, Noé Tits, Mathilde Brousmiche, and Stéphane Dupont. 2020. A Transformer-based joint-encoding for Emotion Recognition and Sentiment Analysis. In *Second Grand-Challenge and Workshop on Multimodal Language (Challenge-HML)*, Amir Zadeh, Louis-Philippe Morency, Paul Pu Liang, and Soujanya Poria (Eds.). Association for Computational Linguistics, Seattle, USA, 1–7. <https://doi.org/10.18653/v1/2020.challengehml-1.1>
- [5] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*. 1459–1462.
- [6] Yunfeng Fan, Wenchao Xu, Haozhao Wang, Junxiao Wang, and Song Guo. 2023. PMR: Prototypical Modal Rebalance for Multimodal Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20029–20038.
- [7] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2019. Slow-fast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*. 6202–6211.
- [8] Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Alexander Gelbukh. 2019. DialogueGCN: A Graph Convolutional Neural Network for Emotion Recognition in Conversation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 154–164. <https://doi.org/10.18653/v1/D19-1015>
- [9] Devamanyu Hazarika, Soujanya Poria, Rada Mihalcea, Erik Cambria, and Roger Zimmermann. 2018. ICON: Interactive Conversational Memory Network for Multimodal Emotion Detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, Brussels, Belgium, 2594–2604. <https://doi.org/10.18653/v1/D18-1280>
- [10] Devamanyu Hazarika, Soujanya Poria, Amir Zadeh, Erik Cambria, Louis-Philippe Morency, and Roger Zimmermann. 2018. Conversational Memory Network for Emotion Recognition in Dyadic Dialogue Videos. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, Marilyn Walker, Heng Ji, and Amanda Stent (Eds.). Association for Computational Linguistics, New Orleans, Louisiana, 2122–2132. <https://doi.org/10.18653/v1/N18-1193>
- [11] Dou Hu, Xiaolong Hou, Lingwei Wei, Lianxin Jiang, and Yang Mo. 2022. MM-DFN: Multimodal dynamic fusion network for emotion recognition in conversations. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 7037–7041.
- [12] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2017. Densely Connected Convolutional Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [13] Yu Huang, Chenzhuang Du, Zihui Xue, Xuanyao Chen, Hang Zhao, and Longbo Huang. 2021. What makes multi-modal learning better than single (provably). *Advances in Neural Information Processing Systems* 34 (2021), 10944–10956.
- [14] Yu Huang, Junyang Lin, Chang Zhou, Hongxia Yang, and Longbo Huang. 2022. Modality competition: What makes joint training of multi-modal network fail in deep learning?(provably). In *International Conference on Machine Learning*. PMLR, 9226–9259.
- [15] Adrián Javaloy, Maryam Meghdadi, and Isabel Valera. 2022. Mitigating modality collapse in multimodal VAEs via impartial optimization. In *International Conference on Machine Learning*. PMLR, 9938–9964.
- [16] Abhinav Joshi, Ashwani Bhat, Ayush Jain, Atin Singh, and Ashutosh Modi. 2022. COGMEN: Contextualized GNN based multimodal emotion recognition. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 4148–4164.
- [17] Bobo Li, Hao Fei, Lizi Liao, Yu Zhao, Chong Teng, Tat-Seng Chua, Donghong Ji, and Fei Li. 2023. Revisiting disentanglement and fusion on modality and context in conversational multimodal emotion recognition. In *Proceedings of the 31st ACM International Conference on Multimedia*. 5923–5934.
- [18] Jiang Li, Xiaoping Wang, Guoqing Lv, and Zhigang Zeng. 2023. GraphMFT: A graph network based multimodal fusion technique for emotion recognition in conversation. *Neurocomputing* 550 (2023), 126427.
- [19] Zheng Lian, Bin Liu, and Jianhua Tao. 2021. CTNet: Conversational transformer network for emotion recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 29 (2021), 985–1000.
- [20] Paul Pu Liang, Yiwei Lyu, Xiang Fan, Zetian Wu, Yun Cheng, Jason Wu, Leslie Yufan Chen, Peter Wu, Michelle A Lee, Yuke Zhu, et al. 2021. MultiBench: Multiscale Benchmarks for Multimodal Representation Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- [21] Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. 2022. Foundations and Trends in Multimodal Machine Learning: Principles, Challenges, and Open Questions. *arXiv preprint arXiv:2209.03430* (2022).
- [22] Xun Lin, Shuai Wang, Rizhao Cai, Yizhong Liu, Ying Fu, Zitong Yu, Wenzhong Tang, and Alex Kot. 2024. Suppress and Rebalance: Towards Generalized Multimodal Face Anti-Spoofing. *arXiv preprint arXiv:2402.19298* (2024).
- [23] Navonil Majumder, Soujanya Poria, Devamanyu Hazarika, Rada Mihalcea, Alexander Gelbukh, and Erik Cambria. 2019. Dialoguernn: An attentive rnn for emotion detection in conversations. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 6818–6825.
- [24] Cam-Van Thi Nguyen, Tuan Mai, Son The, Dang Kieu, and Duc-Trong Le. 2023. Conversation Understanding using Relational Temporal Graph Neural Networks with Auxiliary Cross-Modality Interaction. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 15154–15167. <https://doi.org/10.18653/v1/2023.emnlp-main.937>
- [25] Xiaokang Peng, Yake Wei, Andong Deng, Dong Wang, and Di Hu. 2022. Balanced multimodal learning via on-the-fly gradient modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8238–8247.
- [26] Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh, and Louis-Philippe Morency. 2017. Context-Dependent Sentiment Analysis in User-Generated Videos. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Regina Barzilay and Min-Yen Kan (Eds.). Association for Computational Linguistics, Vancouver, Canada, 873–883. <https://doi.org/10.18653/v1/P17-1081>
- [27] Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. MELD: A Multimodal Multi-Party Dataset for Emotion Recognition in Conversations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Anna Korhonen, David Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, Florence, Italy, 527–536. <https://doi.org/10.18653/v1/P19-1050>
- [28] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. 2019. wav2vec: Unsupervised pre-training for speech recognition. *arXiv preprint arXiv:1904.05862* (2019).
- [29] Weizhou Shen, Siyue Wu, Yunyi Yang, and Xiaojun Qian. 2021. Directed Acyclic Graph Network for Conversational Emotion Recognition. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 1551–1560. <https://doi.org/10.18653/v1/2021.acl-long.123>
- [30] Aman Shenoy and Ashish Sardana. 2020. Multilogue-Net: A Context-Aware RNN for Multi-modal Emotion Detection and Sentiment Analysis in Conversation. In *Second Grand-Challenge and Workshop on Multimodal Language (Challenge-HML)*, Amir Zadeh, Louis-Philippe Morency, Paul Pu Liang, and Soujanya Poria (Eds.). Association for Computational Linguistics, Seattle, USA, 19–28. <https://doi.org/10.18653/v1/2020.challengehml-1.3>
- [31] Jiajia Tang, Kang Li, Ming Hou, Xuanyu Jin, Wanzeng Kong, Yu Ding, and Qibin Zhao. 2022. MMT: Multi-Way Multi-Modal Transformer for Multimodal Learning. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22, LD Raedt, Ed. International Joint Conferences on Artificial Intelligence Organization*, Vol. 7. 3458–3465.
- [32] Geng Tu, Tian Xie, Bin Liang, Hongpeng Wang, and Ruifeng Xu. 2024. Adaptive Graph Learning for Multimodal Conversational Emotion Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 19089–19097.
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 6000–6010.
- [34] Feng Wang, Xiang Xiang, Jian Cheng, and Alan Loddon Yuille. 2017. Normface: L2 hypersphere embedding for face verification. In *Proceedings of the 25th ACM international conference on Multimedia*. 1041–1049.
- [35] Weiyao Wang, Du Tran, and Matt Feiszli. 2020. What makes training multimodal classification networks hard?. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 12695–12705.

929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044

1045	[36]	Yunxiao Wang, Meng Liu, Zhe Li, Yupeng Hu, Xin Luo, and Liqiang Nie. 2023. Unlocking the Power of Multimodal Learning for Emotion Recognition in Conversation. In <i>Proceedings of the 31st ACM International Conference on Multimedia</i> . 5947–5955.	1103
1046			1104
1047			1105
1048	[37]	Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2019. MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video. In <i>Proceedings of the 27th ACM international conference on multimedia</i> . 1437–1445.	1106
1049			1107
1050			1108
1051	[38]	Nan Wu, Stanislaw Jastrzebski, Kyunghyun Cho, and Krzysztof J Geras. 2022. Characterizing and overcoming the greedy nature of learning in multi-modal deep neural networks. In <i>International Conference on Machine Learning</i> . PMLR, 24043–24055.	1109
1052			1110
1053			1111
1054	[39]	Ruize Xu, Ruoxuan Feng, Shi-Xiong Zhang, and Di Hu. 2023. MMCosine: Multi-Modal Cosine Loss Towards Balanced Audio-Visual Fine-Grained Learning. In <i>ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> . IEEE, 1–5.	1112
1055			1113
1056			1114
1057			1115
1058			1116
1059			1117
1060			1118
1061			1119
1062			1120
1063			1121
1064			1122
1065			1123
1066			1124
1067			1125
1068			1126
1069			1127
1070			1128
1071			1129
1072			1130
1073			1131
1074			1132
1075			1133
1076			1134
1077			1135
1078			1136
1079			1137
1080			1138
1081			1139
1082			1140
1083			1141
1084			1142
1085			1143
1086			1144
1087			1145
1088			1146
1089			1147
1090			1148
1091			1149
1092			1150
1093			1151
1094			1152
1095			1153
1096			1154
1097			1155
1098			1156
1099			1157
1100			1158
1101			1159
1102			1160
	[40]	Dong Zhang, Weisheng Zhang, Shoushan Li, Qiaoming Zhu, and Guodong Zhou. 2020. Modeling both intra-and inter-modal influence for real-time emotion detection in conversations. In <i>Proceedings of the 28th ACM International Conference on Multimedia</i> . 503–511.	
	[41]	Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Qiao Yu. 2016. Joint Face Detection and Alignment Using Multitask Cascaded Convolutional Networks. <i>IEEE Signal Processing Letters</i> 23, 10 (2016), 1499–1503.	
	[42]	Qibin Zhao, Guoxu Zhou, Shengli Xie, Liqing Zhang, and Andrzej Cichocki. 2016. Tensor ring decomposition. <i>arXiv preprint arXiv:1606.05535</i> (2016).	
	[43]	Yutong Zheng, Dipan K Pal, and Marios Savvides. 2018. Ring loss: Convex feature normalization for face recognition. In <i>Proceedings of the IEEE conference on computer vision and pattern recognition</i> . 5089–5097.	
	[44]	Yuwei Zhou, Xin Wang, Hong Chen, Xuguang Duan, and Wenwu Zhu. 2023. Intra-and Inter-Modal Curriculum for Multimodal Learning. In <i>Proceedings of the 31st ACM International Conference on Multimedia</i> . 3724–3735.	