

M2FNet: Multi-modal Fusion Network for Emotion Recognition in Conversation

Vishal Chudasama, Purbayan Kar, Ashish Gudmalwar, Nirmesh Shah, Pankaj Wasnik*, Naoyuki Onoe
 Media Analysis Group, Sony Research India, Bangalore, India

{vishal.chudasama1, purbayan.kar, ashish.gudmalwar, nirmesh.shah,
 pankaj.wasnik, naoyuki.onoe}@sony.com

Abstract

Emotion Recognition in Conversations (ERC) is crucial in developing sympathetic human-machine interaction. In conversational videos, emotion can be present in multiple modalities, i.e., audio, video, and transcript. However, due to the inherent characteristics of these modalities, multi-modal ERC has always been considered a challenging undertaking. Existing ERC research focuses mainly on using text information in a discussion, ignoring the other two modalities. We anticipate that emotion recognition accuracy can be improved by employing a multi-modal approach. Thus, in this study, we propose a Multi-modal Fusion Network (M2FNet) that extracts emotion-relevant features from visual, audio, and text modality. It employs a multi-head attention-based fusion mechanism to combine emotion-rich latent representations of the input data. We introduce a new feature extractor to extract latent features from the audio and visual modality. The proposed feature extractor is trained with a novel adaptive margin-based triplet loss function to learn emotion-relevant features from the audio and visual data. In the domain of ERC, the existing methods perform well on one benchmark dataset but not on others. Our results show that the proposed M2FNet architecture outperforms all other methods in terms of weighted average F1 score on well-known MELD and IEMOCAP datasets and sets a new state-of-the-art performance in ERC.

1. Introduction

Emotions are the unseen mental states that are linked to thoughts and feelings [23]. In the absence of physiological indications, they could only be detected by human actions such as textual utterances, visual gestures, and acoustic signals. Emotion Recognition in Conversations (ERC) seeks to recognize the human emotions in conversations depending on their textual, visual, and acoustic cues. Recently, ERC

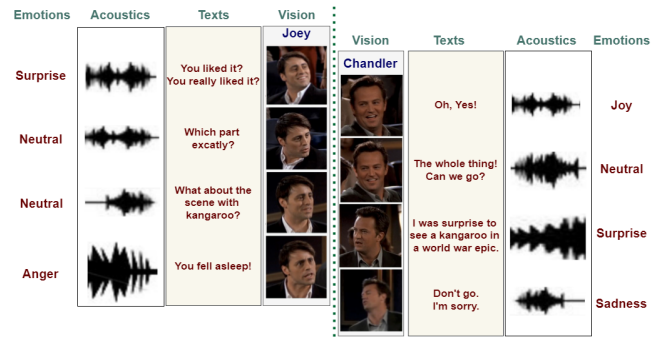


Figure 1. Multi-modal data as input

has become an essential task in multimedia content analysis and moderation. It is a prominent trait to understand the nature of the interaction between users and the content. It has applications in various tasks, namely, AI interviews, personalized dialog systems, sentiment analysis, and understanding the user's perception of the content from the platforms like YouTube, Facebook, and Twitter [23].

In literature, we can see that many state-of-the-art methods adopt text-based processing to perform robust ERC [6, 11], such methods do not take into consideration the vast amount of information present in the acoustic and visual modalities. Since the ERC data mainly consists of all three modalities, i.e., text, visual, and acoustic, we hypothesize that the robust fusion of these modalities can improve the performance and robustness of the existing systems. A sample of emotional expressions in three different modalities is presented in Figure 1 where the ERC system takes each modality as input and predicts the associated emotion.

In this paper, we propose a multi-modal fusion network (M2FNet) that takes advantage of the multi-modal nature of real-world media content by introducing a novel multi-head fusion attention layer. This layer combines features from different modalities to generate rich emotion-relevant representations by mapping the information from acoustic and visual features to the latent space of the textual features. In addition, we propose a new feature extractor model to ex-

*Pankaj Wasnik is the corresponding author.




Figure 2. Quantitative analysis on MELD and IEMOCAP datasets in terms of weighted average F1 score.

tract the deeper features from the audio and visual contents. Here, we introduce a new adaptive margin-based triplet loss function, which helps the proposed extractor to learn representations more effectively. Additionally, we propose a dual network inspired from [13] to combine the emotional content from the scene by taking into account the multiple people present in it. Furthermore, from the literature, we can see that state-of-the-art ERC methods perform well on one benchmark dataset, for example, IEMOCAP [2] while their performance degrades on more complex datasets like MELD [22]. This motivates us to propose a robust multi-modal ERC system.

In order to verify the robustness of the proposed network, one experiment is carried out to compare its performance with existing text-based and multi-modal ERC methods. This comparison is visualized in Figure 2 where results are given in terms of weighted average F1 score on MELD [2] and IEMOCAP [22] datasets. Here, it can be observed that the proposed M2FNet model obtains a higher weighted average F1 score than other models. Followings are our major contributions:

- A novel multi-modal fusion network called M2FNet is proposed for emotion recognition in conversation.
- A multi-head attention-based fusion layer is introduced, which aids the proposed system to combine latent representations of the different inputs.
- To extract deeper relevant features from audio and visual modality utterances, we introduce a new feature extractor model.
- In the feature extractor model, we propose a new adaptive margin-based triplet loss function that helps the proposed model to learn emotion-relevant features.
- To take advantage of the scene’s emotional content, we also present a weighted face model that considers multiple people present in the scene.

2. Related Works

Emotion recognition in conversations (ERC) is different from traditional emotion recognition. Rather than treating emotions as static states, ERC involves emotional dynamics of a conversation, in which the context plays a vital role. Prior works on ERC mainly which use text, and audio features are proposed in [4, 12]. In the past few years, datasets with visual, acoustic and textual cues have been made publicly available [2, 22]. On these datasets, several deep learning methods are applied to recognize emotion. These techniques can be classified based on the type of data; either they merely utilize text or use multi-modal data (i.e. text, visual and audio).

2.1. Text-based methods

With the advent of the Transformer [31], the focus on text-based methods has recently increased. Due to the vast amount of information present in text data, current methods approach ERC as a purely text-based problem. Li *et al.* [16] use BERT [5] to encode the individual sentences and then uses a dialog level network for multitask learning on auxiliary tasks to generate better latent representations of the dialog as a whole. Furthermore, Li *et al.* [14] build on this by incorporating transformer at the dialog end. Jiangnan *et al.* [15] took this one step further by using the contextual representations from a BERT and transformer dialog network by designing three types of masks and utilizing them in three independent transformer blocks. The three designed masks learn the conventional context, Intra-Speaker, and Inter-Speaker dependency.

In [6], Ghosal *et al.* incorporates different elements of commonsense such as mental states, events, and causal relations to learn interactions between interlocutors participating in a conversation. Authors in [7] and [28] use Graph Neural networks to encode inter utterance and inter speaker relationships. Kim *et al.* [11] model contextual information by simply prepending speaker names to utterances and inserting separation tokens between the utterances in a dialogue. To generate contextualized utterance representations, Wang *et al.* [32] uses LSTM-based encoders to capture self and inter-speaker dependency of interlocutors. A directed acyclic graph (DAG) based ERC was introduced by Shen *et al.* in [27] which is an attempt to combine the strengths of conventional graph-based and recurrence-based neural networks. In [38], Zhu *et al.* propose a new model in which the transformer model fuses the topical and commonsense information to predict the emotion label. Recently, Song *et al.* [33] proposed the EmotionFlow model, which encodes the user’s utterances via concatenating the context with an auxiliary question, and then, a random field is applied to capture the sequential information at the emotion level.

2.2. Multi-modal Methods

Prior literature on using previous utterances to provide context with respect to the utterance in hand has set the benchmark for dyadic conversations. In [8, 21], authors use previous utterances of both parties in a dyadic conversation and the contextual information from the same to predict the emotional state of any given utterance. Majumder *et al.* [18] build on this by separately modeling the uni-modal contextual information and then using a hierarchical tri-modal feature level fusion for obtaining a rich feature representation of the utterance. DialogueRNN [19] tracks the contextual information of each speaker and the global state as separate entities. It uses the global and speaker's emotional context to produce accurate predictions. Zhang *et al.* [36] introduced a model called ConGCN, which uses Graph Convolution networks on both audio and Text utterance features to model Speaker-Utterance and Utterance-Utterance relationships concurrently in a single network. Mao *et al.* [20] investigate the differentiated multi-modal emotional behaviors from the intra- and inter-modal perspectives. On a similar dataset, the CMU-Mosei methods like Loshchilov et al [3] and Tsai et al [30] use multi-head attention based fusion [31] for multi-modal emotion recognition.

In most of the previous work, methods do not consider distinct facial features that play a significant role in determining the emotional context of the conversation. They use the frames as a whole entity but do not extract the essential part of the frame (i.e., face). Additionally, most of these methods do not have an active fusion strategy other than simple concatenation to take advantage of the wealth of the information present in the form of visual and acoustic data.

3. Proposed Framework

This section first introduces the problem statement and then provide details of the architecture design of the proposed framework briefly.

3.1. Problem Statement

A dialog consists of k number of utterances (U) along with their respective labels (Y) arranged together with respect to time where each utterance is accompanied by its respective video clip, speech segment and text transcript. Mathematically, a dialog for k number of utterances can be formulated as follows:

$$\{U, Y\} = \{\{x^i = \langle x_t^i, x_a^i, x_v^i \rangle, y^i\} | i \in [1, k]\}, \quad (1)$$

Here, x^i denotes the i^{th} utterance made up of corresponding x_t (text), x_a (audio) and x_v (visual) component, while y_i indicates respective i^{th} utterance's emotion label. The

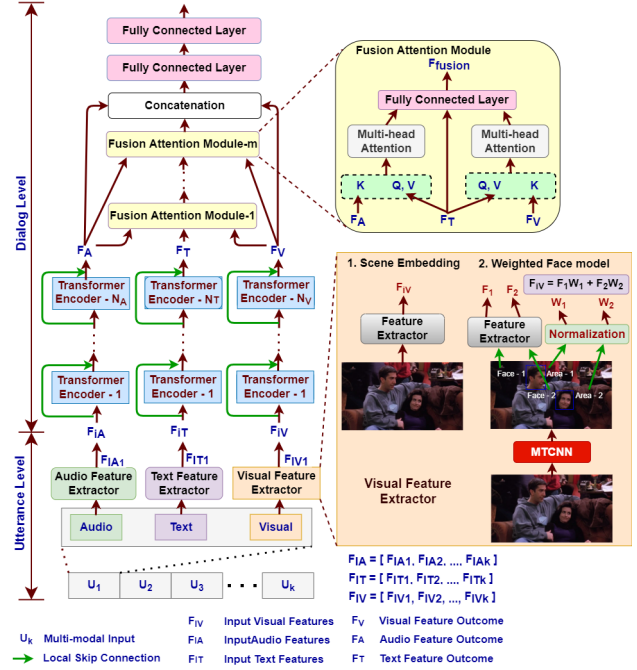


Figure 3. Network design of the proposed framework.

proposed network takes this data as input and assigns the right emotion to any given utterance.

3.2. Multi-modal Fusion Network: M2FNet

We propose a hierarchical framework called Multi-modal Fusion network (i.e., M2FNet) which is illustrated in Figure 3. The network is designed based on two levels of feature extraction:

- Utterance level feature extraction
- Dialog level feature extraction

Initially, the features are extracted by the utterance level module independently. Then, at the dialog level extraction network, the model learns to predict the right emotion for each utterance by using the contextual information from the dialog as a whole. In the subsequent subsections, we briefly discuss the design steps occurring in both utterance and dialog level feature extraction.

3.2.1 Utterance level feature extraction

From Figure 3, one can see that there are k number of utterances and each utterance is made up of x_t (text), x_a (audio) and x_v (visual) content. In this step, the features from each modality are extracted for each utterance separately before passing to the dialog level feature extraction network. Each modality's input signals are passed through

their corresponding feature extractors for generating their embeddings.

Text: In order to provide deeper inter utterance context, the text modality data (i.e., x_t) are passed through the Text Feature Extractor module. Here, we employ a modified RoBERTa model ($\phi_{M-RoBERTa}$) proposed by Kim *et al.* [11] as feature extractor. Every utterance's x_t is accompanied by its preceding and next utterance text separated by the separator token $\langle S \rangle$. The modified RoBERTa model is fine-tuned on this transcript and the respective utterance's labels. The last layer activations $\{F_{IT} : F_{1,T}, F_{2,T}, \dots, F_{k,T}\}$ obtained from the modified RoBERTa model by passing the utterance's text data can be represented by:

$$F_{IT}^i = \phi_{M-RoBERTa}(x_t^i) \mid i \in [1, k], \forall F_{IT} \in \mathbb{R}^{k \cdot D_T}. \quad (2)$$

Here, F_{IT}^i denotes i^{th} utterance's embeddings and D_T denotes the size of the embeddings of text utterance.

Audio: On the audio end, we introduce a new feature extractor model. The network design of the proposed feature extractor module is discussed briefly in Subsection 3.3. Initially, the audio contents are transformed into 2D Mel Spectrogram in RGB format and then passed through the feature extractor model. Here, the audio signal is first processed via different augmentation techniques like time warping and Additive White Gaussian Noise (AWGN) noise. Then the augmented signals are transformed into the corresponding Mel Spectrograms [24]. For computing the Mel Spectrogram, the Short Time Fourier transform (STFT) is used with the frame length of 400 samples (25 ms) and hop length of 160 samples (10ms). We also use 128 Mel filter banks to generate the Mel Spectrogram [29].

The proposed extractor takes the Mel Spectrograms (i.e., x_a) as input and generate the corresponding feature embeddings $\{F_{IA} : F_{1,A}, F_{2,A}, \dots, F_{k,A}\}$. The functionality of the proposed audio feature extractor module can be mathematically expressed as,

$$F_{IA}^i = \phi_{AFE}(x_a^i) \mid i \in [1, k], \forall F_{IA} \in \mathbb{R}^{k \cdot D_A} \quad (3)$$

where, F_{IA}^i is the i^{th} utterance's embeddings and D_A indicates the size of embeddings of audio utterance and the ϕ_{AFE} denotes the function of introduced audio feature extractor module.

Visual: In order to extract rich emotion-relevant features from the visual signal, we propose a dual network inspired from [13] that exploit not only human facial expression but also context information in a joint and boosting manner. For both tasks, we use our proposed extractor model (as discussed in Subsection 3.3) and train on the CASIA webface database [35] to extract the deeper features from the visual image. Following are the details of steps involved in the dual network:

- To encode the context information of the scene as a whole our trained feature extractor model is performed on 15 successive frames of the utterance clip. Following which the features are max pooled over the frame axis to obtain the scene embeddings of the utterance.
- To extract facial emotion-related features from the same 15 successive frames of the utterance clip, we propose a weighted Face Model (i.e., as visualized in Figure 3) which works as follows:
 - Given a frame, it is passed through a Multi-task Cascaded Convolutional Network (MTCNN) [37] to detect the faces present in the frame. This returns the bounding box of each face along with its confidence. Then each of the respective faces is passed through our trained feature extractor model to obtain emotion-relevant features of each respective face. Now, the areas of each bounding box accompanying the faces are normalized to bring their values between 0 and 1. Following this, a weighted sum is performed using the features of each face and their respective normalized areas to obtain the facial emotion feature of a frame.
 - The same process is followed for each of the 15 frames, and similarly, upon extraction, the features are max-pooled over the frame axis to obtain the facial features of the utterance.

Upon extraction of features from each network, the scene embeddings are concatenated with the facial features to obtain a more comprehensive representation of the visual data in an utterance. Finally, this visual feature extractor's output $\{F_{IV} : F_{1,V}, F_{2,V}, \dots, F_{k,V}\}$ can be formulated as:

$$F_{IV}^i = \phi_{WF}(x_v^i) \mid i \in [1, k], \forall F_{IV} \in \mathbb{R}^{k \cdot D_V} \quad (4)$$

where, F_{IV}^i is the i^{th} utterance's embeddings and ϕ_{WF} denotes the function operation of weighted face model while D_V is the size of the feature embedding of the visual utterance. Finally, these embeddings (i.e., text, acoustic and visual) are then sent to the dialog level feature extractor as an input to learn the correct prediction of emotions for each utterance.

3.2.2 Dialog level feature extraction

The design diagram at dialog level feature extraction of the proposed M2FNet model is illustrated in Figure 3. Each modality embeddings (i.e., F_{IT} , F_{IA} and F_{IV}) are passed through their corresponding network with a variable stack of transformer encoders [31] to learn the inter utterance context. The number of transformer encoders for text, audio, and visual modality is denoted by N_T , N_A , and N_V ,

respectively. We also employ a local skip connection between each encoder to prevent the model from ignoring the lower-level features. The corresponding feature maps obtained from the encoders can be mathematically expressed as:

$$\begin{aligned} F_T^i &= Tr_{N_T}(\dots(Tr_{N_2}(Tr_{N_1}(F_{IT}^i)))), \\ F_A^i &= Tr_{N_A}(\dots(Tr_{N_2}(Tr_{N_1}(F_{IA}^i)))), \\ F_V^i &= Tr_{N_V}(\dots(Tr_{N_2}(Tr_{N_1}(F_{IV}^i)))), \end{aligned} \quad (5)$$

where, $i \in [1, k]$.

Here, Tr is the operation function of the transformer encoder.

The corresponding feature maps associated with the text, visual and audio (i.e., F_T, F_V, F_A) are passed to a novel Multi-Head Attention Fusion module that helps the network in incorporating visual and acoustic information. The network architecture of the attention fusion module is also depicted in Figure 3. Here, the text features F_T are used as input to fusion module as **Query** (Q) and **Value** (V) for the multi-head attention operation, and then the visual F_V and acoustic F_A features for the dialog are used as **Key** (K) in order to modulate the attention given to each utterance at any time-step. Hence, each individual modality is now mapped to the text vector space, and the respective features are concatenated and passed to a fully connected layer which outputs a vector $\in \mathbb{R}^{k \cdot D_T}$. The output of the fusion layer is passed through the next fusion layer along with the previous F_A and F_V feature maps. Here, the m number of multi-head attention fusion layers stacked together in order to generate the final feature outcome (F_{fusion}) as demonstrated below:

$$\begin{aligned} F_{fusion_1}^i &= \Phi_1(F_A^i, F_T^i, F_V^i), \\ F_{fusion_m}^i &= \Phi_m(\dots(\Phi_2(F_A^i, F_{fusion_1}^i, F_V^i))), \end{aligned} \quad (6)$$

where, $i \in [1, k]$.

In the above equation, Φ indicates the learning function of the proposed Multi-Head Attention Fusion layer. The main difference in our Fusion strategy compared to the previous work using Multi-Head Attention is that our strategy involves changing the key across modalities while keeping the Query and Value the same to better modulate inter utterance attention and incorporate inter-modal information.

The feature outcome of the last multi-head attention fusion layer (i.e., $F_{fusion_m}^i$) is concatenate with visual and acoustic feature maps (i.e., F_V and F_A) as expressed below:

$$F_{final}^i = \text{Concate}(F_{fusion_m}^i, F_A^i, F_V^i) \quad (7)$$

Finally, we append two fully connected layers (FC) which generates the desired predicted output (i.e., $Y_p^i = \langle y_1^i, y_2^i, \dots, y_p^i \rangle$, where, $i \in [1, k]$) of our proposed system.

$$y_p^i = FC(FC(F_{final}^i)). \quad (8)$$




Figure 4. Network design of the proposed Extractor network.

3.3. Feature Extractor Module

In order to fetch deep features from audio and visual contents, we propose a new feature extractor model and the same is illustrated in Figures 4. The proposed extractor is designed based on triplet network to leverage the importance of triplet loss function [25]. Initially, the anchor, positive and negative samples have been generated as suggested in [25] for audio and visual modalities. Then these samples are passed through encoder network followed by projector module. Here, we use standard ResNet18 [9] as a backbone of the encoder network while the projector consists a linear fully connected layer which project the embedding of encoder network to desired representations (i.e., $Z = [z_1, \dots, z_N] \in \mathcal{R}^{N \times d}$) composed of N representations with dimension d .

The proposed extractor model is trained using weighted combination of three loss function i.e., adaptive margin triplet loss (i.e., L_{AMT}), covariance loss (i.e., L_{Cov}) and variance loss (i.e., L_{Var}) functions. It can be expressed as follow:

$$L_{FE} = \lambda_1 \cdot L_{AMT} + \lambda_2 \cdot L_{Cov} + \lambda_3 \cdot L_{Var}, \quad (9)$$

where, λ_1, λ_2 and λ_3 are weighting factors that controls the distribution of different loss functions.

In [25], authors design the triplet loss function used to learn good representations of faces based on anchor, positive and negative samples. Here, authors have used a fixed margin value in their triplet loss function that helps to separate out the representations of positive and negative samples. However, in some cases where the positive or negative samples have the same distance with the anchor or the positive sample is only a bit closer to the anchor than the negative sample, the loss would be zero, and there would be no correction even though it should still be pulling the positive sample closer and pushing the negative sample away from the anchor. To overcome this issue, we propose a triplet loss function based on a adaptive margin value. This can be mathematically written as

$$L_{AMT} = D_s^{a,p} - \frac{D_s^{a,n} + D_s^{p,n}}{2} + m_{AM}. \quad (10)$$

Here, $D_s^{a,p}$, $D_s^{a,n}$ and $D_s^{p,n}$ denotes the euclidean distance based similarity metric between representations of anchor and positive, anchor and negative, positive and negative samples, respectively. m_{AM} is the adaptive margin which is calculated based on similarity and dissimilarity measures as

$$\begin{aligned} m_{AM} &= m_{AM}^{sim} + m_{AM}^{dissim} \\ m_{AM}^{sim} &= 1 + \frac{2}{e^{4 \cdot D_s^{a,p}}} \\ m_{AM}^{dissim} &= 1 + \frac{2}{e^{-4 \cdot D_s^{a,n} + 4}} \end{aligned} \quad (11)$$

In addition, we also utilize the variance loss function proposed by Bardes et al. [1] which helps the proposed model to tackle the mode collapse issue. Mathematically, the variance loss function can be presented as,

$$\begin{aligned} L_{Var} &= \sum_{k=1}^3 L_{Var}(Z_k); \quad Z_k = Z_a, Z_p, Z_n \\ L_{Var}(Z_k) &= \frac{1}{d} \sum_{j=1}^d 1 - \sqrt{Var(Z_{:,j})} + \epsilon \end{aligned} \quad (12)$$

Here, $Var(Z)$ (i.e., $\frac{1}{N-1} \sum_{i=1}^N (Z^i - \hat{Z})^2$) denotes the variance obtained from the corresponding representations, while \hat{Z} is mean of the corresponding representation.

To decorrelate the different dimensions of the representations, we adopt the covariance loss function [1] and the same can be expressed mathematically for representation as

$$\begin{aligned} L_{Cov} &= \sum_{k=1}^3 L_{Cov}(Z_k); \quad Z_k = Z_a, Z_p, Z_n \\ L_{Cov}(Z_k) &= \frac{1}{d} \sum_{i \neq j} Cov(Z_k)_{i,j}^T, \end{aligned} \quad (13)$$

where, $Cov(Z) = \frac{1}{N-1} \sum_{i=1}^N (Z^i - \hat{Z})(Z^i - \hat{Z})^T$ indicates the covariance matrix of corresponding representation.

4. Experimental Analysis and Discussion

4.1. Dataset

For fair comparison with state-of-the-art methods, we evaluate our proposed network (M2FNet) on Interactive Emotional Dyadic Motion Capture (IEMOCAP) [2], and Multimodal EmotionLines Dataset (MELD) [22] benchmark datasets. The statistics of these datasets are reported in Table 1 and details are provided in Section 1 of the supplementary material. Both IEMOCAP and MELD are multimodal datasets with textual, visual, and acoustic data.

MELD: The MELD [22] is a multimodal and multi-party dataset containing more than 1,400 conversations and

Table 1. Statistics of the testing benchmark datasets: MELD and IEMOCAP

Statistics	MELD		IEMOCAP	
Splitting	# Dialog	# Utterance	# Dialog	# Utterance
Train	1098	9989	100	4778
Dev	114	1109	20	980
Test	280	2610	31	1622
No. of Classes	7		6	

13,000 utterances from the Friends TV series. The utterances are annotated with one of the seven emotion labels (anger, disgust, sadness, joy, surprise, fear, and neutral). We use the pre-defined train/val split provided in the MELD dataset. The details of this dataset are given in Table 1.

IEMOCAP: The IEMOCAP database [2] is an acted, multimodal and multi-speaker database consisting of videos of dyadic sessions having approximately 12 hours of audio-visual data with text transcriptions. Each video contains a single conversation, which is segmented into multiple utterances. Each utterance is annotated with one of six emotion labels, i.e., happy, sad, neutral, angry, excited, and frustrated. The database statistics are given in Table 1. We randomly select 10% of training conversations as evaluation split for computing the hyperparameters.

4.2. Training Setups and Hyper-parameter Tuning

All experiments are carried out using a single NVIDIA GeForce RTX 3090 card. We adopt AdamW [17] as the optimizer with an initial learning rate $5e-4$ with L_2 weight decay ranges between $5e-4$ and $5e-5$. Dropout is used with a rate between 0.4 and 0.5. The number of encoder layers in each modality's encoder (i.e., N_T, N_A, N_V) is tuned using a greedy scheme and set to 1 and 5 for MELD and IEMOCAP validation datasets, respectively. The number of multi-head attention fusion layers is set to 5 (i.e., $= m$) for both dataset. The proposed M2FNet framework is trained using the categorical cross-entropy on each utterances softmax output for each of \mathbf{M} dialogs and their \mathbf{k} utterances. Section 2 of the supplementary material provides training and validation performance details.

$$Loss = -\frac{1}{M \cdot k} \cdot \sum_{i=1}^M \sum_{j=1}^k \sum_{t=1}^C y^{i,j,t} \cdot \log(y_p^{i,j,t}). \quad (14)$$

To extract deeper features from audio and visual contents, we introduce a new feature extractor model. Here, ResNet18 [9] is used as encoder module while the projector module consists a fully connected layer which projects the embeddings of encoder network to desired representations (i.e., Z). Here we set the number of representations as $Z = 300$. For audio task, the extractor model is trained on Mel Spectrograms obtained from the corresponding audio

Table 2. Ablation Studies based comparison to validate the impact of each modality.

Models	Remarks	IEMOCAP		MELD	
		Accuracy	Weighted Average F1	Accuracy	Weighted Average F1
Only Audio	—	26.56	21.79	49.04	39.63
Only Visual	—	20.39	13.10	45.63	32.44
Visual + Audio	Concat	35.12	31.35	48.35	35.74
Only Text	—	66.30	66.20	67.24	66.23
Text + Audio	Concat	66.52	66.48	67.80	66.32
Text + Visual	Concat	66.64	66.67	67.81	66.35
Text + Visual + Audio	Concat	67.16	67.12	67.28	66.81
Text + Visual + Audio	Fusion	69.69	69.86	67.85	66.71

Table 3. Ablation Studies based comparison to validate the impact of dual network.

Models	IEMOCAP		MELD	
	Accuracy	Weighted Average F1	Accuracy	Weighted Average F1
Scene Embeddings	67.65	67.70	65.23	64.23
Weighted Face Embeddings	67.32	67.28	66.37	65.67
Proposed	69.69	69.86	67.85	66.71

Table 4. Ablation Studies based comparison to validate the impact of no. of transformer encoders.

No. of Transformer Encoders (i.e., $N_A = N_T = N_V$)	IEMOCAP		MELD	
	Accuracy	Weighted Average F1	Accuracy	Weighted Average F1
1	69.32	69.36	67.85	66.71
2	69.07	69.21	67.16	65.82
3	69.32	69.44	66.86	66.06
4	69.13	69.22	67.20	66.06
5	69.69	69.86	67.47	66.38
6	69.13	69.26	67.24	65.48

signals while in case of visual feature extraction, it is trained on well-known CASIA webface database [35]. Here, the extractor model is trained using the loss function mentioned in Equation no. 9 in which the weighting factors λ_1 , λ_2 and λ_3 are set to 20, 5 and 1, respectively. The proposed extractor model is trained upto 60 and 100 epochs for audio and visual task, respectively using Adam optimizer with learning rate of 1e-4 and decay rate of 1e-6.

We mainly employ weighted average F1 score as evaluation metric due to its suitability to test with imbalance dataset. Additionally, we present our results in terms of classification accuracy to evaluate the model performance.

4.3. Ablation studies

To better understand the contribution of different modules in the proposed M2FNet model, we have conducted several ablation studies on both IEMOCAP and MELD datasets. The corresponding results are compared in terms of accuracy and weighted average F1 scores for MELD and IEMOCAP testing datasets.

To validate the impact of each modality, we train the proposed network with and/or without Text, Video, and Audio as input and without using the proposed Fusion module.

Table 5. Ablation Studies based comparison to validate the impact of multi-head attention fusion layers.

No. of Attention Fusion layers	IEMOCAP		MELD	
	Accuracy	Weighted Average F1	Accuracy	Weighted Average F1
1	67.22	67.34	66.63	66.11
2	68.08	68.19	67.47	66.50
3	69.13	69.25	67.47	66.65
4	69.01	68.95	67.59	66.51
5	69.69	69.86	67.85	66.71
6	68.95	69.08	66.82	65.32

From Table 2, one can observe that concatenation of multi-modal input with all three modalities obtain higher accuracy and weighted average F1 score than other scenarios such as using only one or two modalities. Furthermore, our fusion mechanism helps to enhance the accuracy by 2.53% and 0.57% for IEMOCAP and MELD datasets, respectively. However, in the case of weighted average F1 score, it obtains slightly inferior performance for the MELD dataset while improving it by 2.74% for the IEMOCAP dataset.

On the visual end, we utilize scene and weighted face embeddings. For understanding the importance of both embeddings, two more experiments have been carried out where the proposed network with individual embedding has been trained, and the corresponding results are given in Table 3. From the table, it can be observed here that the weighted face model or the scene encoding network on their own do not improve the results; however, when the network can access both, it significantly improves the results. This shows that both the context from the scene and the people in the scene are equally important for emotion recognition.

We also observe the effect of transformer encoders in the proposed framework. In these experiments, we set same number of transformer encoders (i.e., $N_A = N_V = N_T$) for each modality and observe its effect for different numbers. The corresponding results are presented in Table 4 where it is observed that $N_A = N_V = N_T = 1$ gives best performance for MELD dataset while the $N_A = N_T = N_V = 5$ setting helps the proposed framework to obtain higher performance for IEMOCAP testing dataset.

In the proposed model, we have set $m = 5$ number of Multi-Head Attention Fusion modules. To validate this, we train the proposed model with different numbers of the Multi-Head Attention Fusion modules and observe the corresponding accuracy and weighted average F1 score. This analysis is demonstrated in Table 5 where one can observe that the proposed model with five Multi-Head Attention Fusion modules (i.e., $m = 5$) obtains higher quantitative measures on both datasets.




Figure 5. Predictions made by the network on the MELD and IEMOCAP test sets

Table 6. Quantitative comparison with text-based state-of-the-art methods in terms of weighted average F1 score.

Name	Year	MELD	IEMOCAP
BERT+MTL [16]	2020	61.90	—
Hi-Trans [14]	2020	61.94	64.50
TRMSM-Att [15]	2020	62.36	65.74
DialogXL [26]	2020	62.41	65.94
COSMIC [6]	2020	65.21	65.28
SumAggGIN [28]	2020	58.45	66.61
DialogueGCN [7]	2019	58.10	66.76
CESTa [32]	2020	58.36	67.10
DialogueCRN [10]	2021	58.39	66.20
EmoBERTa [11]	2021	65.61	67.42
DAG-ERC [27]	2021	63.65	68.03
TODKAT [38]*	2021	65.47	61.33
EmotionFlow [33]	2022	66.50	—
M2FNet	2022	66.71	69.86

* The performance of TODKAT model is updated by its author in their official repository.

4.4. Model Performance

Figure 5 presents the performance of our model in terms of the weighted average F1 score of different emotions. Here, we can see that our model has obtained the highest F1 score of 82.11% for the Sad emotion of the IEMOCAP dataset. Similarly, for MELD dataset, it obtained highest score for Neutral emotion.

4.5. Comparative Analysis

To validate the robustness of the proposed network, we compare our proposed network with state-of-the-art text-based ERC systems in terms of weighted average F1 score and the same is presented in Table 6. Here, one can notice that the proposed network has a state-of-the-art performance by obtaining superior quantitative results than previous methods on both datasets (i.e., 0.21% higher than previous best EmotionFlow [33] model on MELD dataset while 1.83% higher than that of the previous best DAG-ERC [27] model on IEMOCAP testing dataset).

When compared with the existing multi-modal methods, our proposed M2FNet network shows a substantial im-

Table 7. Quantitative comparison with multimodal-based state-of-the-art methods on MELD and IEMOCAP datasets. Here, top two performances are highlighted with bold font texts.

Name of Model	MELD		IEMOCAP	
	Accuracy	Weighted Average F1	Accuracy	Weighted Average F1
BC-LSTM-Att [21]	57.50	56.44	56.32	56.19
DialogRNN [19]	59.54	57.03	63.40	62.75
ConGCN [36]	—	57.40	64.18	64.18
Xie et al. [34]	65.00	64.00	—	—
DialogueTRM [20]	65.66	63.55	68.92	69.23
M2FNet	67.85	66.71	69.69	69.86

provement compared to state-of-the-art multi-modal ERC methods. The obtained results are presented in Table 7 where one can notice that the proposed M2FNet model obtains 2.19% and 2.71% higher accuracy and weighted average F1 score on MELD dataset than that of previous best performance. Similarly, it set 0.77% and 0.63% higher accuracy and weighted average F1 score than that of previous best DialogueTRM model [20] on IEMOCAP testing dataset.

5. Limitations

Our model sometimes gets confused and miss-classifies similar or close emotions such as Frustration and Anger, Happy and Excited. We can also observe that for highly imbalance data, our model misclassifies many emotions as the emotion with higher number of data samples. For example, many emotions are overwhelmingly predicted as *Neutral* for MELD dataset.

6. Conclusion

In this paper, we propose a robust multi-modal fusion network called M2FNet for the task of Emotion Recognition in Conversation. In M2FNet, we propose a multi-head fusion attention module that helps the network to extract rich features from multiple modalities. A new feature extractor model is introduced in the proposed design to learn the audio and visual features effectively. Here, a new adaptive margin triplet loss function is introduced, which helps the extractor module to learn representations effectively. A new weighted face model is proposed in our framework to learn the rich facial features. Detailed analysis shows that encoding both scene and face-related information is essential for emotion recognition. Similarly, we observed that multi-modal fusion is necessary to leverage information from multiple modalities present in an utterance. Finally, our experiments validate the robustness of the proposed network quantitatively on both benchmark datasets.

References

- [1] Adrien Bardes, Jean Ponce, and Yann LeCun. Vireg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- [2] Carlos Busso, Murtaza Bulut, et al. Iemocap: interactive emotional dyadic motion capture database. *Language Resources and Evaluation*, 42:335–359, 2008.
- [3] Jean-Benoit Delbrouck, Noé Tits, Mathilde Brousmiche, and Stéphane Dupont. A transformer-based joint-encoding for emotion recognition and sentiment analysis. *CoRR*, abs/2006.15955, 2020.
- [4] Laurence Devillers and Laurence Vidrascu. Real-life emotions detection with lexical and paralinguistic cues on human-human call center dialogs. In *Ninth International Conference on Spoken Language Processing*, 2006.
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [6] Deepanway Ghosal, Navonil Majumder, Alexander F. Gelbukh, Rada Mihalcea, and Soujanya Poria. COSMIC: commonsense knowledge for emotion identification in conversations. *CoRR*, abs/2010.02795, 2020.
- [7] Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Alexander Gelbukh. Dialoguecn: A graph convolutional neural network for emotion recognition in conversation. *arXiv preprint arXiv:1908.11540*, 2019.
- [8] Devamanyu Hazarika, Soujanya Poria, Amir Zadeh, Erik Cambria, Louis-Philippe Morency, and Roger Zimmermann. Conversational memory network for emotion recognition in dyadic dialogue videos. In *Proceedings of the conference. Association for Computational Linguistics. North American Chapter. Meeting*, volume 2018, page 2122. NIH Public Access, 2018.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] Dou Hu, Lingwei Wei, and Xiaoyong Huai. Dialoguecrn: Contextual reasoning networks for emotion recognition in conversations. *arXiv preprint arXiv:2106.01978*, 2021.
- [11] Taewoon Kim and Piek Vossen. Emoberta: Speaker-aware emotion recognition in conversation with roberta. *arXiv preprint arXiv:2108.12009*, 2021.
- [12] Chul Min Lee and Shrikanth S Narayanan. Toward detecting emotions in spoken dialogs. *IEEE transactions on speech and audio processing*, 13(2):293–303, 2005.
- [13] Jiyoung Lee, Seungryong Kim, Sunok Kim, Jungin Park, and Kwanghoon Sohn. Context-aware emotion recognition networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10143–10152, 2019.
- [14] Jingye Li, Donghong Ji, Fei Li, Meishan Zhang, and Yijiang Liu. HiTrans: A transformer-based context- and speaker-sensitive model for emotion detection in conversations. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4190–4200, Barcelona, Spain (Online), Dec. 2020. International Committee on Computational Linguistics.
- [15] Jiangnan Li, Zheng Lin, Peng Fu, Qingyi Si, and Weiping Wang. A hierarchical transformer with speaker modeling for emotion recognition in conversation. *arXiv preprint arXiv:2012.14781*, 2020.
- [16] Jingye Li, Meishan Zhang, Donghong Ji, and Yijiang Liu. Multi-task learning with auxiliary speaker identification for conversational emotion recognition. *arXiv preprint arXiv:2003.01478*, 2020.
- [17] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [18] Navonil Majumder, Devamanyu Hazarika, Alexander Gelbukh, Erik Cambria, and Soujanya Poria. Multimodal sentiment analysis using hierarchical fusion with context modeling. *Knowledge-based systems*, 161:124–133, 2018.
- [19] Navonil Majumder, Soujanya Poria, Devamanyu Hazarika, Rada Mihalcea, Alexander F. Gelbukh, and Erik Cambria. Dialoguecrn: An attentive RNN for emotion detection in conversations. *CoRR*, abs/1811.00405, 2018.
- [20] Yuzhao Mao, Qi Sun, Guang Liu, Xiaojie Wang, Weiguo Gao, Xuan Li, and Jianping Shen. Dialoguetrm: Exploring the intra-and inter-modal emotional behaviors in the conversation. *arXiv preprint arXiv:2010.07637*, 2020.
- [21] Soujanya Poria, Erik Cambria, et al. 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)*, pages 873–883, Vancouver, Canada, 2017. Association for Computational Linguistics.
- [22] Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. Meld: A multimodal multi-party dataset for emotion recognition in conversations. *arXiv preprint arXiv:1810.02508*, 2018.
- [23] Soujanya Poria, Navonil Majumder, Rada Mihalcea, and Edward H. Hovy. Emotion recognition in conversation: Research challenges, datasets, and recent advances. *IEEE Access*, 7:100943–100953, 2019.
- [24] Lawrence Rabiner and Ronald Schafer. *Theory and applications of digital speech processing*. Prentice Hall Press, 2010.
- [25] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [26] Weizhou Shen, Junqing Chen, Xiaojun Quan, and Zhixian Xie. Dialogxl: All-in-one xlnet for multi-party conversation emotion recognition. *arXiv preprint arXiv:2012.08695*, 2020.
- [27] Weizhou Shen, Siyue Wu, Yunyi Yang, and Xiaojun Quan. Directed acyclic graph network for conversational emotion recognition. *arXiv preprint arXiv:2105.12907*, 2021.
- [28] Dongming Sheng, Dong Wang, Ying Shen, Haitao Zheng, and Haozhuang Liu. Summarize before aggregate: A global-to-local heterogeneous graph inference network for conversational emotion recognition. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4153–4163, Barcelona, Spain (Online), Dec. 2020.

- [29] Stanley Smith Stevens, John Volkman, and Edwin Broomell Newman. A scale for the measurement of the psychological magnitude pitch. *The journal of the acoustical society of america*, 8(3):185–190, 1937.
- [30] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2019, page 6558. NIH Public Access, 2019.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [32] Yan Wang, Jiayu Zhang, et al. Contextualized emotion recognition in conversation as sequence tagging. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 186–195, 1st virtual meeting, July 2020. Association for Computational Linguistics.
- [33] Rong Zhang Songlin Hu Xiaohui Song, Liangjun Zang and Longtao Huang. Emotionflow: Capture the dialogue level emotion transitions. In *Accepted in ICASSP*, 2022.
- [34] Baijun Xie, Mariia Sidulova, and Chung Hyuk Park. Robust multimodal emotion recognition from conversation with transformer-based crossmodality fusion. *Sensors*, 21(14):4913, 2021.
- [35] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014.
- [36] Dong Zhang, Liangqing Wu, Changlong Sun, Shoushan Li, Qiaoming Zhu, and Guodong Zhou. Modeling both context- and speaker-sensitive dependence for emotion detection in multi-speaker conversations. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5415–5421, 7 2019.
- [37] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016.
- [38] Lixing Zhu, Gabriele Pergola, Lin Gui, Deyu Zhou, and Yulan He. Topic-driven and knowledge-aware transformer for dialogue emotion detection. *arXiv preprint arXiv:2106.01071*, 2021.