

CROSS-MODAL FUSION TECHNIQUES FOR UTTERANCE-LEVEL EMOTION RECOGNITION FROM TEXT AND SPEECH

Jiachen Luo¹, Huy Phan^{2*}, Joshua Reiss¹

Centre for Digital Music, Queen Mary University of London, UK
²Amazon Alexa, Cambridge, MA, USA

ABSTRACT

Multimodal emotion recognition (MER) is a fundamental complex research problem due to the uncertainty of human emotional expression and the heterogeneity gap between different modalities. Audio and text modalities are particularly important for a human participant in understanding emotions. Although many successful attempts have been designed multimodal representations for MER, there still exist multiple challenges to be addressed: 1) bridging the heterogeneity gap between multimodal features and model inter- and intra-modal interactions of multiple modalities; 2) effectively and efficiently modeling the contextual dynamics in the conversation sequence. In this paper, we propose Cross-Modal RoBERTa (CM-RoBERTa) model for emotion detection from spoken audio and corresponding transcripts. As the core unit of the CM-RoBERTa, parallel self- and cross- attention is designed to dynamically capture inter- and intra-modal interactions of audio and text. Specially, the mid-level fusion and residual module are employed to model long-term contextual dependencies and learn modality-specific patterns. We evaluate the approach on the MELD dataset and the experimental results show the proposed approach achieves the state-of-art performance on the dataset.

Index Terms— speech emotion recognition, multimodal fusion, deep learning

1. INTRODUCTION

Emotion recognition in conversations (ERC) is vital and very challenging in the natural human machine interaction [1], intelligent education tutoring [2], and mental health analysis applications [3]. In daily life, humans utter a multi-turn conversation in a natural way which conveys emotion state through language and nonverbal content (e.g., facial expression and body language) [4]. Different modalities all carry emotion-relevant information and how to efficiently and effectively fuse heterogeneous inputs across multiple modalities has been an active research focus.

In daily life, text modality is often accompanied by audio modality. The sentiment content contained in text features is complemented by the multiple information such as pitch, loudness, and other frequency-related measures [5,6]. The effective interaction between text and audio data enables more comprehensive content and provides more emotional information [7~9]. As the example shown in Fig. 1, a single word “Okay” is ambiguous, and it can express different emotions. It is challenging to determine the associated emotion according to one turn of textual utterance. After

introducing corresponding low voices and sobs, it is not difficult to discern the emotion sentiment of this sentence is negative. However, fusing different modalities in an effective manner is a non-trivial task that researchers often have to face.

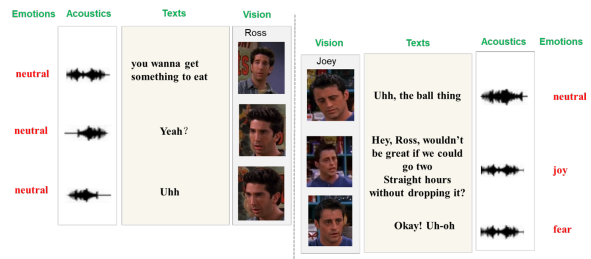


Fig. 1. Emotion dynamic of speakers in a dialogue in comparison.

Previous methods have achieved good performance [8,9], there still exist key challenges in multimodal emotion recognition: different modalities depend on independent preprocessing and feature excavation designs due to heterogeneous space; and to make an applicable and generalizable model for both the individual modalities and fusion model, it is necessary to learn intra- and cross-modal interactions to reveal discriminative emotion content; and emotion is a subjective concept [10~12]. Moreover, textual information and the associated context conveys more influential information in inferring the speaker's emotions [13], and plays a crucial role in inferring emotions in conversations. Unfortunately, the existing approaches cannot efficiently capture the textual emotional-relevant content in the fusion processing, and often the learning of intra- and cross-modal information will lose some semantic content [12,14].

To tackle the above problems, we proposed a Cross-Modal RoBERTa to integrate the information of audio modality in fine-tuning the pre-trained RoBERTa. As the unit of the CM-RoBERTa, cross- and self-attention layers were designed to learn inter- and intra-modality information, and could adaptively extend the interactions and correlations of information from different sources. In addition, we designed a temporal calibration module to effectively preserve the local information (text and audio) from a single modality and the overall information generated from the contextual level. We evaluate it on the public multimodal sentiment dataset MELD [15]. The experiment results showed a significant improvement in model performance and efficiency. The main contributions of this paper could be summarized as follows:

- We propose a CM-RoBERTa model that introduces the information audio modality to help text-modality fine-tuning of the pre-trained RoBERTa model.
- We employ parallel self- and cross-attention that can dynamically adjust the attentive weights through

*The work was done when H. Phan was at Centre for Digital Music, Queen Mary University of London, UK and prior to joining Amazon.

Thanks to the China Scholarship Council and Queen Mary University of London for funding.

the interaction between audio and text to capture inter- and intra-modal information. • We utilize mid-level fusion and residual connection to capture long-term contextual dependencies, and learn modality-specific patterns. • Our proposed approach is superior to state-of-the-art methods for conversational emotion recognition.

2. RELATED WORKS

Emotion recognition in conversations is a popular research area because of their applications in various areas. Since human emotions are expressed in many ways, and there are some limitations in the recognition of emotions by single modal, multimodal emotion recognition has gradually become a new trend.

To date, multimodal emotion recognition is roughly divided into four categories [14~21]: early fusion, late fusion, model-level fusion, and hybrid fusion. Early fusion is to fuse input features from different modalities through concatenating and other ways [15, 16]. Poria et al. leveraged context-aware utterance representations for emotion recognition [15]. Due to the complex inter-modality dynamics at the input level, this feature fusion approach was incapable of capturing the intra-modality information. Late fusion, instead, trains an independent classification model of each modality and utilizes decision voting or other classifier combination strategies [17,22]. Xie et al. employed crossmodal transformer fusion for audio, text and video late-level fusion, avoiding the difficulty of fusing heterogeneous information [22]. However, such approaches ignored associations across multiple modalities and failed to model mutual correlations.

On the other hand, model-level fusion is a compromise between the former two where the fusion happens between the intermediate representations of the multimodal features [20, 23]. Akhtar et al. designed a context-level inter-modal attention module for learning the joint association between textual-audio-visual information of utterances [20]. Hybrid fusion combines early fusion and late fusion [14, 24]. Wollmer et al. fused audio-video information with Bi-LSTM for audio-visual emotion recognition[14]. How to design an effective and efficient structure to capture inter- and intra- modalities features is necessary in multimodal emotion recognition.

With the popularity of the attention mechanism, it has recently attracted some in the field for its ability to fuse multimodal information. Devamanyu et al. used a multi-attention recurrent network framework that learned multimodal interactions [25]. Also, recently transfer learning techniques that utilized pre-trained networks for feature extraction have become popular [26]. RoBERTa, a BERT-based model, has achieved better performance by fine-tuning from pre-trained weights for emotion recognition tasks [27]. In our work, we proposed an effective and efficient fusion approach for fine-tuning by fusing heterogeneous nonverbal information (audio) that complements the linguistic expressions of RoBERTa.

3. METHODOLOGY

We leveraged the multi-modal and contextual information for discerning the discriminative emotion for an utterance. The architecture of the model is shown in Fig. 2.

3.1. Audio and Text Feature Extraction

The input of the CM-RoBERTa model consists of the effective heterogeneous representation (handcrafted and bottleneck features) for

each audio and text from the data (see Fig. 2). For each audio utterance, we extract 6552 low-level handcrafted features (a^{LLF}) constituting several descriptors and various statistical functionals of varied vocal and prosodic features using the openSMILE toolkit [28]. In addition, a pre-trained OpenL3 model [29] is employed to extract bottleneck embedding (a^{OpenL3}). a^{LLF} and a^{OpenL3} are then concatenated to obtain audio presentation $x^a = (a^{LLF} \oplus a^{OpenL3})$. Where \oplus denotes concatenation. For text modality, the RoBERTa encoder ($t^{RoBERTa}$) received the transcript of the utterance to produce activations from the final four layers. Then, these four were concatenated to obtain the context-independent utterance feature vector x^t with a dimension of 4096 [27].

3.2. Cross-Modal RoBERTa

As shown in Fig. 2, CM-RoBERTa first uses a bidirectional LSTM to encode acoustic (H_a) and textual features (H_t). We input them (H_a and H_t) into the self- and cross-attention unit (SCA) to capture intra- and inter- modal interactions between each pair of acoustic frames and textual words. The SCA module consists of two stacked parallel cross- and self-attention layers and feed-forward layers. The main insight of the module is to learn modality associations and modality-specific patterns, and then propagate information from intra- and inter-modal perspectives according to the attentive weights.

Following [30], we estimate the associations between audio and text modalities in a crossed way, whose query (Q_m), key (K_m), value (V_m) are the representation of modality m , H_m , under different projection space, where m is either audio (a) or text (t):

$$\nabla H_{a-t} = softmax(Q_a K_t^T / \sqrt{d}) V_t, \quad (1)$$

$$\nabla H_{t-a} = softmax(Q_t K_a^T / \sqrt{d}) V_a, \quad (2)$$

where ∇H_{a-t} and ∇H_{t-a} represent the propagated information from audio to text and text to audio, respectively. d denotes the key dimensionality. Later, we update the features of one modality with the information propagating from the other modality:

$$H_a^{L^N} = LayerNorm(H_a \oplus \nabla H_{t-a}), \quad (3)$$

$$H_t^{L^N} = LayerNorm(H_t \oplus \nabla H_{a-t}). \quad (4)$$

To further enhance the representation capacity, we feed the learned weighed information into the full connection to obtain the final cross-attention information between audio and text modalities:

$$H^{c^a} = LayerNorm(H_a^{L^N} \oplus \nabla FeedForward(H_a^{L^N})), \quad (5)$$

$$H^{c^t} = LayerNorm(H_t^{L^N} \oplus \nabla FeedForward(H_t^{L^N})). \quad (6)$$

We denote the final cross-attention information between audio and text modalities as H_a^c and H_t^c , respectively. Where \oplus is concatenation operation.

The self-attention mechanism shares the similar principle with the cross-attention mechanism. The only different is the query, key, and value are from the same modality. The whole process for the one stacked layer in the self-attention module can be summarized as:

$$\nabla H_m = softmax(Q_m K_m^T / \sqrt{d}) V_m, \quad (7)$$

$$H_m^{LN^s} = LayerNorm(H_m \oplus \nabla H_m), \quad (8)$$

$$H_m^s = LayerNorm(H_m^{LN^s} \oplus \nabla FeedForward(H_m^{LN^s})), \quad (9)$$

where ∇H_m is the propagated information within modality m , $m \in \{a, t\}$. We denote the outputs of the last stacked layer in two self-attention modules as H_a^s and H_t^s , respectively. In parallel, we then

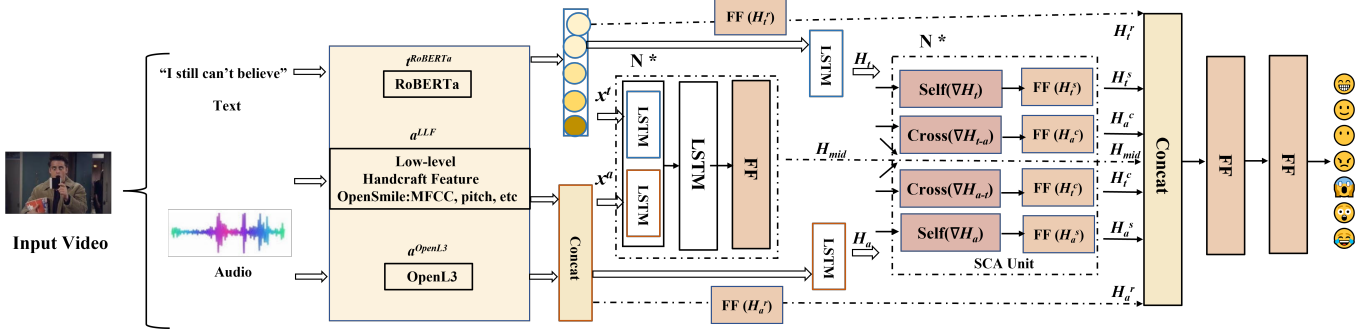


Fig. 2. Overview architecture of the Cross-Modal RoBERTa Fusion Network N represents two layers, and the first two parallel LSTM are exactly the same as the last two parallel LSTM.

perform mid-level fusion of the joint features with modality-specific patterns. They help in capturing the inter- and intra-modality dynamics. Individual modalities (x_a, x_t) are passed in parallel through two stacked LSTM layers separately to store single modality feature (H_a^m, H_t^m) efficiently, then combined via a bidirectional LSTM to explore modality associations:

$$H_{mid} = FeedForward(LSTM([H_a, H_t])). \quad (10)$$

Following prior works [13], we utilized a residual connection on audio and textual representation to keep the original structure of the data (H_a^r, H_t^r). Then it is passed through a linear and a normalization layer. Finally, we concatenate them ($H_a^c, H_t^c, H_a^s, H_t^s, H_{mid}, H_a^r, H_t^r$) as the aggregating representation and presented it to two fully-connected layers for classification.

4. EXPERIMENTS

4.1. Dataset and Feature Sets

4.1.1. Dataset

Multi-modal Emotion Dataset (MELD) is a multi-modal (visual, audio, and text) and multi-speaker emotion recognition dataset consisting of dialogues from the Friends TV show [15]. MELD has more than 13000 utterances and 1400 conversations. Each utterance is annotated using the following discrete categories: anger, joy, sadness, neutral, disgust, fear and surprise. The dataset is split into the training set, validation set and test set which contains 9989, 1109, and 2610 utterances [15]. Weighted F1-score was used to evaluate the performance, as the class distribution is highly imbalanced.

4.1.2. Audio Features

The audio features consisted of two types of features for each utterance: handcrafted and bottleneck features. As handcraft features, we extracted 6552 features using the openSMILE toolkit [28]. They encapsulated features like MFCC, pitch, etc. In addition, we extracted deep-learning-based bottleneck features using the OpenL3 model [29]. OpenL3 are self-supervised embeddings learned from AudioSet [29]. The OpenL3 network contains four blocks stacked together, each block consists of three convolutional layers and one pooling layer. We processed audio data to conform to the network input and extracted 6144 high-level bottleneck features.

4.1.3. Textual Features

We fine-tuned the pre-trained RoBERTa Large model (Uncased: hidden-1024, heads-24, layer-24) as our language model for extracting the features from the text modality [27]. RoBERTa took the transcript of the utterance as input and produced rich contextual textual representation from the final four layers. This produced four 1024-dimension vectors for every token in the input. Then, we concatenated them to obtain the context independent utterance feature vector with a dimension of 4096.

4.2. Baselines

We compared our model with the following baselines and state-of-the-art models in multi-modal emotion recognition (Table 1).

- SMIN integrated semi-supervised learning with multimodal interactions for conversational emotion recognition [31].
- EmoCaps utilized the transformer structure for contextual emotion classification [17].
- EmbraceNet employed transformer-based crossmodality fusion to estimate the emotion [22].
- CMN used attention-based hops to capture inter-speaker dependencies [25].
- ICON employed a memory network to store contextual summaries for classification [32].
- DialogueRNN employed three GRUs to model the speaker, the context and the emotion of the preceding utterances [16].
- bc-LSTM performed context-dependent sentiment analysis and emotion recognition [15].

4.3. Experimental Setup

The model is trained with batch-size 32 for 200 epochs and stop training if the validation loss does not decrease for 15 consecutive epochs. The number of hidden neurons in the model is 128. We employed Adam optimizer with the learning rate of $1e-4$. The L_2 weight decay was set to $3e-4$. Cross-entropy loss with softmax activation in the last layer was used for network training.

5. RESULTS AND DISCUSSION

Table 1 summarizes the performance of unimodal and multimodal variants of the state-of-the-art text-based ERC systems along with the proposed CM-RoBERTa model. It should be highlighted that the text modality performs noticeably better than the audio modality (see Tables 1 and 2). However, for the MELD dataset, several factors make emotion detection using audio considerably hard. First, a lot of background noise exist. Second, the length of utterances is shorter

Table 1: Overall performance comparison between the proposed method, the state-of-the-art, and the baselines. The weighted F1-scores were used as the metric. Note that our unimodal audio-based and text-based models are comprised of LSTM layer, and residual block.

model	audio	text	bimodal	bimodal - text	reference
SMIN	45.2	63.65	64.5	0.85	[31]
EmoCaps	31.26	63.51	63.73	0.22	[17]
EmbraceNet	40	61	63.1	2.1	[22]
CMN	38.3	54.5	55.5	1	[25]
ICON	37.7	54.6	56.3	1.7	[32]
DialogueRNN	34	55.1	57	1.9	[16]
bc-LSTM	36.4	54.3	56.8	2.5	[15]
Proposed method without SCA unit	43.10	63.83	65.66	1.83	
Proposed method	43.10	63.83	66.28	2.45	

Table 2: The class-wise weighted F1 scores on the MELD dataset

Proposed method	neutral	surprise	fear	sadness	joy	disgust	anger	w-average F1
audio modality	62.9	20.12	0.0	6.7	22.65	0.0	53.21	43.10
text modality	76.39	60.47	30.5	42.21	63.21	31.74	45.82	63.83
bimodality	78.67	63.42	28.7	42.19	63.65	30.91	53.56	66.28

than other benchmark datasets (e.g., IEMOCAP [33]). Third, there are more than 2 speakers in most of the conversations.

Text tends to effectively alleviate the above problems by providing rich-emotional features in the joint representation. The word-level textual content carries a wealth of information, such as viewpoint, intent, and so on. For short communication, a specific word (e.g., amazing) may directly represent the emotional state and dominate the final decision. In our work, we employed the fine-tuned RoBERTa model to extract textual utterance level features. Compared with other text embedding (see Table 2), RoBERTa attains marginally better results [27], which may be attributed to dynamically changing mask pattern and training the model longer with more data, enabling better marriage with the large-scale language model.

However, emotion recognition in conversations sometimes is dependent to its audio, rather than text, for inferring emotions. We notice the openL3 embeddings play significantly better than low-level traditional acoustic features: pitch, MFCC, and other audio features for prediction (Table 3). Our proposed model with openL3 embedding exceeds the baseline bc-LSTM model by 4.9% in terms of F1-score. The pre-trained OpenL3 model is useful in limited and unbalanced MELD dataset as knowledge transfer should contribute to better capture high-level emotion-relevant audio features. Additionally, by combing low-level descriptors and high-level OpenL3 acoustic representation, surprise and anger can gain 2.95%, and 7.74% on F1-score, respectively, over the text-based model (Table 2). It is showed that anger is best recognized from speech.

Introduction of the audio information to text modality can empower significant advantages, as it allows for improved emotion detection by helping to fuse additional audio features, disambiguate linguistic meaning and bridge a gap to real-world environments. Multimodality extracts and fuses vital information from the respective modality source and captures richer representation and performance than the individual modalities. The inter, intra and cross-modal interactions, correlations and relationships between multiple modalities are release more emotion-relevant content for better performance and prediction.

Table 3: Ablation audio features study on the MELD dataset

Method	w-average F1
w/o low-level descriptors	37.3
w/o openL3 embedding	41.3
our audio features	43.1

Our proposed multimodal model gained 2.45% on F1-score over the text-based unimodal model, and SCA unit improvements about 1.83% on F1-score (see Table 1). The superior performances of CM-RoBERTa might benefit for the following reasons. First, the early multi-modal cross-attention fused the underlying temporal interaction between text and audio features. We introduced cross-attention calibration by temporal-awareness to deduce the mutually exclusive and dependent between heterogeneous knowledge. By giving more weight to relevant audio modality, we could introduce the information of the audio modality to help the text modality effectively adjust the weight of words. We speculate that our cross-attention fusion unit can effectively employ the different knowledge into a new space to strengthen the details embedded in a single modality and adaptively fuse the implicit complementary content to magnify the interactions and correlations. Second, the self-attention mechanism was also employed to propagate information within each modality. By modeling the intra-modal interactions, we can capture learn modality-specific patterns, and distribute the self-attention weights over the entire dialog sequence for each modality. Third, mid-level fusion and residual blocks were targeted at explicitly modeling both inter- and intra-modal interaction of audio and text. This aspect has the power to fuse heterogeneous space, retain contextual intra-modality-specific representation and facilitate synchronization.

6. CONCLUSION

In this paper, we proposed the CM-RoBERTa model that effectively fused multi-modal emotion recognition using a fine-tuned RoBERTa model for languages and heterogeneous features unification for handcrafted and bottleneck audio features. Thanks to the parallel self- and cross-attention modules, mid-level fusion and residual blocks that could explicitly model inter- and intra-modal interactions within/between audio and text, capture long-term contextual dependencies, and learn modality-specific patterns. They achieved state-of-the-art utterance-level recognition performance when evaluated on the MELD dataset.

However, the work relied on the controlled database with non-realistic recording conditions and error-less text transcriptions. Thus, in practice, more robust models need to explore for the problem of multimodal emotion recognition in conversations that are: (i) performed under unexpected automatic speech recognition errors; (ii) explored inference in cases of noisy or absent modalities.

7. REFERENCES

- [1] Farhoudi Z, Setayeshi S, “Fusion of deep learning features with mixture of brain emotional learning for audio-visual emotion recognition,” *Speech Communication*, 2021, 127:92-103.
- [2] Yang K, Xu H, Gao K, “Cross-Modal BERT for Text-Audio Sentiment Analysis,” *Paper presented at Proceedings of the 28th ACM International Conference on Multimedia*, 2020.
- [3] Lee S, Han D, Ko H, “Multimodal Emotion Recognition Fusion Analysis Adapting BERT With Heterogeneous Feature Unification,” *IEEE Access*, 2021, 9:94557-94572.
- [4] Sun L, et al., “Multimodal Cross- and Self-Attention Network for Speech Emotion Recognition,” *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, 4275-4279.
- [5] Chi P-H, et al., “Audio Albert: A Lite Bert for Self-Supervised Learning of Audio Representation,” *Paper presented at 2021 IEEE Spoken Language Technology Workshop (SLT)*, 2021.
- [6] Ancilin J, Milton A, “Improved speech emotion recognition with Mel frequency magnitude coefficient,” *Applied Acoustics*, 2021, 179.
- [7] Aneesh M, Martin R, “Speech Emotion Recognition Using Quaternion Convolutional Neural Networks,” *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, 6309-6313.
- [8] Kumar P, et al., “End-to-end Triplet Loss based Emotion Embedding System for Speech Emotion Recognition,” *Paper presented at 25th International Conference on Pattern Recognition (ICPR)*, 2021.
- [9] Neil S, et al., “SERAB: A Multi-Lingual Benchmark for Speech Emotion Recognition,” *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, 7697-7701.
- [10] Sailunaz K, et al., “Emotion detection from text and speech: a survey,” *Social Network Analysis and Mining*, 2018, 8.
- [11] Khalil R, et al., “Speech Emotion Recognition Using Deep Learning Techniques: A Review,” *IEEE Access*, 2019, 7:117327-117345.
- [12] Kim T, Vossen P, “EmoBERTa: Speaker-Aware Emotion Recognition in Conversation with RoBERTa,” *arXiv*, <https://arxiv.org/abs/2108.12009>, 2021.
- [13] KrishnaD N, Ankita P, “Multimodal Emotion Recognition Using Cross-Modal Attention and 1D Convolutional Neural Networks,” *Interspeech*, 2020, 4243-4247.
- [14] Wang Y, et al., “A Systematic Review on Affective Computing: Emotion Models, Databases, and Recent Advances,” *Information Fusion*, 2022, 19-52.
- [15] Soujanya P, et al., “MELD: A multimodal multi-party dataset for emotion recognition in conversations,” *In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, 527-536.
- [16] Navonil M, et al., “DialogueRNN: An Attentive RNN for Emotion Detection in Conversations,” *In Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, 6818-6825.
- [17] Li Z, et al., “EmoCaps: Emotion Capsule based Model for Conversational Emotion Recognition,” *In Findings of the Association for Computational Linguistics: ACL*, 2022, 1610-1618.
- [18] Abbaschian BJ, Sierra-Sosa D, Elmaghraby A, “Deep Learning Techniques for Speech Emotion Recognition, from Databases to Models,” *Sensors (Basel)*, 2021, 21(4).
- [19] Zhou J, et al., “Inferring Emotion from Conversational Voice Data: A Semi-Supervised Multi-Path Generative Neural Networks Approach,” *The Thirty-Second AAAI Conference on Artificial Intelligence*, 2018, 32:579-586.
- [20] Akhtar MS, et al., “Multi-task Learning for Multi-modal Emotion Recognition and Sentiment Analysis,” *arXiv*, <https://arxiv.org/abs/1905.05812>, 2019.
- [21] Zadeh A, et al., “Multi-attention Recurrent Network for Human Communication Comprehension,” *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- [22] Xie B, Sidulova M, Park CH, “Robust Multimodal Emotion Recognition from Conversation with Transformer-Based Crossmodality Fusion,” *Sensors*, 2021, 21(14):4913-4932.
- [23] Asif IM, et al. “Deep learning based multimodal emotion recognition using model-level fusion of audio-visual modalities”, *Knowledge-Based Systems*, 2022, 244, 1-14.
- [24] Yucel C, et al. “Cross-subject multimodal emotion recognition based on hybrid fusion”, *IEEE Access*, 2020, 8, 168865-168878.
- [25] Devamanyu H, et al., “Conversational Memory Network for Emotion Recognition in Dyadic Dialogue Videos,” *In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019, 1:2122-2132.
- [26] Feng K, Chaspari T, “A Review of Generalizable Transfer Learning in Automatic Emotion Recognition,” *Frontiers in Computer Science*, 2020, 2.
- [27] Ghosal D, et al., “COSMIC: CommonSense knowledge for eMotion Identification in Conversations,” *Findings of the Association for Computational Linguistics (EMNLP)*, 2020.
- [28] Eyben F, Willmer M, Schuller B, “Opensmile: the munich versatile and fast open-source audio feature extractor,” *Acm International Conference on Multimedia*, 2010.
- [29] Cramer J, et al., “Look, Listen, and Learn More: Design Choices for Deep Audio Embeddings,” *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [30] Vaswani A, et al., “Attention Is All You Need,” *Advances in Neural Information Processing Systems* 30, 2017.
- [31] Zheng Lian BL, and Jianhua Tao, “SMIN: Semi-supervised Multi-modal Interaction Network for Conversational Emotion Recognition,” *IEEE Transactions on Affective Computing*, 2022, 1-15.
- [32] Hazarika D, et al., “ICON: Interactive Conversational Memory Network for Multimodal Emotion Detection,” *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.
- [33] Busso C, et al., “IEMOCAP: Interactive emotional dyadic motion capture database,” *Journal of Language Resources and Evaluation*, 2008, 42(4), 335-359.