



# Enhancing Modal Fusion by Alignment and Label Matching for Multimodal Emotion Recognition

Qifei Li, Yingming Gao, Yuhua Wen, Cong Wang, Ya Li

School of Artificial Intelligence, Beijing University of Posts and Telecommunications, China

{liqifei, yingming.gao, 2020212213, congwang, yli01}@bupt.edu.cn

## Abstract

To address the limitation in multimodal emotion recognition (MER) performance arising from inter-modal information fusion, we propose a novel MER framework based on multi-task learning where fusion occurs after alignment, called Foal-Net. The framework is designed to enhance the effectiveness of modality fusion and includes two auxiliary tasks: audio-video emotion alignment (AVEL) and cross-modal emotion label matching (MEM). First, AVEL achieves alignment of emotional information in audio-video representations through contrastive learning. Then, a modal fusion network integrates the aligned features. Meanwhile, MEM assesses whether the emotions of the current sample pair are the same, providing assistance for modal information fusion and guiding the model to focus more on emotional information. The experimental results conducted on IEMOCAP corpus show that Foal-Net outperforms the state-of-the-art methods and emotion alignment is necessary before modal fusion. The code is open-source<sup>1</sup>.

**Index Terms:** multimodal emotion recognition, multitask learning, contrastive learning, cross-attention

## 1. Introduction

Emotion recognition is an important part of human-computer interaction (HCI). In order to improve the interaction experience, it is necessary to fully utilize the information between different modalities to improve the recognition performance of the system [1]. Hence, efficient modal fusion methods is one of the current research hotspots.

In the MER task, there are many challenges that need to be investigated. This challenges include the need for improved feature representation, developing reasonable model structure to fit the feature representation, exploring effective modal fusion methods to realize inter-modal information complementary and improve the performance of MER, and addressing or alleviating the problem of model performance degradation caused by missing modal information.

The commonly used modalities for emotion recognition include audio, video, and text. In recent years, most of researchers have extracted deep representations of pre-trained models to represent modal information [2, 3, 4, 5], such as HuBERT [6], WavLM [7], BERT [8], RoBERT [9], Resnet-FER2013 [10] and MAE [11]. Compared with the conventional features, such as Mel Frequency Cepstrum Coefficient, one-hot representation and Face Action Unit, the deep representations have better performance and generalization. For modeling the modal information, Wang et al. [12] proposed a modality-sensitive MER framework to exploring complementary features. Mitra

et al. [13] utilized TC-GRU which consists of time convolution network and GRU to capture the spatio-temporal properties of the inputs. Maji et al. [14] proposed a cross-model Transformer model to model the information of audio and text modalities and achieved remarkable performance. The cross-attention mechanism is currently the most commonly used for modal fusion [15]. The above studies all made use of cross-attention to fuse the audio and lexical information to realize MER. Lian et al. [16] exploited graph neural networks to fuse three modalities and attained excellent performance. In order to alleviate the effects of modal absence, Lian et al.[16] and Wang et al. [12] employed different strategies individually to simulate scenarios of modality absence during the training process, enhancing the robustness of the emotion recognition model.

Recently, remarkable performance have been achieved in research related to utilizing auxiliary tasks for modality fusion. In many image-text tasks, such as image-text retrieval, visual entailment and visual question answering (VAQ), many researchers always take advantage of auxiliary tasks to assist image-text information to fuse [17, 18, 19, 20, 21]. These auxiliary tasks include image-text matching, image-text contrastive learning, maskd language model, etc. In particular, the authors of ALBEF [17] point out that it is very effective and necessary to perform modal alignment prior to modal fusion. Similarly, in the MER task, Ghosh et al. [5] and Sun et al. [22] have respectively proposed different auxiliary tasks for multimodal emotion recognition, both achieving outstanding performance.

Inspired by the above studies, we propose a novel MER framework based on multitask learning for MER, named Foal-Net, which consists of two unimodal sub-networks, multiple layers of cross-attention and two auxiliary tasks. The outputs of two sub-networks first need to undergo audio-visual emotion alignment, where the audio-visual emotion alignment contrastive loss (AVEL) is computed. Subsequently, they are fed into the cross-attention network for modalities fusion. Simultaneously, we calculate negative samples for each current modality sample based on AVEL's similarity matrix and hard negative technique, forming pairs of positive-negative samples. Next, these sample pairs, along with the initially matched pairs, are used to implement cross-modal emotion label matching (MEM). The MEM serves two purposes. Firstly, it promotes modality fusion, preventing insufficient integration of modal information and recognizing the emotion of the current sample solely based on a single modality. Secondly, it directs the model to focus more on the emotional information within the current input. For modal information representation, we use CLIP embeddings instead of traditional facial expression features. This is because these embeddings not only encompass facial information but also include body language information. We have validated the effectiveness of our proposed method on

<sup>1</sup><https://github.com/ASolitaryMan/Foal-Net>

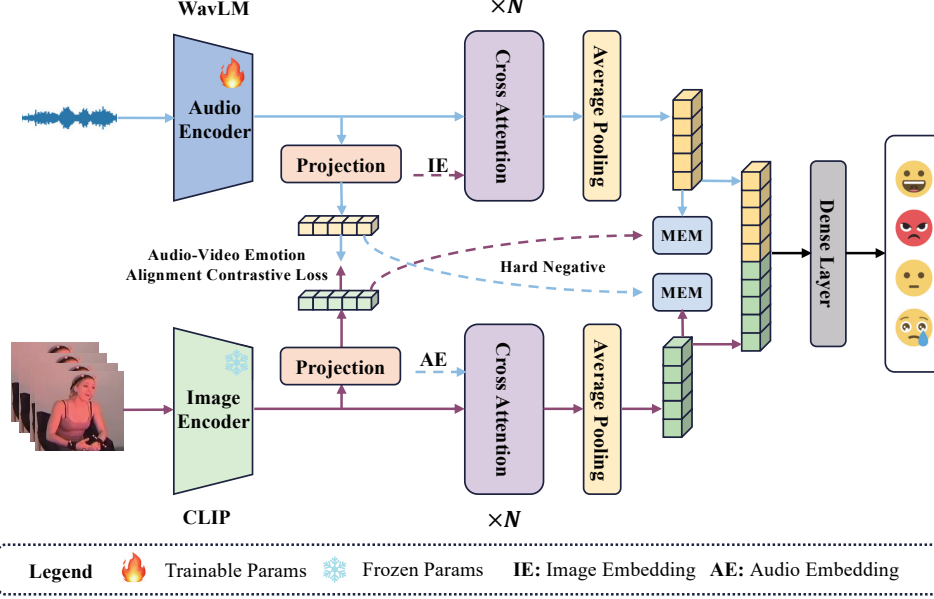


Figure 1: An overview of Foyal-Net architecture, which mainly consists of two auxiliary tasks (AVEL, MEM), and modal fusion module.

the IEMOCAP corpus. Experimental results indicate that the performance of Foyal-Net surpasses other SOTA methods. Our main contributions can be summarized as follows:

- Unlike other existing studies that mainly focus on modal fusion methods, we propose a method that involves alignment before fusion for multimodal emotion recognition, which fills the gap between current researches.
- We introduce an emotion label matching auxiliary task to assist modal information fusion. In addition, it can pay more attention to emotional information in the fusion process.
- The unweighted accuracy (UA) and weighted accuracy (WA) of Foyal-Net are 80.1% and 79.45%, outperforming the other state-of-the-art (SOTA) methods.

## 2. Proposed Method

In this section, we will provide a detailed introduction to the proposed AVEL, MEM auxiliary tasks, and the modal fusion module used.

### 2.1. Audio-Video Emotion Aligning

Inspired by ALBEF [17], in order to achieve better modal fusion for complementary emotional information between modalities in the later stage, we first align the emotional information between the audio and video modalities. In other words, the goal of AVEL is to enhance the similarity between sample pairs with the same emotional category through contrastive learning, while reducing the similarity between sample pairs with different emotional categories.

The input audio-video sample pairs are  $\{X_i^a, X_i^v\}$ , where  $i \in [0, N]$ ,  $N$  is batch size and  $a, v$  represent audio, video respectively. First, Foyal-Net leverages WavLM and CLIP to extract audio embeddings  $Z^a \in \mathbb{R}^{N \times T \times D_a}$  and video embeddings  $Z^v \in \mathbb{R}^{N \times F \times D_v}$ , where  $T, F$  represent the number of frames in the audio and video, respectively,  $D_a, D_v$  are the dimension of audio and video embeddings. Then, we conduct average pooling along with time dimension of  $Z^a$  and  $Z^v$  to ob-




Figure 2: The example for AVEL. The horizontal and vertical axis represents the video and audio modality respectively. As indicated in the subplot on the left, the value is 1 when their emotional labels are same; otherwise, it is 0. CL denotes contrastive loss.

tain  $\bar{Z}^a \in \mathbb{R}^{N \times D_a}$  and  $\bar{Z}^v \in \mathbb{R}^{N \times D_v}$ . The  $\bar{Z}^a$  and  $\bar{Z}^v$  are fed into projection blocks, which are designed to map their inputs to feature vectors of the same dimension, enabling the calculation of inter-modal similarity matrices. These operations are measured as:

$$E^a = MLP_a(\bar{Z}^a); E^v = MLP_v(\bar{Z}^v) \quad (1)$$

$$C^{a2v} = \epsilon \times (E^a \cdot E^{vT}); C^{v2a} = \epsilon \times (E^v \cdot E^{aT}) \quad (2)$$

where the  $MLP_a(\cdot)$  and  $MLP_v(\cdot)$  are projection modules, which consist of two linear layers.  $E^a \in \mathbb{R}^{N \times D}$  and  $E^v \in \mathbb{R}^{N \times D}$  denote the scaled audio and video feature vectors, the  $\epsilon$  is temperature hyper-parameter. The  $C^{a2v} \in \mathbb{R}^{N \times N}$  and  $C^{v2a} \in \mathbb{R}^{N \times N}$  represent inter-modal similarity matrices.

The labels  $\tilde{C}^e \in \mathbb{R}^{N \times N}$  for the contrastive loss, as shown in Figure 2, are set such that if the true labels are the same for different sample pairs within the same batch, the corresponding ground truth values are set to 1. Otherwise, these ground truth values are set to 0. Finally, the inter-modal emotion alignment

loss is defined as:

$$L_a = -\frac{1}{2} \left( \frac{\sum_{i=1}^N \sum_{j=1}^N (\sigma(C_{i,j}^{a2v}) \tilde{C}_{i,j}^e + \sigma(C_{i,j}^{v2a}) \tilde{C}_{i,j}^e)}{N} \right) \quad (3)$$

where the  $\sigma$  is log-softmax function.

## 2.2. Audio-Video Emotion Matching

The research by Sun et al. [22] indicates that during the training process, for some simple samples, the model can identify emotions using only the information from a single modality. This may result in insufficient fusion of modal information. Therefore, errors are more likely to occur when the model predicts complex samples.

Hence, we propose the MEM auxiliary task to alleviate this issue. MEM is a binary classification task designed to enable the model to fully leverage information from both modalities to determine whether the emotional labels for the current input sample pair are consistent. We adopt the hard negative technique to generate the negative samples required for binary classification. Hard negative refers to identifying samples in other modalities that have inconsistent emotional information with the current modality sample but the highest similarity. First, we will set the value of  $C^{a2v}$  and  $C^{v2a}$  to negative infinity at the corresponding position where  $\tilde{C}^e$  has a value of 1. Then, the steps for finding the most similar negative sample from other modalities for the current sample are as follows:

$$id_i^{a2v} = \operatorname{argmax}(\operatorname{softmax}(C_i^{a2v})); i \in [0, N] \quad (4)$$

$$id_i^{v2a} = \operatorname{argmax}(\operatorname{softmax}(C_i^{v2a})); i \in [0, N] \quad (5)$$

$$Z_{neg}^a = Z^a[id_i^{v2a}]; Z_{neg}^v = Z^v[id_i^{a2v}] \quad (6)$$

where  $id_i^{a2v}$  and  $Z_{neg}^v$  respectively represent the index and CLIP embeddings of samples in the video modality. They are selected to serve as negative samples for the audio modality. Similarly,  $id_i^{v2a}$  and  $Z_{neg}^a$  represent the index and WavLM embedding of samples from the audio modality, serving as negative samples for the video modality. Then, they are fed into fusion network for MEM. The operations are as follows:

$$M_p^a = \delta(f_a(Z^a, Z^v)); M_n^a = \delta(f_a(Z^a, Z_{neg}^v)) \quad (7)$$

$$M_p^v = \delta(f_v(Z^v, Z^a)); M_n^v = \delta(f_v(Z^v, Z_{neg}^a)) \quad (8)$$

$$L_m = \frac{1}{2} \sum_k^{k \in \{a,v\}} CE(\operatorname{cat}(M_p^k, M_n^k), \tilde{M}_e) \quad (9)$$

where  $f_a(\cdot)$  and  $f_v(\cdot)$  denote fusion networks for audio and video modalities respectively. The  $\delta$  means average pooling along with time dimension. The  $CE$  represents *Cross Entropy* loss.  $M_p^k$  and  $M_n^k$  ( $k \in \{a, v\}$ ) denote the output of paired positive and unpaired negative samples from the fusion network, respectively.  $\tilde{M}_e \in \mathbb{R}^{2N \times 1}$  represents the labels used for MEM, where the first half takes the value of 1, and the second half takes the value of 0.

Finally, the loss function of Foal-Net is represented as Equation 10, where  $L_{ce}$  is the emotion classification loss, and the value of  $\lambda$  is 0.01.

$$L_{Total} = L_{ce} + L_a + \lambda L_m \quad (10)$$

## 2.3. Modal Fusion Module

The fusion module we used is based on multi-head cross-attention [15]. Here, we use the representation of the speech modality as the query to illustrate the calculation method of the cross-attention mechanism. The calculation process of multi-head cross-attention is as follows:

$$F_{out}^a = LN(MH(Q^a, K^v, V^v) + F^a) \quad (11)$$

$$MH(Q^a, K^v, V^v) = C(head_1, \dots, head_h) \quad (12)$$

$$\begin{aligned} head_h &= \operatorname{Attention}(Q_h^a, K_h^v, V_h^v) \\ &= \operatorname{softmax}\left(\frac{Q_h^a (K_h^v)^T}{\sqrt{d_k}}\right) (V_h^v) \end{aligned} \quad (13)$$

$$Q_h^a = F^a W_h^Q + b_h^Q \quad (14)$$

$$K_h^v = F^v W_h^K + b_h^K \quad (15)$$

$$V_h^v = F^v W_h^V + b_h^V \quad (16)$$

where  $F^a$  and  $F^v$  denotes the embeddings from pretrained audio and video models respectively.  $LN$ ,  $MH$  and  $C$  mean Layer Normalization, Multi-Head and Concatenate. The  $h$  represents the  $h$ th cross-attention. When calculating  $F_{out}^v$ , the method remains consistent with  $F_{out}^a$ , but the query is now based on the representation of the video modality. In this paper, we utilize two multi-head cross-attention layers for modal fusion, with each layer consisting of four heads. The  $F_{out}^a$  and  $F_{out}^v$  will undergo an average pooling layer along with time dimension, and then be concatenated together for MER.

## 3. Experiments and Results

### 3.1. Experimental Setup

**Dataset** The IEMOCAP [23] multimodal corpus is one of the most well-known databases for emotion recognition, encompassing data from three modalities: audio, video, and text. In total, it comprises 5 sessions, each including one male and one female speaker. To be consistent and compare with previous studies, we conduct experiments with 5,531 audio utterances of four emotion categories *happy* (*happy & excited*, 1,636), *angry* (1,103), *sad* (1,084) and *neutral* (1,708). We perform five-fold cross-validation using a leave-one-session-out strategy on the corpus to evaluate the effectiveness of our proposed method. The weighted accuracy (WA) and unweighted accuracy (UA) are used as metrics in line with previous methods.

**Experimental Details** The feature dimensions of the WavLM<sup>2</sup> and CLIP<sup>3</sup> image encoders are 1024 and 768, respectively. The Projection module has 512 neurons, with a dropout rate of 0.5. During training, the batch size is set to 64, the learning rate is a constant 1e-4, and the optimizer is AdamW. The dropout rate for cross-attention is 0.1. The input for the audio modality consists of 6 seconds of speech with a sampling rate

<sup>2</sup><https://huggingface.co/microsoft/wavlm-large>

<sup>3</sup><https://huggingface.co/openai/clip-vit-large-patch14>

Table 1: Performance comparison between the proposed model and other SOTA methods on IEMOCAP corpus. The A means Audio, V denotes Video.

Methods	Type	Year	UA(%)	WA(%)
FTSLLM [24]	A+T	2023	78.50	77.70
MSMSER [12]	A+T	2023	76.40	75.20
ATMF [22]	A+T	2023	79.71	78.42
RMER [25]	A+T	2023	77.00	76.00
MTTX [26]	A+T	2023	75.00	74.50
CRNN-SAN [14]	A+T	2023	79.95	78.82
BAM [27]	A+T	2023	77.00	75.50
MSER [28]	A+T	2024	76.56	77.20
MMAN [29]	A+V	2020	-	73.94
AM-FBP [30]	A+V	2021	75.49	-
MCWSA-CMHA [31]	A+V	2022	78.90	-
GNet [16]	A+V	2023	78.36	-
<b>Foal-Net</b>	<b>A+V</b>	<b>-</b>	<b>80.10</b>	<b>79.45</b>

of 16 kHz. Due to the simultaneous appearance of two speakers in the video recording, it is necessary to separate the speakers in the frame. In the end, the input for the video modality consists of 180 individual images. Any questions regarding the parameters in the method we proposed, please refer to the github repository.

### 3.2. Performance Comparison with Previous Methods

The effectiveness of the method we proposed can be underscored by contrasting it with the latest significant findings derived from the IEMOCAP corpus, as show in Table 1. It demonstrates that the best UA (80.10%) and WA (79.45%) are achieved by Foal-Net. Research on facial expression recognition in the IEMOCAP corpus is relatively scarce, primarily due to two reasons. Firstly, the video frames are relatively blurry, making it challenging to extract facial expression features. Secondly, the issue of speakers appearing in the same frame in the videos complicates the data processing. However, based on the experimental results, the performance of the audio-video combination is not inferior to that of the audio-text combination. This suggests that besides facial information, other details in the images also contribute to emotion recognition. Meanwhile, it highlights the outstanding performance of the Foal-Net.

### 3.3. Experiments and Analysis

We conducted a series of ablation experiments to validate the effectiveness of the features we used, the proposed auxiliary tasks, and the final model. The baseline is the Foal-Net without two auxiliary tasks. To validate the effectiveness of global image encoding for emotion recognition compared to local facial information encoding, we extracted deep representations from three pretrained models (EmoNet [32], Resnet-FER2013 [10], SENet-FER2013 [10]) for facial expression recognition and utilized them for emotion recognition. As shown in Table 2, their performance is significantly lower than that of the CLIP model’s features. This indicates that, in addition to facial expressions, body movements and other information in the images also contribute to emotion recognition. Furthermore, another reason for the superior performance of CLIP features is the inclusion of semantic information.

In addition, we observe that after incorporating the AVEL

Table 2: The results of ablation experiments related to feature input, auxiliary tasks, and modal fusion conducted on the IEMOCAP corpus.

Type	Methods	UA(%)	WA(%)
V	<b>Face-only</b>		
	EmoNet	37.22	36.50
	SENet-FER2013	36.18	35.53
	Resnet-FER2013	40.74	40.16
A	<b>All Image</b>		
	CLIP	47.90	46.61
A+V	WavLM	77.68	77.25
	Baseline	77.85	77.31
	Baseline+AVEL	79.31	78.20
	Baseline+MEM	78.10	76.94
	<b>Baseline+AVEL+MEM</b>	<b>80.10</b>	<b>79.45</b>

task, UA and WA improved by 2.2% and 1.36%, respectively, compared to the baseline. This indicates the crucial necessity of cross-modal emotion alignment before modal fusion. On the contrary, the performance improvement is not very significant when introducing the MEM task alone. This is because the use of hard negative for finding negative samples heavily relies on the similarity matrix from the AVEL task. When the similarity matrix is not optimized, it cannot bring positive benefits to the MEM task. This also indirectly confirms the necessity of alignment before fusion. When we introduce both the AVEL and MEM tasks simultaneously, the model’s performance is further improved compared to using AVEL or MEM alone. This demonstrates that with the assistance of AVEL, MEM can facilitate the fusion of modal information and enhance the model’s capability for emotion recognition.

## 4. Conclusions

In this paper, we propose a novel framework called Foal-Net for MER, which includes two auxiliary tasks: AVEL and MEM. Our research on the AVEL task has demonstrated the effectiveness and necessity of initially aligning emotional information across modalities before integrating inter-modality information. The MEM task can guide modality fusion and make fusion module focus more on emotional information during the fusion process with the assistance of the AVEL task. Under the influence of both tasks, Foal-Net achieves SOTA performance. Moreover, we show that the embeddings of CLIP outperform facial expression features in MER. In the future work, we will optimize the MEM task as its performance relies on AVEL and validate the universality of our method by substituting the video modality with the text modality.

## 5. Acknowledgements

The work was supported by the National Natural Science Foundation of China (No. 62271083), the Key Project of the National Language Commission (No. ZDI145-81), the Fundamental Research Funds for the Central Universities (No. 2023RC13, No. 2023RC73), BUPT Excellent Ph.D. Students Foundation (No. 2023116), and partly supported by the Major Program of the National Social Science Fund of China (13&ZD189).

## 6. References

- [1] M. K. Tellamekala, S. Amiriparian, B. W. Schuller, E. André, T. Giesbrecht, and M. Valstar, "Cold fusion: Calibrated and ordinal latent distribution fusion for uncertainty-aware multimodal emotion recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [2] J. Tavornor, M. Perez, and E. Mower Provost, "Episodic Memory For Domain-Adaptable, Robust Speech Emotion Recognition," in *Proc. INTERSPEECH 2023*, 2023, pp. 656–660.
- [3] M. Tran, Y. Yin, and M. Soleymani, "Personalized Adaptation with Pre-trained Speech Encoders for Continuous Emotion Recognition," in *Proc. INTERSPEECH 2023*, 2023, pp. 636–640.
- [4] Z. Zhao, T. Gao, H. Wang, and B. W. Schuller, "SWRR: Feature Map Classifier Based on Sliding Window Attention and High-Response Feature Reuse for Multimodal Emotion Recognition," in *Proc. INTERSPEECH 2023*, 2023, pp. 2433–2437.
- [5] S. Ghosh, U. Tyagi, S. Ramaneswaran, H. Srivastava, and D. Manocha, "MMER: Multimodal Multi-task Learning for Speech Emotion Recognition," in *Proc. INTERSPEECH 2023*, 2023, pp. 1209–1213.
- [6] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [7] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [8] J. D. M.-W. C. Kenton and L. K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [9] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [10] Y. Zhong, S. Qiu, X. Luo, Z. Meng, and J. Liu, "Facial expression recognition based on optimized resnet," in *2020 2nd World Symposium on Artificial Intelligence (WSAI)*. IEEE, 2020, pp. 84–91.
- [11] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.
- [12] S. Wang, Y. Ma, and Y. Ding, "Exploring complementary features in multi-modal speech emotion recognition," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [13] V. Mitra, V. Kowtha, H.-Y. S. Chien, E. Azemi, and C. Avendano, "Pre-trained model representations and their robustness against noise for speech emotion analysis," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [14] B. Maji, M. Swain, R. Guha, and A. Routray, "Multimodal emotion recognition based on deep temporal features using cross-modal transformer and self-attention," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [16] Z. Lian, L. Chen, L. Sun, B. Liu, and J. Tao, "Gcnet: graph completion network for incomplete multimodal learning in conversation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [17] J. Li, R. Selvaraju, A. Gotmare, S. Joty, C. Xiong, and S. C. H. Hoi, "Align before fuse: Vision and language representation learning with momentum distillation," *Advances in Neural Information Processing Systems*, vol. 34, pp. 9694–9705, 2021.
- [18] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8748–8763.
- [19] J. Li, D. Li, C. Xiong, and S. Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *International Conference on Machine Learning*. PMLR, 2022, pp. 12 888–12 900.
- [20] W. Wang, H. Bao, L. Dong, J. Bjorck, Z. Peng, Q. Liu, K. Aggarwal, O. K. Mohammed, S. Singhal, S. Som *et al.*, "Image as a foreign language: Beit pretraining for vision and vision-language tasks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 19 175–19 186.
- [21] H. Bao, W. Wang, L. Dong, Q. Liu, O. K. Mohammed, K. Aggarwal, S. Som, S. Piao, and F. Wei, "Vlmo: Unified vision-language pre-training with mixture-of-modality-experts," *Advances in Neural Information Processing Systems*, vol. 35, pp. 32 897–32 912, 2022.
- [22] D. Sun, Y. He, and J. Han, "Using auxiliary tasks in multimodal fusion of wav2vec 2.0 and bert for multimodal emotion recognition," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [23] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language Resources and Evaluation*, vol. 42, pp. 335–359, 2008.
- [24] W. Yang, S. Fukayama, P. Heracleous, and J. Ogata, "Exploiting fine-tuning of self-supervised learning models for improving bi-modal sentiment analysis and emotion recognition," in *INTER-SPEECH*, 2022, pp. 1998–2002.
- [25] B. Lin and L. Wang, "Robust multi-modal speech emotion recognition with asr error adaptation," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [26] J. He, M. Wu, M. Li, X. Zhu, and F. Ye, "Multilevel transformer for multimodal emotion recognition," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [27] Z. Zhao, Y. Wang, and Y. Wang, "Knowledge-aware bayesian co-attention for multimodal emotion recognition," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [28] M. Khan, W. Gueaieb, A. El Saddik, and S. Kwon, "Mser: Multimodal speech emotion recognition using cross-attention with deep fusion," *Expert Systems with Applications*, vol. 245, p. 122946, 2024.
- [29] Z. Pan, Z. Luo, J. Yang, and H. Li, "Multi-Modal Attention for Speech Emotion Recognition," in *Proc. Interspeech 2020*, 2020, pp. 364–368.
- [30] H. Zhou, J. Du, Y. Zhang, Q. Wang, Q.-F. Liu, and C.-H. Lee, "Information fusion in attention networks using adaptive and multi-level factorized bilinear pooling for audio-visual emotion recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2617–2629, 2021.
- [31] J. Zheng, S. Zhang, Z. Wang, X. Wang, and Z. Zeng, "Multi-channel weight-sharing autoencoder based on cascade multi-head attention for multimodal emotion recognition," *IEEE Transactions on Multimedia*, 2022.
- [32] S. E. Kahou, X. Bouthillier, P. Lamblin, C. Gulcehre, V. Michalski, K. Konda, S. Jean, P. Froumenty, Y. Dauphin, N. Boulanger-Lewandowski *et al.*, "Emonets: Multimodal deep learning approaches for emotion recognition in video," *Journal on Multimodal User Interfaces*, vol. 10, pp. 99–111, 2016.