

Recursive Joint Cross-Modal Attention for Multimodal Fusion in Dimensional Emotion Recognition

R. Gnana Praveen Jahangir Alam

Computer Research Institute of Montreal (CRIM), Canada

gnana-praveen.rajasekhar@crim.ca jahangir.alam@crim.ca

Abstract

Though multimodal emotion recognition has achieved significant progress over recent years, the potential of rich synergic relationships across the modalities is not fully exploited. In this paper, we introduce Recursive Joint Cross-Modal Attention (RJCMA) to effectively capture both intra- and inter-modal relationships across audio, visual, and text modalities for dimensional emotion recognition. In particular, we compute the attention weights based on cross-correlation between the joint audio-visual-text feature representations and the feature representations of individual modalities to simultaneously capture intra- and inter-modal relationships across the modalities. The attended features of the individual modalities are again fed as input to the fusion model in a recursive mechanism to obtain more refined feature representations. We have also explored Temporal Convolutional Networks (TCNs) to improve the temporal modeling of the feature representations of individual modalities. Extensive experiments are conducted to evaluate the performance of the proposed fusion model on the challenging Affwild2 dataset. By effectively capturing the synergic intra- and inter-modal relationships across audio, visual, and text modalities, the proposed fusion model achieves a Concordance Correlation Coefficient (CCC) of 0.585 (0.542) and 0.674 (0.619) for valence and arousal, respectively, on the validation set (test set). This shows a significant improvement over the baseline of 0.240 (0.211) and 0.200 (0.191) for valence and arousal, respectively, in the validation set (test set), achieving second place in the valence-arousal challenge of the 6th Affective Behavior Analysis in-the-Wild (ABAW) competition. The code is available on GitHub: <https://github.com/praveena2j/RJCMA>.

1. Introduction

Emotion recognition is a challenging problem due to the diverse nature of expressions associated with emotional states

across individuals and cultures [1]. It has a wide range of applications in various fields such as health care (assessing pain, fatigue, depression) [2, 3], autonomous driving (assessing driver emotional states) [4], robotics (for realistic human-computer interaction) [5], etc. Emotion recognition is often explored in the literature as a classification problem of categorizing emotions into seven classes: anger, disgust, fear, happiness, sadness, surprise, and contempt [6]. Recently, compound expression datasets have also been introduced to capture human emotions beyond seven classes [7]. However, they also do not capture the diverse range of emotions expressed by humans. Therefore, a dimensional model of emotions has been introduced to capture the diverse range of emotions, predominantly using valence and arousal. Valence denotes the range of emotions from very sad (negative) to very happy (positive), whereas arousal represents the intensities of emotions from being very passive (sleepiness) to very active (high excitement) [8]. Dimensional Emotion Recognition (DER) is more challenging than categorical emotion recognition, as they are highly prone to label noise due to the complex process of obtaining dimensional annotations.

Multimodal learning has been recently gaining a lot of attention as it can offer rich complementary information across multiple modalities, which can play a crucial role in outperforming unimodal approaches [9]. Human emotions are communicated in complex ways through various modalities such as face, voice, and language. Multimodal emotion recognition aims to effectively fuse audio, visual, and text by capturing the rich intra- and inter-modal complementary relationships across the modalities. Early approaches of multimodal fusion for DER either rely on Long Short Term Memory (LSTM)-based fusion [10, 11] or early feature concatenation [12, 13]. With the advent of transformers [14], attention models using multimodal transformers have attained much interest to effectively combine multiple modalities for DER [15–17]. Recently, cross-modal attention was found to be quite promising in capturing the rich inter-modal complementary relationships across the modalities, which has been successfully applied for several appli-

cations such as action localization [18], emotion recognition [19], and person verification [20]. Praveen et al. [21, 22] explored joint cross-attention by introducing joint feature representation in the cross-attentional framework and achieved significant improvement over vanilla cross-attention [19]. They have further improved the performance of the model by introducing recursive fusion, and LSTMs for temporal modeling of individual feature representations as well as the audio-visual representations [23]. Recursive fusion of the cross-attention models has also been successfully explored for other audio-visual tasks such as event localization [24] and person verification [25]. However, most of the existing cross-attention based models are focused on audio-visual fusion for DER.

In this work, we have investigated the prospect of efficiently capturing the synergic intra- and inter-modal relationships across audio, visual, and text modalities for DER. By deploying cross-correlation between the joint audio-visual-text feature representation and feature representations of individual modalities, we can simultaneously capture intra- and inter-modal relationships across the modalities. Inspired by the performance of recursive attention models [23, 24], we have also incorporated a recursive fusion of audio, visual, and text modalities in the context of joint cross-attentional fusion to obtain more refined feature representations. The major contributions of the paper can be summarized as follows.

- Joint cross-modal attention is explored among audio, visual, and text modalities using a joint audio-visual-text feature representation to simultaneously capture intra- and inter-modal relationships across the modalities.
- Recursive fusion is used to further improve the feature representations of individual modalities. TCNs are also used to improve the temporal modeling of individual feature representations.
- Extensive experiments are conducted to evaluate the robustness of the proposed approach on the challenging AFWild2 dataset.

2. Related Work

2.1. Multimodal Emotion Recognition

One of the early approaches using deep learning architectures for DER was proposed by Tzirakis et al. [10], where they explored 1D Convolutional Neural Networks (CNN) for audio and Resnet-50 [26] for visual modality. The deep features are then fused using LSTMs for estimating predictions of valence and arousal. With the advancement of 3D CNN models, Kuhnke et al. [12] showed performance improvement using R3D [27] for visual modality and Resnet-18 [26] for audio modality with simple feature concatenation. Kollias et al. [28–30] explored DER along with other tasks of classification of expressions and action units in a

unified framework. Another widely explored line of research for DER is based on knowledge distillation (KD) [31]. Schoneveld et al. [11] explored KD for obtaining robust visual representations, while Wang et al. [32] and Deng et al. [33] attempted to leverage KD to deal with label uncertainties. Recently, KD is also explored along with the paradigm of Learning Under Privileged Information (LUPI) to efficiently exploit multiple modalities for DER [34]. Li et al. [35] proposed Decoupled Multimodal Distillation (DMD) to mitigate the issues of multimodal heterogeneities by dynamically distilling the modality-relevant information across the modalities. Although these methods have shown promising performance by exploiting multiple modalities, they do not focus on capturing the synergic relationships pertaining to intra- and inter-modal relationships across the modalities.

2.2. Attention Models for Multimodal Emotion Recognition

Inspired by the performance of transformers [14], several approaches have been proposed to investigate the potential of transformers for DER. Most of the existing works explored transformers to encode the concatenated version of the feature representations of individual modalities [16, 17, 36, 37]. Tran et al. [38] showed that fine-tuning transformers with cross-modal attention trained on a large-scale voxceleb2 [39] dataset helps to improve the performance of multimodal emotion recognition. Huang et al. [40] explored multimodal transformers along with self-attention modules, where audio modality is used to attend to visual modality to produce robust multimodal feature representations. Parthasarathy et al. [15] further extended their idea by employing cross-modal attention in a bidirectional fashion, where audio modality is used to attend to visual modality and vice-versa. Karas et al. [41] provided a comprehensive evaluation of the fusion models based on self-attention, cross-attention and LSTMs for DER. Zhang et al. [42] proposed leader-follower attention, where audio, visual, and text modalities are combined to obtain the attended feature representations using the modality-wise attention scores, which is further concatenated to visual features for final predictions.

Praveen et al. [19] explored cross-attention based on cross-correlation across the feature representations of audio and visual modalities. They further extended their approach by employing joint audio-visual feature representation in the cross-attentional framework and showed significant performance improvement [21]. Praveen et al. [23] improved the performance by introducing recursive fusion and LSTMs for temporal modeling of individual modalities. In this work, we further extend the idea of [23] by incorporating text modality, and TCNs to effectively capture the intra-modal relationships. The proposed approach primar-

ily differs from [23] in three aspects : (1) In this work, we employ text modality in the framework of recursive joint cross-attention in addition to audio and visual modalities, whereas [23] uses only audio and visual modalities. (2) In [23], LSTMs are used, whereas we have deployed TCNs as they are found to be effective in improving the temporal modeling of individual modalities. (3) In [23], R3D [27] is used for visual backbone and Resnet 18 for audio modality, whereas we have used Resnet-50 fine-tuned on FER+ [43] for visual and VGG [44] pretrained on Audioset for audio modality.

3. Proposed Approach

3.1. Visual Network

Facial expressions play a significant role in conveying the emotional state of a person. In videos, spatial information provide the semantic regions of face pertinent to the expressions, whereas temporal dynamics convey the evolution of expressions in videos across the frames. Therefore, effectively modeling the spatial and temporal dynamics of facial expressions in videos is crucial to obtain robust visual feature representations. Several approaches have been explored using 2D CNNs in conjunction with LSTMs, where 2D CNNs are used to encode the spatial information and LSTMs for temporal dynamics of facial expressions [45, 46]. With the advent of 3D CNN models [27], they are successfully explored for DER [12, 21, 22] and showed improvement over 2D CNNs with LSTMs. It has also been shown that 3DCNNs in combination with LSTMs are effective in capturing the temporal dynamics [23], where 3DCNNs are efficient in capturing short-term dynamics and LSTMs are robust in modeling long-term dynamics. Recently, TCNs are found to be promising to effectively capture the temporal dynamics for DER [36, 42]. In this work, we have used Resnet-50 [26] pretrained on MS-CELEB-1M dataset [47], which is further fine-tuned on FER+ [43] dataset similar to that of [42]. We have further used TCNs to effectively capture the temporal dynamics of facial expressions.

3.2. Audio Network

Speech-based emotion recognition is another promising research area due to the rich emotion-relevant information in the vocal expressions. With the advancement of deep learning models, vocal expressions are encoded using 1D CNNs with raw speech signals [10] or 2D CNNs with spectrograms [11, 12]. Some of the works also explored deep features in combination with conventional hand-crafted features to encode the vocal expressions [16, 48]. Recently, spectrograms have been widely explored as they are found to be efficient in capturing the affective states of the vocal expressions [21, 42]. Therefore, we have also explored

spectrograms with 2D CNNs to encode the vocal expressions. Specifically, we have used VGG-Net architecture pretrained on large-scale audioset dataset [49]. Similar to visual modality, we have also used TCNs to encode the temporal dynamics of frame-level vocal embeddings.

3.3. Text Network

Text modality is another predominantly explored modality for emotion detection, which carries semantic emotion-relevant information in the text data [50]. Effectively leveraging the textual data can boost the performance of multi-modal fusion as they can offer significant emotion-relevant information and complement audio and visual modalities. Based on transformers, BERT features are predominantly explored text encoders for emotion recognition in the literature [51]. Therefore, we also used BERT as text encoder, followed by TCNs to encode the temporal information across the word embeddings.

3.4. Recursive Joint Cross-Modal Attention

Given the input video sub-sequence S of K frames, the audio, visual and text data are preprocessed and fed to the corresponding encoders, followed by TCNs to obtain the feature representations of the respective modalities as shown in Fig 1. The feature representations of the audio, visual and text modalities are denoted by $\mathbf{X}_a = \{\mathbf{x}_a^1, \mathbf{x}_a^2, \dots, \mathbf{x}_a^K\} \in \mathbb{R}^{d_a \times K}$, $\mathbf{X}_v = \{\mathbf{x}_v^1, \mathbf{x}_v^2, \dots, \mathbf{x}_v^K\} \in \mathbb{R}^{d_v \times K}$ and $\mathbf{X}_t = \{\mathbf{x}_t^1, \mathbf{x}_t^2, \dots, \mathbf{x}_t^K\} \in \mathbb{R}^{d_t \times K}$, respectively, where d_a , d_v , and d_t are dimensions of audio, visual and text features. \mathbf{x}_a^k , \mathbf{x}_v^k and \mathbf{x}_t^k represents the feature vectors of the individual frames of audio, visual, and text modalities.

Given the audio (A), Visual (V), and Text (T) feature representations, \mathbf{X}_a , \mathbf{X}_v , and \mathbf{X}_t the joint feature representation (\mathbf{J}) of audio, visual, and text modalities is obtained by concatenating the feature vectors of all modalities, followed by fully connected layer as

$$\mathbf{J} = FC([\mathbf{X}_a; \mathbf{X}_v; \mathbf{X}_t]) \in \mathbb{R}^{d \times K} \quad (1)$$

where $d = d_a + d_v + d_t$ denotes the dimensionality of \mathbf{J} and FC denote fully connected layer.

Now \mathbf{J} is fed to the joint cross-attentional framework of the respective modalities as shown in Fig 2 to attend to the feature representations of individual modalities. This helps to simultaneously encode the intra- and inter-modal relationships within the same modalities as well as across the modalities in obtaining the attention weights. The cross-correlation between the \mathbf{J} and \mathbf{X}_a are obtained as joint cross-correlation matrix \mathbf{C}_a , which is given by

$$\mathbf{C}_a = \tanh\left(\frac{\mathbf{X}_a^\top \mathbf{W}_{ja} \mathbf{J}}{\sqrt{d}}\right) \quad (2)$$

where $\mathbf{W}_{ja} \in \mathbb{R}^{d_a \times d}$ represents learnable weight matrix across the \mathbf{X}_a and \mathbf{J} .

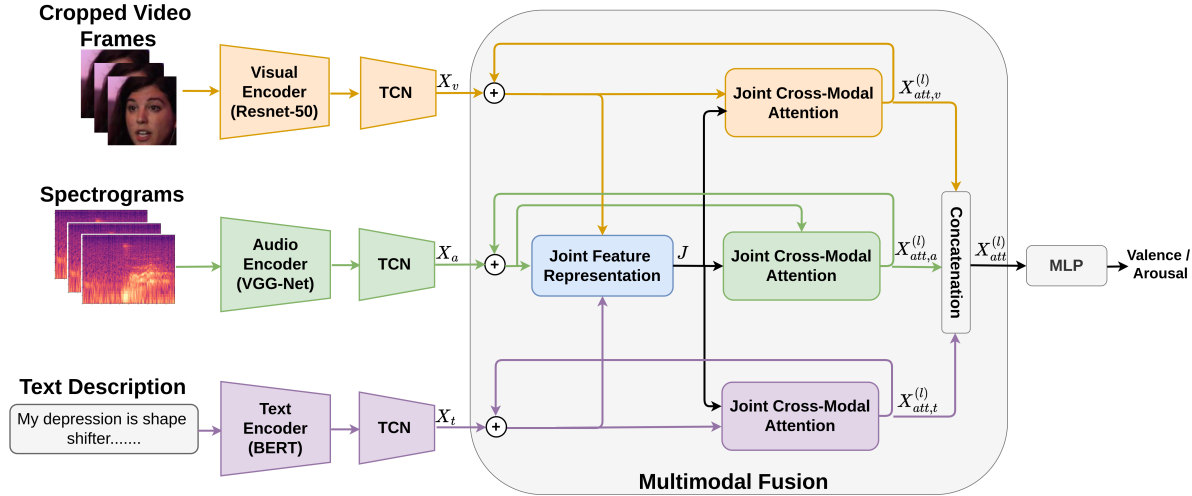


Figure 1. Illustration of the proposed framework with the recursive joint cross-modal attention. Best viewed in color.

Similarly, the cross-correlation matrices of other two modalities C_v and C_t are obtained as

$$C_v = \tanh \left(\frac{\mathbf{X}_v^\top \mathbf{W}_{jv} \mathbf{J}}{\sqrt{d}} \right) \quad (3)$$

$$C_t = \tanh \left(\frac{\mathbf{X}_t^\top \mathbf{W}_{jt} \mathbf{J}}{\sqrt{d}} \right) \quad (4)$$

where $\mathbf{W}_{jv} \in \mathbb{R}^{d_v \times d}$, $\mathbf{W}_{jt} \in \mathbb{R}^{d_t \times d}$ represents learnable weight matrices for visual and text modalities, respectively.

The obtained joint cross-correlation matrices of the individual modalities are used to compute the attention weights, thereby capturing the semantic relevance of both across and within the same modalities. Higher correlation coefficient of the joint cross-correlation matrices denote higher semantic relevance pertinent to the intra- and inter-modal relationships of the corresponding feature vectors. Now the joint cross-correlation matrices are used to compute the attention maps of the individual modalities. For the audio modality, the joint correlation matrix C_a and the corresponding audio features \mathbf{X}_a are combined using the learnable weight matrix \mathbf{W}_{ca} , followed by $ReLU$ activation function to compute the attention maps \mathbf{H}_a , which is given by

$$\mathbf{H}_a = ReLU(\mathbf{X}_a \mathbf{W}_{ca} C_a) \quad (5)$$

where $\mathbf{W}_{ca} \in \mathbb{R}^{K \times K}$ denote learnable weight matrix for audio modality. Similarly the attention maps of visual and text modalities are obtained as

$$\mathbf{H}_v = ReLU(\mathbf{X}_v \mathbf{W}_{cv} C_v) \quad (6)$$

$$\mathbf{H}_t = ReLU(\mathbf{X}_t \mathbf{W}_{ct} C_t) \quad (7)$$

where $\mathbf{W}_{cv} \in \mathbb{R}^{K \times K}$, $\mathbf{W}_{ct} \in \mathbb{R}^{K \times K}$ are the learnable weight matrices of visual and text modalities, respectively. Now the attention maps are used to compute the attended features of the individual modalities as:

$$\mathbf{X}_{att,a} = \mathbf{H}_a \mathbf{W}_{ha} + \mathbf{X}_a \quad (8)$$

$$\mathbf{X}_{att,v} = \mathbf{H}_v \mathbf{W}_{hv} + \mathbf{X}_v \quad (9)$$

$$\mathbf{X}_{att,t} = \mathbf{H}_t \mathbf{W}_{ht} + \mathbf{X}_t \quad (10)$$

where $\mathbf{W}_{ha} \in \mathbb{R}^{K \times K}$, $\mathbf{W}_{hv} \in \mathbb{R}^{K \times K}$, and $\mathbf{W}_{ht} \in \mathbb{R}^{K \times K}$ denote the learnable weight matrices for audio, visual and text modalities, respectively.

In order to obtain more refined feature representations, the attended features of each modality are again fed as input to the respective joint cross-modal attention module, which is given by

$$\mathbf{X}_{att,a}^{(l)} = \mathbf{H}_a^{(l)} \mathbf{W}_{ha}^{(l)} + \mathbf{X}_a^{(l-1)} \quad (11)$$

$$\mathbf{X}_{att,v}^{(l)} = \mathbf{H}_v^{(l)} \mathbf{W}_{hv}^{(l)} + \mathbf{X}_v^{(l-1)} \quad (12)$$

$$\mathbf{X}_{att,t}^{(l)} = \mathbf{H}_t^{(l)} \mathbf{W}_{ht}^{(l)} + \mathbf{X}_t^{(l-1)} \quad (13)$$

where $\mathbf{W}_{ha}^{(l)} \in \mathbb{R}^{K \times K}$, $\mathbf{W}_{hv}^{(l)} \in \mathbb{R}^{K \times K}$, and $\mathbf{W}_{ht}^{(l)} \in \mathbb{R}^{K \times K}$ denote the learnable weight matrices of audio, visual and text modalities, respectively, and l refers to the recursive step.

The attended features of the individual modalities after l iterations are concatenated to obtain the multimodal feature representation $\mathbf{X}_{att}^{(l)}$, shown as:

$$\mathbf{X}_{att}^{(l)} = [\mathbf{X}_{att,a}^{(l)}; \mathbf{X}_{att,v}^{(l)}; \mathbf{X}_{att,t}^{(l)}] \quad (14)$$

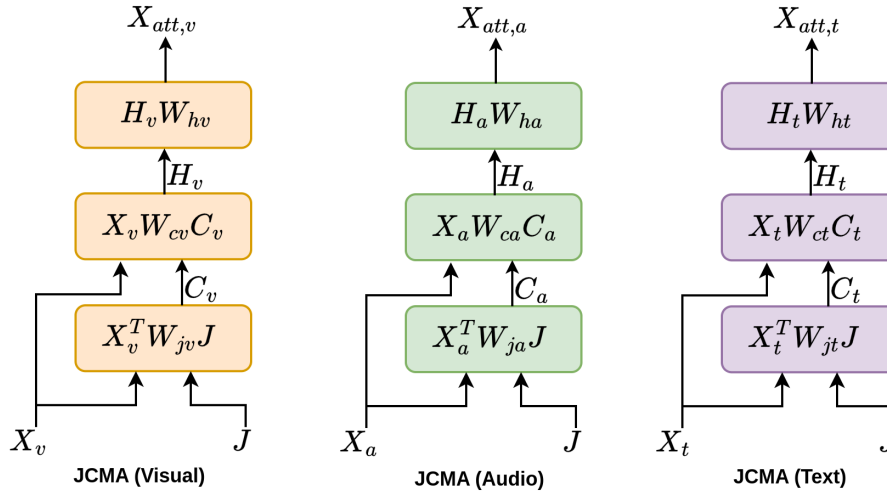


Figure 2. Joint Cross-Modal Attention blocks of audio, visual and text modalities

Finally the multimodal feature representation $X_{att}^{(l)}$ is fed to regression layers (Multi Layer Perceptron) for the final prediction of valence or arousal.

4. Experimental Setup

4.1. Dataset

Affwild is the largest database in the field of affective computing, consisting of 298 videos captured under extremely challenging conditions from YouTube videos [52]. The database is extended substantially to foster the development of robust models for estimating valence and arousal by introducing 260 more videos, resulting in a total of 558 videos with 1.4M frames [53]. The database has been further expanded by adding new videos in the subsequent series of ABAW challenges [54–59]. In the 6th ABAW challenge [60] for the valence-arousal track, the dataset is provided with 594 videos of around 2, 993, 081 frames from 584 subjects. Sixteen of these videos display two subjects, both of which have been annotated. The final annotations are obtained by taking an average of annotations provided by four experts, using a joystick. The annotations for valence and arousal are provided continuously in the range of $[-1, 1]$. The dataset is partitioned into the training, validation, and test sets in a subject-independent manner to ensure that each subject appears exclusively in one partition. The partitioning resulted in 356, 76, and 162 videos for train, validation, and test partitions, respectively.

4.2. Implementation Details

4.2.1 Preprocessing

For the visual modality, we have used the cropped and aligned images provided by the challenge organizers [60].

For the missing faces in the video frames, we have considered black frames (i.e., zero pixels), and the video frames with no annotations of valence and arousal i.e., with annotations of -5 are excluded. The given video sequences are divided into sub-sequences of 300 (i.e., $K = 300$) with a stride of 200 frames and the facial images are resized to 48×48 .

For audio modality, the speech signals are extracted from the corresponding videos with a sampling rate of 16KHz. The log melspectrograms are then obtained using the preprocessing code provided by the Vggish repository¹. To ensure that the audio modality is properly synchronized with the sub-sequences of other modalities, we have used a hop length of $1/fps$ of the raw videos to extract the spectrograms.

For the text modality, the extracted speech signals from audio preprocessing are fed to the pretrained speech recognition model of Vosk toolkit² to obtain the recognized words along with word-level timestamps. Next, a pretrained punctuation restoration and capitalization model³ is used to restore the punctuations of the recognized words, which carries semantic information pertinent to emotional states. Now BERT features are extracted at word-level using a pre-trained BERT model⁴. The word-level features are computed by taking a summation of the last four layers of the BERT model [61]. The recognized words may usually span a larger time window of multiple frames. In order to synchronize the word-level BERT features of text modality with audio and visual modalities, the word-level

¹<https://github.com/harritaylor/torchvggish>

²<https://alphacephei.com/vosk/models/vosk-model-en-us-0.22.zip>

³<https://pypi.org/project/deepmultilingualpunctuation/>

⁴<https://pypi.org/project/pytorch-pretrained-bert/>

text embedding is populated as per the timestamp of each word by reassigning the same word-level feature to all the frames within the time span of the corresponding word.

4.2.2 Training Details

For visual modality, random flipping, and random crop with a size of 40 are used for data augmentation in training, while only center crop is used for validation. For audio and visual features, the input data is normalized in order to have a mean and standard deviation of 0.5. For text modality, the BERT features are normalized to ensure the mean as 0 and standard deviation as 1. Adam optimizer is used with a weight decay of 0.001 and batch size is set to be 12. The models are trained separately for valence and arousal. The maximum number of epochs is set to be 100 and early stopping is employed to avoid over-fitting. The hyperparameters of the initial learning rate and minimum learning rate are set to be $1e - 5$ and $1e - 8$, respectively. In our training strategy, we have deployed a warm-up scheme using *ReduceLROnPlateau* scheduler with patience of 5 and a factor of 0.1 based on the CCC score of the validation partition. It has been shown that gradual training of the backbones of individual modalities along with the fusion model by gradually fine-tuning the layers of the backbones helps to improve the performance of the system [42]. Therefore, we have deployed a similar strategy in our training framework, where three groups of layers for visual (Resnet-50) and audio (VGG) backbones are progressively selected for fine-tuning. Initially at epoch 0, the first group is unfrozen and the learning rate is linearly warmed up to $1e - 5$ within an epoch. Then repetitive warm-up is employed until epoch 5, after which *ReduceLROnPlateau* is used to update the learning rate. The learning rate is gradually dropped with a factor of 0.1 until validation CCC does not improve over 5 consecutive epochs. After which the second group is unfrozen and the learning rate is reset to $1e - 5$, followed by a warm-up scheme with *ReduceLROnPlateau*. The procedure is repeated till all the layers are fine-tuned for audio and visual backbones. Also, note that the best model state dictionary over prior epochs is loaded at the end of each epoch to mitigate the issues of over-fitting. To further control the problem of over-fitting, we have employed cross-validation with 6 folds, where the fold 0 partition is the same as the original partition provided by the organizers [60]. The results obtained from the 6-fold cross-validation are shown in Table 1. In all these experiments, we have used 3 iterations in the fusion model (i.e., $l=3$).

4.2.3 Loss Function

The concordance correlation coefficient (ρ_c) is the widely-used evaluation metric in the literature for DER to measure the level of agreement between the predictions (x) and

Table 1. CCC of the proposed approach on the six folds of cross-validation. Bold indicates the highest scores. Fold 0 is the official validation set.

Validation Set	Valence	Arousal	Mean
Fold 0	0.455	0.652	0.553
Fold 1	0.585	0.674	0.629
Fold 2	0.467	0.631	0.549
Fold 3	0.536	0.647	0.591
Fold 4	0.432	0.629	0.526
Fold 5	0.463	0.615	0.539

ground truth (y) annotations [10]. Let μ_x , and μ_y represent the mean of predictions and ground truth, respectively. Similarly, σ_x^2 and σ_y^2 denote the variance of predictions and ground truth, respectively, then ρ_c between the predictions and ground truth can be obtained as:

$$\rho_c = \frac{2\sigma_{xy}^2}{\sigma_x^2 + \sigma_y^2 + (\mu_x - \mu_y)^2} \quad (15)$$

where σ_{xy}^2 denotes the covariance between predictions and ground truth. Though Mean Square Error (MSE) is the commonly used loss function for regression models, we use the CCC-based loss function as it is a standard loss function in the literature for DER [10, 21], which is given by

$$\mathcal{L} = 1 - \rho_c \quad (16)$$

5. Results and Discussion

5.1. Ablation Study

In order to understand the impact of the recursive mechanism, we have conducted a series of experiments by varying the number of recursions as shown in Table 3. First, we did an experiment with a single recursion, which is the same as joint cross attention [21] for audio, visual, and text modalities. Now we executed multiple experiments by slowly increasing the number of recursions and found that the performance of the system gradually increases with multiple recursions. This shows that the recursive mechanism helps in obtaining more robust feature representations by progressively refining the features. We have achieved the best results at 3 iterations, beyond that the performance of the system declines. We hypothesize that this can be attributed to the fact that though recursive fusion initially helps to improve the performance, it may result in over-fitting with more iterations, resulting in a performance decline on the validation set. A similar trend of performance improvement

Table 2. CCC of the proposed approach compared to state-of-the-art methods for multimodal fusion on the original Affwild2 validation set (fold 0). * indicates that results are presented with the implementation of our experimental setup. Highest scores are shown in bold.

Method	Type of Fusion	Validation Set		Test Set	
		Valence	Arousal	Valence	Arousal
Zhang et al. [37]	Transformers	0.464	0.640	0.648	0.625
Zhang et al. [42]	Leader-Follower	0.441	0.645	0.552	0.629
Zhou et al. [36]	Transformers	0.550	0.681	0.500	0.632
Zhang et al. [17]	Transformers	0.554	0.659	0.523	0.545
Meng et al. [16]	Transformers	0.588	0.668	0.606	0.596
Praveen et al [21]	JCA	0.663	0.584	0.374	0.363
Praveen et al [23]	RJCA	0.703	0.623	0.467	0.405
Praveen et al [23]*	RJCA	0.443	0.639	0.537	0.576
RJCMA (Ours)	RJCMA	0.455	0.652	0.542	0.619

Table 3. CCC of the proposed fusion model by varying the number of recursions. Bold indicates the highest scores. Fold 1 is used for experiments with multiple recursions.

Num. of recursions (l)	Valence	Arousal	Mean
$l = 1$	0.568	0.656	0.607
$l = 2$	0.576	0.669	0.618
$l = 3$	0.585	0.674	0.629
$l = 4$	0.579	0.652	0.615

with multiple recursions is also observed for other audio-visual tasks such as person verification [25] and event localization [24], thereby underscoring our hypothesis.

5.2. Comparison to State-of-the-art

Most of the approaches evaluated on the Affwild2 dataset have been submitted to previous ABAW challenges. Therefore, we have compared the performance of the proposed approach with some of the relevant state-of-the-art models of previous ABAW challenges as shown in Table 2. Several approaches explored ensemble-based methods using multiple encoders for each modality, followed by transformers to encode the concatenated multimodal feature representation [16, 17, 36]. By exploiting multiple backbones and large-scale training with external datasets, Meng et al. [16] significantly improved the test set performance on both valence and arousal. Similarly, Zhou et al. [36] also explored multiple backbones and showed better improvement in the arousal performance of the test set than [16]. Even though ensemble-based methods show better performance on test set, it is often cumbersome and computationally expensive. Zhang et al. [37] showed that exploring Masked Auto-Encoders (MAE) is a promising line of research to

achieve better generalization and consistently improved performance in both valence and arousal. Zhang et al. [42] explored the leader-follower attention model for multimodal fusion, where audio and text modalities are leveraged to attend to the visual modality, and showed good performance on the test set without the need for multiple backbones for each modality.

Praveen et al. [21] proposed Joint Cross Attention (JCA) by introducing a joint feature representation in the cross-attention framework and showed significant improvement in the validation set, especially for valence. They further improved the performance of their approach by deploying a recursive mechanism and LSTMs for the temporal modeling of individual modalities and multimodal feature representations [23]. However, they do not seem to have a better generalization ability as they fail to show improvement in the test set. We hypothesize that this may be due to naive audio and visual backbones, fine-tuned on Affwild2 dataset. Therefore, to have a fair comparison with [23], we have reimplemented their fusion model by replacing their audio and visual backbones with Vggish [44] and Resnet-50 pre-trained on MSCELEB-1M and FER+ [42], followed by TCNs. By replacing the visual and audio backbones of [23] and progressive fine-tuning of the backbones similar to that of [42], the problem of over-fitting has been mitigated and showed improvement in test set performance also. We can observe that performance has been improved in both validation and test sets by introducing text modality into the RJCA framework [23], especially for arousal. Note that even though the performance of the valence for [37, 42] as well as our approach on the official validation set is lower, they achieved better performance on other folds of cross-validation, thus improving the performance of the test set.

Table 4. CCC of the proposed approach on Affwild2 test set compared to other methods submitted to 6th ABAW competition. Bold indicates best results of our approach on the test set.* denote citations are not available.

Method	Modalities	Ensemble	Valence	Arousal	Mean
Netease Fuxi AI Lab [62]	Audio, Visual	Yes	0.6873	0.6569	0.6721
RJCMA (Ours)	Audio, Visual, Text	No	0.5418	0.6196	0.5807
CtyunAI [63]	Visual	No	0.5223	0.6057	0.5640
SUN-CE [64]	Audio, Visual	Yes	0.5355	0.5861	0.5608
USTC-IAT-United [65]	Audio, Visual	No	0.5208	0.5748	0.5478
HSEmotion [66]	Visual	No	0.4925	0.5461	0.5193
KBS-DGU*	Audio, Visual	No	0.4836	0.5318	0.5077
ETS-LIVIA [67]	Audio, Visual	No	0.4198	0.4669	0.4434
CAS-MAIS*	Audio, Visual, Text	No	0.4245	0.3414	0.3830
IMLAB [68]	Visual	No	0.2912	0.2456	0.2684
Baseline [60]	Visual	No	0.2110	0.1910	0.2010

5.3. Results on Test Set

Multimodal learning is gaining a lot of attention for affective behavior analysis as most of the methods submitted for the valence-arousal challenge of the 6th ABAW competition [60] employed multimodal fusion [62, 64, 65, 67]. Some of these methods explored ensemble-based fusion for better generalization ability [62, 64]. Another widely explored strategy to improve the test set performance is to exploit pre-training with multiple large-scale datasets using Masked Auto-Encoders (MAE) [62, 63] or multiple backbones for each modality [64–66]. Netease Fuxi AI Lab [62] used MAEs pretrained with 5 external data sets of around 262M images and Vggish model [44] for the audio modality, followed by ensemble-based fusion, and showed a significant improvement for both valence and arousal, achieving first place in the challenge. Similarly, CtyunAI [63] also used MAE for visual modality pre-trained with 4 external datasets, followed by TCNs for temporal modeling, and achieved very good performance for arousal than valence. Sun-CE [64] explored multiple backbones for both audio and visual modalities, followed by ensemble-based fusion, while USTC-IAT-United [65] used multiple backbones only for the audio modality and achieved decent performance for both valence and arousal. KBS-DGU and ETS-LIVIA [67] exploited combining features from multiple fusion models of self-attention and cross-attention models for improving the performance of audio-visual fusion. KBS-DGU explored both self-attention and cross-attention models, while ETS-LIVIA [67] used joint feature representation with cross-attention models across audio and visual modalities. Unlike other approaches, we have explored the prospect of simultaneously capturing the intra- and inter-modal relationships across audio, visual, and text modal-

ities recursively and achieved very good performance for both valence and arousal without using multiple backbones or pre-training with large-scale external datasets. Therefore, the performance of our approach can be solely attributed to the robustness of the sophisticated fusion model, providing a cost-effective solution for affective behavior analysis in-the-wild.

6. Conclusion

In this work, we showed that effectively capturing the synergic relationships pertinent to the intra- and inter-modal characteristics among the audio, visual, and text modalities can significantly improve the performance of the system. By introducing the joint representation in the cross-attentional framework, we can simultaneously capture both intra- and inter-modal relationships across the audio, visual, and text modalities. The performance of the system is further enhanced using recursive fusion by progressively refining the features to obtain robust multimodal feature representations. Experimental results in the challenging Affwild2 dataset indicate that the proposed model can achieve a better multimodal fusion performance, outperforming most methods for both valence and arousal. The performance of the system can also be improved by leveraging advanced text encoders and sophisticated backbones for the individual modalities.

Acknowledgements The authors wish to acknowledge the funding from the Government of Canada’s New Frontiers in Research Fund (NFRF) through grant NFRFR-2021-00338 and Ministry of Economy and Innovation (MEI) of the Government of Quebec for the continued support.

References

- [1] Sicheng Zhao, Xingxu Yao, Jufeng Yang, Guoli Jia, Guiguang Ding, Tat-Seng Chua, Björn W. Schuller, and Kurt Keutzer. Affective image content analysis: Two decades review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6729–6751, 2022. 1
- [2] R Gnana Praveen, Eric Granger, and Patrick Cardinal. Deep weakly supervised domain adaptation for pain localization in videos. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 473–480, 2020. 1
- [3] Gnana Praveen Rajasekhar, Eric Granger, and Patrick Cardinal. Deep domain adaptation with ordinal regression for pain assessment using weakly-labeled videos. *Image and Vision Computing*, 110:104167, 2021. 1
- [4] Zhongshan Chen, Xinning Feng, and Shengwei Zhang. Emotion detection and face recognition of drivers in autonomous vehicles in iot platform. *Image and Vision Computing*, 128:104569, 2022. 1
- [5] Maja Pantic, Nicu Sebe, Jeffrey F. Cohn, and Thomas Huang. Affective multimodal human-computer interaction. In *Proceedings of the 13th Annual ACM International Conference on Multimedia*, page 669–676, 2005. 1
- [6] Paul Ekman. An argument for basic emotions. *Cognition and Emotion*, 6(3-4):169–200, 1992. 1
- [7] Dimitrios Kollias. Multi-label compound expression recognition: C-expr database & network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5589–5598, 2023. 1
- [8] H. Schlosberg. Three dimensions of emotion. *Psychological Review*, 61:81–88, 1954. 1
- [9] Daheng Wang, Tong Zhao, Wenhao Yu, Nitesh V. Chawla, and Meng Jiang. Deep multimodal complementarity learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12):10213–10224, 2023. 1
- [10] Panagiotis Tzirakis, George Trigeorgis, Mihalis A. Nicolaou, Björn W. Schuller, and Stefanos Zafeiriou. End-to-end multimodal er using deep neural networks. *IEEE J. of Selected Topics in Signal Proc.*, 11(8):1301–1309, 2017. 1, 2, 3, 6
- [11] L Schoneveld, A Othmani, and H Abdelkawy. Leveraging recent advances in deep learning for audio-visual emotion recognition. *Pattern Rec. Letters*, 146:1–7, 2021. 1, 2, 3
- [12] Felix Kuhnke, Lars Rumberg, and Jörn Ostermann. Two-stream aural-visual affect analysis in the wild. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 600–605, 2020. 1, 2, 3
- [13] Juan D. S. Ortega, Patrick Cardinal, and Alessandro L. Koerich. Emotion recognition using fusion of audio and video features. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pages 3847–3852, 2019. 1
- [14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. 1, 2
- [15] Srinivas Parthasarathy and Shiva Sundaram. Detecting expressions with multimodal transformers. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 636–643, 2021. 1, 2
- [16] Liyu Meng, Yuchen Liu, Xiaolong Liu, Zhaopei Huang, Wenqiang Jiang, Tenggao Zhang, Chuanhe Liu, and Qin Jin. Valence and arousal estimation based on multimodal temporal-aware features for videos in the wild. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2344–2351, 2022. 2, 3, 7
- [17] Ziyang Zhang, Liuwei An, Zishun Cui, Ao Xu, Tengdong Dong, Yueqi Jiang, Jingyi Shi, Xin Liu, Xiao Sun, and Meng Wang. Abaw5 challenge: A facial affect recognition approach utilizing transformer encoder and audiovisual fusion. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 5725–5734, 2023. 1, 2, 7
- [18] Jun-Tae Lee, Mihir Jain, Hyungwoo Park, and Sungrack Yun. Cross-attentional audio-visual fusion for weakly-supervised action localization. In *International Conference on Learning Representations*, 2021. 2
- [19] R. Gnana Praveen, Eric Granger, and Patrick Cardinal. Cross attentional audio-visual fusion for dimensional emotion recognition. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pages 1–8, 2021. 2
- [20] Gnana Praveen Rajasekhar and Jahangir Alam. Audio-visual speaker verification via joint cross-attention. In *Speech and Computer*, pages 18–31, 2023. 2
- [21] R Gnana Praveen, Wheidima Carneiro de Melo, Nasib Ullah, Haseeb Aslam, Osama Zeeshan, Théo Denorme, Marco Pedersoli, Alessandro L. Koerich, Simon Bacon, Patrick Cardinal, and Eric Granger. A joint cross-attention model for audio-visual fusion in dimensional emotion recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2486–2495, June 2022. 2, 3, 6, 7
- [22] R. Gnana Praveen, Patrick Cardinal, and Eric Granger. Audio-visual fusion for emotion recognition in the valence-arousal space using joint cross-attention. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 5(3):360–373, 2023. 2, 3
- [23] R Gnana Praveen, Eric Granger, and Patrick Cardinal. Recursive joint attention for audio-visual fusion in regression based emotion recognition. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023. 2, 3, 7
- [24] Bin Duan, Hao Tang, Wei Wang, Ziliang Zong, Guowei Yang, and Yan Yan. Audio-visual event localization via recursive fusion by joint co-attention. In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 4012–4021, 2021. 2, 7
- [25] R Gnana Praveen and Jahangir Alam. Audio-visual person verification based on recursive fusion of joint cross-attention. *arXiv preprint arXiv:2403.04654*, 2024. 2, 7
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 2, 3

- [27] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri. A closer look at spatiotemporal convolutions for action recognition. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6450–6459, 2018. 2, 3
- [28] Dimitrios Kollias and Stefanos Zafeiriou. Affect analysis in-the-wild: Valence-arousal, expressions, action units and a unified framework. *arXiv:2103.15792*, 2021. 2
- [29] Dimitrios Kollias, Viktoriia Sharmanska, and Stefanos Zafeiriou. Distribution matching for heterogeneous multi-task learning: a large-scale face study. *arXiv:2105.03790*, 2021.
- [30] Dimitrios Kollias, Viktoriia Sharmanska, and Stefanos Zafeiriou. Face behavior a la carte: Expressions, affect and action units in a single network. *arXiv:1910.11111*, 2019. 2
- [31] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [32] Lingfeng Wang, Shisen Wang, Jin Qi, and Kenji Suzuki. A multi-task mean teacher for semi-supervised facial affective behavior analysis. In *ICCV Workshop*, 2021. 2
- [33] Didan Deng, Liang Wu, and Bertram E. Shi. Iterative distillation for better uncertainty estimates in multitask emotion recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 3557–3566, October 2021. 2
- [34] Muhammad Haseeb Aslam, Muhammad Osama Zeeshan, Marco Pedersoli, Alessandro L. Koerich, Simon Bacon, and Eric Granger. Privileged knowledge distillation for dimensional emotion recognition in the wild. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 3338–3347, 2023. 2
- [35] Yong Li, Yuanzhi Wang, and Zhen Cui. Decoupled multimodal distilling for emotion recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6631–6640, June 2023. 2
- [36] Weiwei Zhou, Jiada Lu, Zhaolong Xiong, and Weifeng Wang. Leveraging tcn and transformer for effective visual-audio fusion in continuous emotion recognition. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 5756–5763, 2023. 2, 3, 7
- [37] Wei Zhang, Bowen Ma, Feng Qiu, and Yu Ding. Multimodal facial affective analysis based on masked autoencoder. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 5793–5802, 2023. 2, 7
- [38] Minh Tran and Mohammad Soleymani. A pre-trained audio-visual transformer for emotion recognition. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4698–4702, 2022. 2
- [39] J. S. Chung, A. Nagrani, and A. Zisserman. Voxceleb2: Deep speaker recognition. In *INTERSPEECH*, 2018. 2
- [40] Jian Huang, Jianhua Tao, Bin Liu, Zheng Lian, and Mingyue Niu. Multimodal transformer fusion for continuous emotion recognition. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3507–3511, 2020. 2
- [41] Vincent Karas, Mani Kumar Tellamekala, Adria Mallol-Ragolta, Michel Valstar, and Björn W. Schuller. Time-continuous audiovisual fusion with recurrence vs attention for in-the-wild affect recognition. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2381–2390, 2022. 2
- [42] Su Zhang, Ziyuan Zhao, and Cuntai Guan. Multimodal continuous emotion recognition: A technical report for abaw5. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 5764–5769, 2023. 2, 3, 6, 7
- [43] Emad Barsoum, Cha Zhang, Cristian Canton Ferrer, and Zhengyou Zhang. Training deep networks for facial expression recognition with crowd-sourced label distribution. In *Proceedings of the 18th ACM international conference on multimodal interaction*, pages 279–283, 2016. 3
- [44] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. Cnn architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*, pages 131–135. IEEE, 2017. 3, 7, 8
- [45] Mihalis A. Nicolaou, Hatice Gunes, and Maja Pantic. Continuous prediction of spontaneous affect from multiple cues and modalities in valence-arousal space. *IEEE Trans. on Affective Computing*, 2:92–105, 2011. 3
- [46] M Wöllmer, M Kaiser, F Eyben, B Schuller, and G Rigoll. Lstm-modeling of continuous emotions in an a-v affect recognition framework. *Image and Vision Computing*, 31(2):153–163, 2013. 3
- [47] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *Computer Vision – ECCV 2016*, pages 87–102, 2016. 3
- [48] S Zhang, Y Ding, Z Wei, and C Guan. Continuous emotion recognition with audio-visual leader-follower attentive fusion. In *ICCV Workshop*, 2021. 3
- [49] S Hershey, S Chaudhuri, D Ellis, J F Gemmeke, A Jansen, C Moore, M Plakal, D Platt, R A Saurous, B Seybold, M Slaney, R Weiss, and K Wilson. Cnn architectures for large-scale audio classification. In *ICASSP 2017*. 3
- [50] Abdullah Al Maruf, Fahima Khanam, Md. Mahmudul Haque, Zakaria Masud Jiyad, M. F. Mridha, and Zeyar Aung. Challenges and opportunities of text-based emotion detection: A survey. *IEEE Access*, 12:18416–18450, 2024. 3
- [51] Francisca Adoma Acheampong, Henry Nunoo-Mensah, and Wenyu Chen. Transformer models for text-based emotion detection: a review of bert-based approaches. *Artificial Intelligence Review*, 54:5789 – 5829, 2021. 3
- [52] D Kollias, P Tzirakis, M A Nicolaou, A Papaioannou, G Zhao, B Schuller, I Kotsia, and S Zafeiriou. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *IJCV*, 127:907–929, 2019. 5
- [53] Dimitrios Kollias and Stefanos Zafeiriou. Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arface. *arXiv:1910.04855*, 2019. 5

- [54] Stefanos Zafeiriou, Dimitrios Kollias, Mihalios A Nicolaou, Athanasios Papaioannou, Guoying Zhao, and Irene Kotzia. Aff-wild: Valence and arousal ‘in-the-wild’ challenge. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pages 1980–1987. IEEE, 2017. 5
- [55] Dimitrios Kollias, Attila Schulc, Elnar Hajiyev, and Stefanos Zafeiriou. Analysing affective behavior in the first abaw 2020 competition. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 637–643, 2020.
- [56] Dimitrios Kollias and Stefanos Zafeiriou. Analysing affective behavior in the second abaw2 competition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3652–3660, 2021.
- [57] Dimitrios Kollias. Abaw: Valence-arousal estimation, expression recognition, action unit detection and multi-task learning challenges. *arXiv:2202.10659*, 2022.
- [58] Dimitrios Kollias. Abaw: learning from synthetic data & multi-task learning challenges. In *European Conference on Computer Vision*, pages 157–172. Springer, 2023.
- [59] Dimitrios Kollias, Panagiotis Tzirakis, Alice Baird, Alan Cowen, and Stefanos Zafeiriou. Abaw: Valence-arousal estimation, expression recognition, action unit detection & emotional reaction intensity estimation challenges. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5888–5897, 2023. 5
- [60] Dimitrios Kollias, Panagiotis Tzirakis, Alan Cowen, Stefanos Zafeiriou, Chunchang Shao, and Guanyu Hu. The 6th affective behavior analysis in-the-wild (abaw) competition. *arXiv preprint arXiv:2402.19344*, 2024. 5, 6, 8
- [61] Licai Sun, Zheng Lian, Jianhua Tao, Bin Liu, and Mingyue Niu. Multi-modal continuous dimensional emotion recognition using recurrent neural network and self-attention mechanism. In *Proceedings of the 1st international on multimodal sentiment analysis in real-life media challenge and workshop*, pages 27–34, 2020. 5
- [62] Wei Zhang, Feng Qiu, Chen Liu, Lincheng Li, Heming Du, Tiancheng Guo, and Xin Yu. Affective behaviour analysis via integrating multi-modal knowledge. *arXiv preprint arXiv:2403.10825*, 2024. 8
- [63] Weiwei Zhou, Jiada Lu, Chenkun Ling, Weifeng Wang, and Shaowei Liu. Boosting continuous emotion recognition with self-pretraining using masked autoencoders, temporal convolutional networks, and transformers. *arXiv preprint arXiv:2403.11440*, 2024. 8
- [64] Denis Dresvyanskiy, Maxim Markitantov, Jiawei Yu, Peitong Li, Heysem Kaya, and Alexey Karpov. Sun team’s contribution to abaw 2024 competition: Audio-visual valence-arousal estimation and expression recognition. *arXiv preprint arXiv:2403.12609*, 2024. 8
- [65] Jun Yu, Gongpeng Zhao, Yongqi Wan, Zhihong Wei, Yang Zheng, Zerui Zhang, Zhongpeng Cai, Guochen Xie, Jichao Zhu, and Wangyuan Zhu. Multimodal fusion method with spatiotemporal sequences and relationship learning for valence-arousal estimation. *arXiv preprint arXiv:2403.12425*, 2024. 8
- [66] Andrey V Savchenko. Hsemotion team at the 6th abaw competition: Facial expressions, valence-arousal and emotion intensity prediction. *arXiv preprint arXiv:2403.11590*, 2024. 8
- [67] Paul Waligora, Osama Zeeshan, Haseeb Aslam, Soufiane Belharbi, Alessandro Lameiras Koerich, Marco Pedersoli, Simon Bacon, and Eric Granger. Joint multimodal transformer for dimensional emotional recognition in the wild. *arXiv preprint arXiv:2403.10488*, 2024. 8
- [68] Seongjae Min, Junseok Yang, Sangjun Lim, Junyong Lee, Sangwon Lee, and Sejoon Lim. Emotion recognition using transformers with masked learning. *arXiv preprint arXiv:2403.13731*, 2024. 8