

Joint Multimodal Transformer for Emotion Recognition in the Wild

Paul Waligora^{*1}, Muhammad Haseeb Aslam^{*1}, Muhammad Osama Zeeshan¹, Soufiane Belharbi¹, Alessandro Lameiras Koerich¹, Marco Pedersoli¹, Simon Bacon² & Eric Granger¹

¹ LIVIA, Dept. of Systems Engineering, ETS Montreal, Canada

² Dept. of Health, Kinesiology & Applied Physiology, Concordia University, Montreal, Canada

* These authors contributed equally

Abstract

Multimodal emotion recognition (MMER) systems typically outperform unimodal systems by leveraging the inter- and intra-modal relationships between, e.g., visual, textual, physiological, and auditory modalities. This paper proposes an MMER method that relies on a joint multimodal transformer (JMT) for fusion with key-based cross-attention. This framework can exploit the complementary nature of diverse modalities to improve predictive accuracy. Separate backbones capture intra-modal spatiotemporal dependencies within each modality over video sequences. Subsequently, our JMT fusion architecture integrates the individual modality embeddings, allowing the model to effectively capture inter- and intra-modal relationships. Extensive experiments on two challenging expression recognition tasks – (1) dimensional emotion recognition on the Affwild2 dataset (with face and voice) and (2) pain estimation on the Biovid dataset (with face and biosensors) – indicate that our JMT fusion can provide a cost-effective solution for MMER. Empirical results show that MMER systems with our proposed fusion allow us to outperform relevant baseline and state-of-the-art methods. Code is available at: <https://github.com/PoloWlg/Joint-Multimodal-Transformer-6th-ABAW>

1. Introduction

Human-computer interaction is applied in a wide range of real-world scenarios, e.g., health care, the Internet of Things (IoT), and autonomous driving. Researchers have classified human emotions in different ways, most notably according to discrete categories, ordinal intensity levels, and the valence/arousal circumplex [1]. With the recent advancements in deep learning and sensor technologies, research in affective computing has evolved from lab-controlled to real-world (in the wild) scenarios. In the latter, human emo-

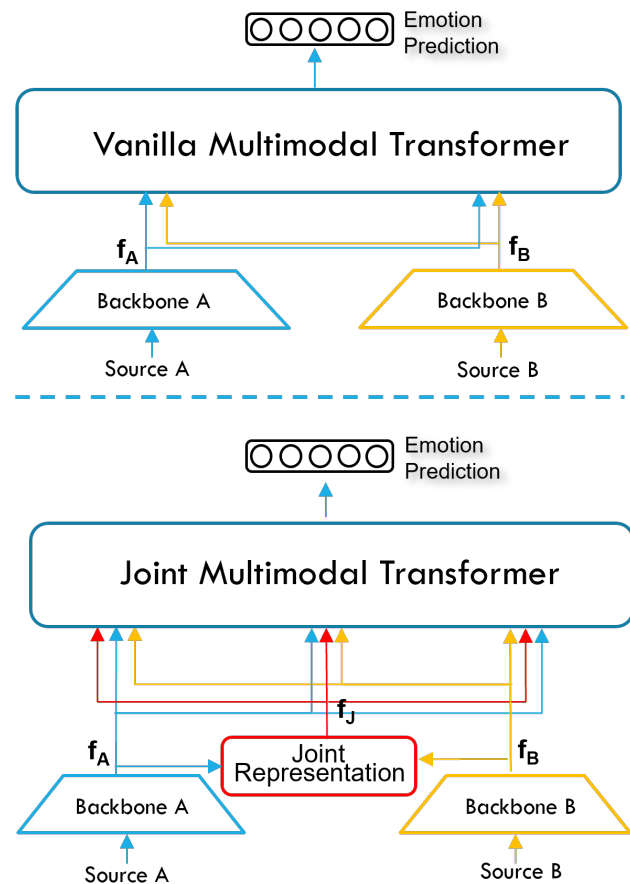


Figure 1. (Top) An illustration of the vanilla multimodal transformer fusion architecture in the case of two input sources, A and B. (Bottom) Our proposed JMT fusion (in red) relies on joint multimodal representations.

tions are usually expressed over a broader spectrum beyond the six basic categorical expressions - anger, disgust, fear, happy, sad, and surprise [2]. Therefore, there is much interest in analyzing and modeling complex and subtle ex-

Emails: paul.waligora.1@ens.etsmtl.ca, eric.granger@etsmtl.ca

pressions of emotions in real-world scenarios. For instance, the spectrum of emotions can be formulated as dimensional emotional recognition (ER), where complex human emotions are represented along arousal (intensity) and valence (positiveness) axes.

Multimodal fusion has been widely explored for the problem of video-based ER in the literature [3, 4]. For instance, audio and visual modalities may provide complementary and redundant information over a video sequence. These relationships must be captured to model the intricacies of human emotions effectively. Furthermore, effectively capturing both the intra-modal temporal dependencies within the audio and visual modalities and the inter-modal association across the audio and visual modalities is crucial in developing an effective AER system [5].

Several methods have been proposed for video-based ER and recurrent networks have been employed to capture the intra-modal temporal dependencies from video sequences [3, 4, 6]. Recently, attention-based methods have been introduced to extract features that are the most relevant to downstream tasks. Cross-attention-based methods have also been [7] employed to capture the inter-modal association between the audio, visual, and other modalities. Lu et al. [8] proposed ViLBERT, the seminal work in multimodal co-attention. Since then, many transformer-based cross-attention methods have been proposed [9, 10]. These methods, however, cannot effectively capture the intra-modal temporal dynamics. Further, they specialize in capturing the complementary information among the modalities but do not include a mechanism to explicitly capture the redundant information.

The proposed method introduces a third branch with the joint representation of the multiple modalities, as shown in Figure 1. By incorporating a joint representation branch, the model can access additional contextual information that may not be fully captured by cross-attention alone. Such a joint representation branch can help improve the model's understanding of complex relationships between the input sequences. Further, the proposed method becomes more robust to noise or irrelevant information present in individual sequences, which helps mitigate the sensitivity of cross-attention to noisy inputs and improves the system's overall performance.

Our main contributions are summarized as follows.

(1) This paper proposes a joint multimodal transformer (JMT) fusion architecture that leverages joint modality representations. It captures inter- and intra-modal information in videos using key-based cross-attention, and exploits the redundant and complementary associations among modalities. (2) An extensive set of experiments on two challenging emotion recognition datasets (pain estimation on BioVid and dimension valence-arousal assessment on Affwild2) indicate that our proposed JMT fusion architecture can out-

perform relevant baseline and state-of-the-art methods.

2. Related Work in Emotion Recognition

2.1. Multimodal Methods

MMER refers to integrating multiple sources of information (modalities) to improve the accuracy and robustness of automated emotion recognition systems at the expense of complexity. These modalities typically include visual, audio, textual, and physiological. The seminal work in multimodal deep learning was proposed by Ngiam et al. [11], where the features from the audio and visual modalities were extracted separately, and then autoencoders and Restricted Boltzmann Machines were used to feature fusion. Tzirakis et al. [12] proposed one of the early approaches for A-V fusion for dimensional emotion recognition, in which the visual features were extracted using a ResNet50 and the audio features were obtained using a 1D convolutional neural network (CNN). The modality-specific features were concatenated and fed to a recurrent net for simultaneous temporal modeling and modality fusion. An empirical study was presented by Juan et al. [13], where the authors studied the impact of fine-tuning multiple layers in a pretrained CNN for the visual modality.

A two-stream autoencoder with a long short-term memory (LSTM) network was proposed by Nguyen et al. [14] to jointly learn and compact representative features from the visual and audio modalities. A knowledge distillation-based approach was investigated by Schonevald et al. [3] for visual modality. For the audio modality, spectrograms were obtained and fed to a CNN model, and the two modalities were fused using a recurrent net. A novel self-distillation scheme was put forward by Deng et al. [15] to overcome the problem of noisy labels in a multitasking setting. A two-stream aural visual (TSAV) network was proposed by Kuhnke et al. [4], in which the audio features were extracted using a ResNet18, and the visual features were extracted using a 3D-CNN. The obtained embeddings were fed to a specially designed TSAV network for information fusion.

Pain estimation is one of the primary problems in affective computing. Researchers have proposed many multimodal datasets for the pain estimation task. The facial activity descriptors method for pain estimation was introduced by Werner et al. [16]. Dragomir et al. [17] propose a subject-independent method from facial images with a residual learning technique. A Sparse LSTM-based method was proposed by Zhi et al. [18] to solve the problem of vanishing gradients in temporal learning. Morabit et al. [19] proposed a data-efficient image transformer. To process multiscale electrodermal activity signals, a SE-Net-based network was proposed by Lu et al. [20]. Multimodal solutions to fuse the physiological and visual modalities were proposed by Werner et al. [16], Kachele et al. [21], and Zhu

et al. [18]. Physiological signals are more discriminative for pain classification than the visual modality.

2.2. Attention-Based and Transformer Methods

Since its inception, attention models have shown extraordinary performance in many applications. These models have been extensively investigated for capturing the inter and intra-modal associations between the audio and visual modalities for tasks like action localization [22], A-V event localization [23], and multimodal emotion recognition [9]. An attention-based fusion mechanism was proposed by Zhang et al. [24], 3D-CNNs and 2D-CNNs were used to extract multi-features in the visual modality, and for the audio modality, a 2D-CNN was used to learn representation from spectrograms. Specialized scoring functions were used to re-weight the audio and visual features.

Recently, cross-modal attention has shown promising results because of its ability to model inter-modal relationships. Srinivas et al. [9] explored a transformer network with encoder layers, where cross-modal attention is used to fuse audio and visual features for continuous arousal/valence prediction in the wild. Tzirakis et al. [25] explored the idea of cross-attention in conjunction with self-attention. The authors proposed a transformer-based fusion architecture. Although the methods mentioned above have used cross-modal attention with transformers, they do not have any explicit mechanism to capture semantic relevance between the A-V features, particularly the intra-modal correlations. Zhang et al. [26] proposed a method for A-V fusion using leader-follower attentive fusion for continuous arousal/valence prediction. Attention weights are combined with the encoded visual and audio features. Cross attention presented in Praveen et al. [7] has shown a substantial increase in performance by using cross-correlation across the individual features. In contrast, our proposed method uses key-based cross-attention in multimodal transformers and explores the idea of feeding the joint A-V feature vector. By feeding the joint A-V feature representation, the proposed method effectively captures the inter- and intra-modal relationships simultaneously by interacting across itself and the other modalities.

Huang et al. [27] investigated the idea of multi-head attention in transformer-based fusion architecture, which was further combined with LSTM to capture the high-level representations. Tran et al. [28] proposed a cross-modal transformer architecture that consisted of a multimodal cross-modal attention block, where the Queries were generated from one modality and the key values were generated from the other modality. Le et al. [29] put forward an end-to-end transformer-based fusion mechanism for multilabel multimodal emotion classification; the model consisted of three parts: i) three backbone networks for visual, audio, and textual feature extractor, ii) a transformer network for informa-

tion fusion, and iii) classification network. Zhou et al. [30] proposed a transformer-based fusion scheme along with the temporal convolutional network (TCN); the audio and visual features were extracted using pretrained backbones followed by a TCN, the output of TCN was concatenated and fed to a transformer encoder block. A multilayer perceptron (MLP) was then used for the final prediction.

All the aforementioned transformer-based fusion architectures primarily focus on intermodality correlation. In contrast, in addition to modeling the intermodality relationships to capture the complementarity between modalities, the proposed method explicitly feeds the joint (combined) features to the multimodal transformer to introduce redundancy. By incorporating this third joint representation branch, the proposed model can access enhanced contextual information that cross-attention might only partially capture. Doing this improves the model's understanding of complex relationships between the input sequences. Further, the proposed method becomes more robust to noise or irrelevant information present in individual sequences. This third joint representation allows the model to dynamically focus on this newly introduced information in sequences where both modalities are simultaneously noisy. This helps mitigate the sensitivity of cross-attention to noisy inputs and improves the system's overall performance.

3. Proposed Approach

The proposed method is a hierarchical fusion mechanism, where the intra-modality features are combined using transformer-based self-attention, and cross-modality features are fused using transformer-based cross-attention. Further, we feed a third joint representation to the joint transformer module (JTM) to enhance robustness. The K (key matrix), V (value matrix), and Q (query matrix) vectors are shared among the six transformer blocks. In the end, the output of these six blocks is again fed to a transformer self-attention block to weigh the most relevant representations dynamically. The final prediction is made using fully connected (FC) layers.

3.1. Modality Specific Feature Extraction

In the first step, modality-specific features are extracted using backbones. The proposed method allows combining multiple backbones for each modality to improve system robustness. The extracted feature vectors from each backbone are fed to a transformer self-attention block. The combined feature vector represents the particular modality. For example, to capture information about a person's emotional state, we can use an R(2+1)D CNN pretrained on the Kinetics-400 dataset [31] to extract visual features. For the audio modality, we could extract features using a ResNet18 [32] CNN with a GRU [33].

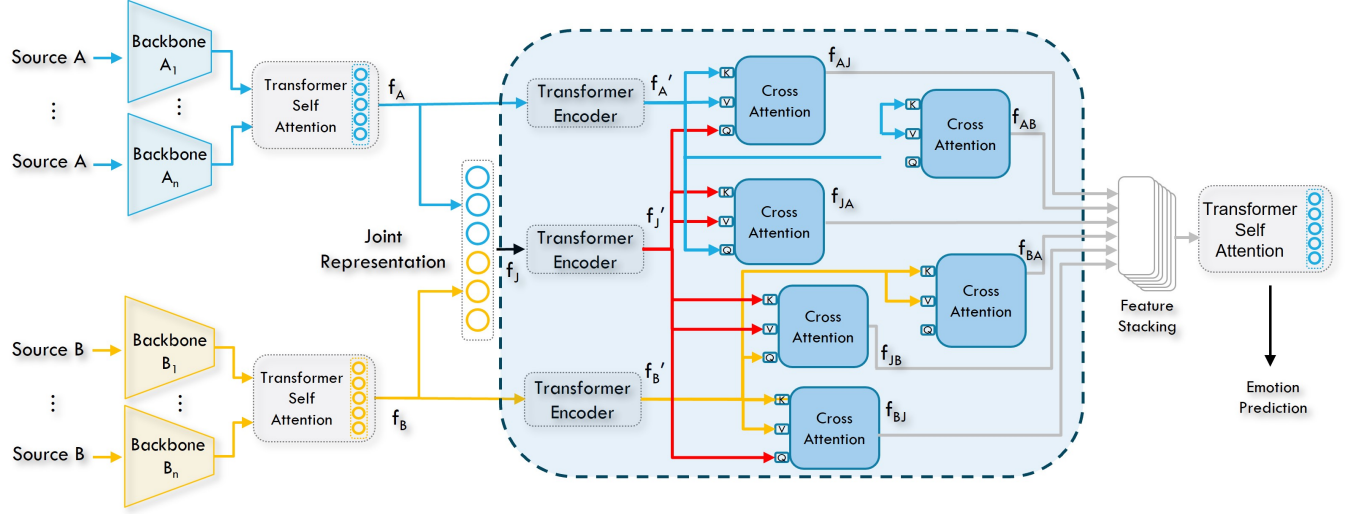


Figure 2. An overview of the proposed joint multimodal transformer model for A-V fusion. The audio and visual modalities are cross-attended using transformer blocks. The JMT block also takes in the joint representation (shown with red arrows). The output of the cross-attended features is concatenated, and an FC layer is used for valence/arousal prediction.

3.2. Multi-transformer Attention Fusion

We define f_A as the deep features extracted from backbone A, and f_B as the deep features extracted from the backbone B in Figure 2. Given f_B and f_A , the joint feature representation is obtained by concatenating f_B and f_A feature vectors:

$$f_J = [f_B; f_A], \quad (1)$$

where $[\cdot; \cdot]$ denotes a concatenation operation. The concatenated features are then fed to an FC layer for dimensionality reduction of the joint feature representation to yield f_J . We now have three key sources of information: f_B , f_A and f_J , which have the same dimensionality.

Each representation is then fed to a specific encoder. Our model is composed of three different encoders, one for each type of feature f_B , f_A , and f_J . Each encoder consists of a multi-head-self-attention (Eq. 2) followed by a fully connected feed-forward network. Residual connection and layer normalization are performed around both of these layers. The key is used to associate a sequence to a key value, the value matrix holds information that is ultimately used to compute the output of the attention mechanism, and the query matrix represents a set of vectors used to query the key-value pairs. The K , V , Q matrix are calculated this way: $K = XW_K$, $Q = XW_Q$, $V = XW_V$. X corresponds to one of the sources of information and W_K , W_Q , W_V are the weights of the key, query, value matrices respectively. The output values of the self-attention layers are given by:

$$\text{Attention}(Q, K, V) = \text{softmax}(KQ^T/d_k)V \quad (2)$$

With self-attention layers, each encoder focuses independently on important cues related to its respective source of information. Afterward, each encoder embedding is combined by utilizing six cross-modal attention layers where the query matrix Q is shared with the key K and value V matrix of the other source of information. Sharing this matrix between each source of information helps the model add redundancy and complement the visual and audio modalities, thus improving its performance. At the output of each of the six cross-attention modules, the feature vector of dimension 512 is output. These six feature vectors are then stacked to form a sequence, which is then fed to a transformer self-attention block. This block dynamically selects and weighs these feature vectors. The final attended features are fed to an FC layer for final prediction.

The model aims to maximize the Concordance Correlation Coefficient (CCC) [34], which is common in dimensional emotion recognition. To achieve this, we minimize the following loss:

$$L_c = 1 - \rho_c = 1 - \frac{2\rho_{xy}^2}{\rho_x^2 + \rho_y^2 + (\mu_x - \mu_y)^2} \quad (3)$$

where ρ_{xy}^2 is the covariance between the predictions and the ground truth, ρ_x^2 and ρ_y^2 the variance of the prediction and the ground truth, μ_x and μ_y the mean of prediction and the ground truth.

4. Experimental Methodology

4.1. Datasets

Affwild2 was put forward by [35]. It is one of the largest and most comprehensive datasets for affective computing focused in-the-wild scenarios. The dataset is also a part of the Affective Behavior Analysis in the Wild (ABAW) challenge [36–50]. The dataset comprises 564 videos that were collected from YouTube. These videos are labeled for three main affective computing tasks: i) categorical expression recognition, ii) continuous valence arousal prediction, and iii) action unit detection. The dataset comes with a train, validation, and test split with 341, 71, and 152 videos, respectively. The continuous valence/arousal annotation set is used to validate the proposed method.

Biovid Heat Pain Database: The BioVid Heat Pain Database [51] comprises 87 subjects where heat pain was induced experimentally at the right arm with four different intensities. The data comprises video and depth map video from a Kinect camera, galvanic skin response (EDA), electromyograph (EMG) on trapezius muscle, and electrocardiogram (ECG). The Biovid dataset has various partitions from Part A through E. These partitions differ in the modalities, annotations, and tasks. We use Part A of the dataset and utilize the video and raw EDA data to validate our proposed method.

4.2. Implementation Details

Affwild2: In the visual modality, cropped and aligned facial images provided with the dataset are used [43]. Black frames (zero pixels) replace missing frames in the visual modality. These images are then fed to a 3D network that takes the input size of 224×224 . A clip length of 8 is used, which makes up a sub-sequence of 64 frames. Each sub-sequence contains eight clips. A dropout with a value of 0.8 was used in the linear layers for network regularization. The network was optimized using the stochastic gradient descent (SGD) iterative method with an initial learning rate (LR) of 10^{-3} . The batch size used for the visual modality was 8. For further generalization, random cropping and random horizontal flips were added as data augmentation. The maximum number of epochs is 50, with early stopping for model selection. In the audio modality, the audio from each video is separated and resampled for 44 kHz. Following the segmentation in the visual modality, small vocal segments are extracted corresponding to the length of the clip. A ResNet18 is used to extract the features from the spectrograms. A Discrete Fourier transform with a length of 1024, a hop length of 10 msec, and a window length of 20 msec is used to obtain spectrograms of each vocal segment. The spectrograms have a resolution of 64×107 pixels, corresponding to a single clip in the visual modality. Other

preprocessing of spectrograms include mean and variance normalization, as well as conversion to log-power spectrum. The first convolutional layer in the pretrained ResNet model is adapted to take in the single-channel spectrograms. The learning rate of 1×10^{-2} is used and Adam optimizer is used to optimize the network. 64 batch size is set for the audio modality.

The audio and visual backbones are frozen to train the A-V fusion network and only train the whole transformer-based fusion model. Each of these backbones outputs deep feature vectors of dimension 512. These features are concatenated (to obtain the joint feature representation, a feature vector of dimension 1024) and then fed to an FC layer to reduce dimensionality to 512. Each of these features is fed to the fusion model, as we can see in Figure 2. At the output of each of the six cross-attention modules, the feature vector of dimension 512 is outputted. These six feature vectors are then stacked to form a sequence, which is then fed to a transformer self-attention block. This block dynamically selects and weighs these feature vectors. The final attended features are fed to an FC layer for final prediction. We perform a grid search to find the optimal learning rate and batch size for our fusion network. Thus, we use a list of learning rates: $[8 \times 10^{-4}, 6 \times 10^{-4}, 3 \times 10^{-4}]$ and an SGD optimizer to optimize the fusion network. We use a batch size of 32, and the maximum number of epochs is 5, with early stopping for fusion model selection. We select the best model among the learning rates listed before.

Biovid: In the visual modality, the faces are cropped and aligned using an MTCNN. We apply the frame retention strategy for missing frames where the face is not visible, and the MTCNN cannot capture any frame. Further, to ensure noise-free input to the visual model, we clip the first 2 seconds of the video and the last 0.5 seconds at the end because the subjects show no signs of pain during this duration. The total number of frames is 75. The extracted faces are fed to an R3D model for a visual feature extractor. The batch size is set to 64. The network is separately optimized using the SGD optimizer and the learning rate of 10^{-3} .

For the physiological modality, the EDA is used. The signal is clipped to correspond to the visual modality and fed to a 1D CNN. The architecture of the custom 1D-CNN is shown in Table 1. The CNN outputs a 512-dimensional feature vector. The model is optimized using the SGD optimizer with a learning rate of 10^{-4} . The batch size for the physiological backbone is set to 1024.

For the modality fusion, the two feature vectors and the joint representation are fed to the joint transformer block, as shown in Figure 3. The FC layers are removed from both backbones that were added in the backbone training phase. 512-dimensional feature vectors from visual and physiological backbones are obtained and fed to the joint multimodal

transformer module. The backbones are frozen, and the joint transformer block is optimized using the ADAM optimizer with a learning rate of 5×10^{-6} , and the batch size is set to 128.

Table 1. Description of the architecture of the custom 1D-CNN for the physiological backbone.

Layer type	No. of Filters	Size of Kernel	Stride	Output
Input	-	-	-	2816×1
1st Conv	32	5	2	1406×32
ReLU	-	-	-	1406×32
Max Pooling	-	2	-	703×32
2nd Conv	64	5	1	699×64
ReLU	-	-	-	699×64
Max Pooling	-	2	-	349×64
1st FC	-	-	-	512
ReLU	-	-	-	512
2nd FC	-	-	-	2

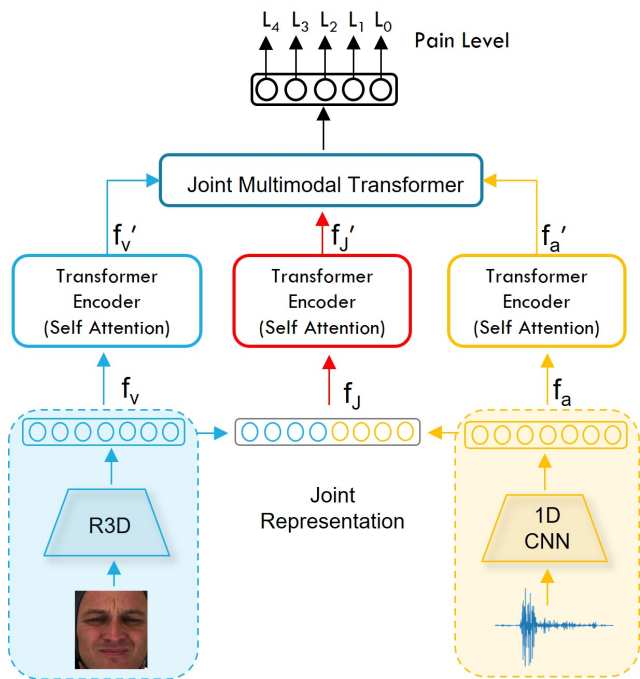


Figure 3. Illustration of the proposed joint multimodal transformer architecture used for the Biovid pain estimation task. The blue branch shows the visual backbone, and the yellow branch is the physiological backbone. The joint representation is shown with a red block. The three feature vectors are fed into the joint transformer block.

5. Results and Discussion

5.1. Comparison with the State-of-the-Art

This section compares the proposed method with the baseline and state-of-the-art on the Affwild2 and Biovid datasets.

Table 2 presents the performance of our fusion model on the Biovid Heat Pain Database. We performed 5-fold cross-validation to pick up the best average fusion model. It can be seen from the table that the proposed model can achieve state-of-the-art performance with multimodal input while using 5-fold cross-validation. We used 5-fold cross-validation instead of LOSO validation method due to the computational cost of processing video. The physiological modality is stronger in the Biovid database. Many studies have validated the models on physiological modality. Our empirical results show that the EDA-only accuracy is 77.2%, whereas the visual-only accuracy is 72.9%. The proposed model can improve over unimodal performance and achieves state-of-the-art performance on the Biovid dataset. We also compare it with standard fusion techniques like feature concatenation. For a fair comparison, we keep all the parameters the same. The proposed model improves 6% over simple feature concatenation and 1.3% over a vanilla multimodal transformer i.e. without a joint representation.

Figure 4 shows the visualization of the attention weights generated by the joint transformer model. It can be seen that the model gives more weightage to the physiological modality.

Table 3 shows the valence and arousal CCC values on the Affwild2 official validation set and custom-defined folds to increase generalizability. On the official validation set, the proposed method achieves a 0.666 average with 0.717 va-

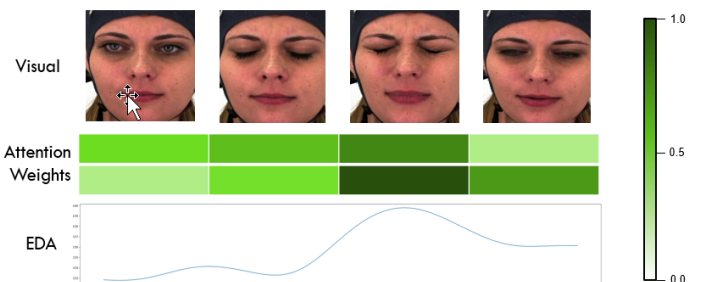


Figure 4. Visualization of attention weights for visual and physiological modalities on the Biovid heat pain database. The facial frames are taken 1400 msec each.

Table 2. A comparison of the proposed JMT and state-of-the-art fusion models for pain estimation on the Biovid Heat Pain Database. The highest score is indicated in bold.

Method	Modality	XV Scheme	Accuracy (%)
Werner et al. [51] <i>ICPR 2014</i>	EDA, ECG, EMG + Video	5-FOLD	80.6
Werner et al. [52] <i>IEEE TAC 2016</i>	Video	LOSO	72.4
Kachele et al. [53] <i>IEEE IJSTSP 2016</i>	EDA, ECG, EMG	LOSO	82.73
Lopez et al. [54] <i>ACIWI 2017</i>	EDA, ECG	10-FOLD	82.75
Lopez et al. [55] <i>EMBC 2018</i>	EDA	LOSO	74.21
Thiam et al. [56] <i>Sensors 2019</i>	EDA	LOSO	84.57
Wang et al. [57] <i>EMBC 2020</i>	EDA, ECG, EMG	LOSO	83.3
Pouromran et al. [58] <i>PLoS ONE 2021</i>	EDA	LOSO	83.3
Thiam et al. [59] <i>Frontiers 2021</i>	EDA, ECG, EMG	LOSO	84.25
Phan et al. [60] <i>IEEE Access 2023</i>	EDA, ECG, EMG	LOSO	84.8
Audio backbone: 1D CNN only	EDA	5-FOLD	77.2
Visual backbone: R3D model only	Video	5-FOLD	72.9
Fusion: feature concatenation	EDA + Video	5-FOLD	83.5
Fusion: vanilla transformer	EDA + Video	5-FOLD	87.8
Fusion: JMT (ours)	EDA + Video	5-FOLD	89.1

lence and 0.614 arousal. In addition, we performed ensemble by considering a dual modeling. We observed that different training configurations lead to models that can be best at once case: 'Valence', or 'Arousal'. Therefore, we used a dual model for prediction by taking the best at each category. This lead to performance increase.

Table 3. CCC for valence and arousal of the fusion model trained on different folds of the Affwild2 validation set. Highest scores are indicated in bold.

Validation Set	Valence	Arousal	Average
Official	0.717	0.614	0.666
fold-1	0.705	0.683	0.694
fold-2	0.741	0.623	0.682
fold-3	0.657	0.637	0.647
fold-4	0.760	0.666	0.713
fold-5	0.684	0.629	0.657
Ensembling	0.769	0.692	0.731

Table 4 compares the proposed JMT fusion method against state-of-the-art fusion methods on the Affwild2 test set. The proposed method can achieve 0.472 valence and 0.443 arousal. The average of valence and arousal is 0.458. The proposed method significantly improves over the baseline (provided by the challenge organizers). It is essential to mention here that the other methods that achieve higher performance are due to extensive pretraining, the use of additional modalities like text, and the use of more robust back-

bones. For a fair comparison, we use a similar setting to Joint Cross Attention [61], which includes similar pretraining and identical backbones for audio and visual modalities. They used joint cross-attention to fuse the two modalities and achieve a 0.369 average. On the other hand, the proposed model uses JMT fusion and can achieve a 0.443 average. The proposed method improves 7% over the joint cross-attention-based method.

Table 4. CCC performance of the proposed JMT and state-of-the-art methods for A-V fusion on the Affwild2 test set.

Method	Valence	Arousal	Mean
Baseline	0.180	0.170	0.175
Joint Cross Attention [61]	0.374	0.363	0.369
AU-NO [62]	0.418	0.407	0.413
HSE-NN [63]	0.417	0.454	0.436
PRL [64]	0.450	0.445	0.448
Joint Multimodal Transformer (ours)	0.472	0.443	0.458
FlyingPigs [65]	0.520	0.602	0.561
Situ-RUCAIM3 [66]	0.606	0.596	0.601

5.2. Ablations

Table 5 shows that the proposed model improves performance over the vanilla multimodal transformer by 1.8%

Table 5. CCC performance of the backbones alone, and combined using baselines and the proposed A-V fusion method on the Affwild2 test set. Experiments are performed on the default training-validation split. R2D1 (visual) and ResNet18 (audio) backbones are used for all cases.

Fusion Model	Valence	Arousal	Mean
R2D1 model only (visual)	0.194	0.310	0.252
ResNet18 model only (audio)	0.273	0.246	0.260
Concat + FC layers	0.320	0.327	0.323
Vanilla Transformer	0.376	0.334	0.355
JMT (ours)	0.366	0.379	0.373

Table 6. CCC performance of visual and audio backbones alone, and their A-V fusion using the proposed and baseline methods on the Affwild2 test set. Experiments are performed on the default training-validation split. I3D (visual) and ResNet18 (audio) backbones are used for all cases. FC: fully connected layers for feature-level fusion of backbones from the same modality. TR: transformer for fusing features from backbones of the same modality.

Fusion Model	Valence	Arousal	Mean
Joint Cross Attention [61]	0.374	0.363	0.369
I3D model only (visual)	0.336	0.422	0.379
ResNet18 model only (audio)	0.273	0.246	0.260
Concatenation + FC	0.387	0.453	0.420
Vanilla Transformer - I3D (visual) - ResNet18 (audio)	0.432	0.410	0.421
JMT (ours) - I3D (visual) - ResNet18 (audio)	0.425	0.450	0.438
JMT (ours) - I3D+R2D1 (visual, FC) - ResNet18 (audio)	0.472	0.443	0.458
JMT (ours) - I3D+R2D1 (visual, TR) - ResNet18 (audio)	0.458	0.445	0.452

with an R2D1 and ResNet18 backbone on the Affwild2 dataset. With an I3D and a ResNet18 backbone, results are improved by 1.7%, as shown in Table 6. Notice that the results on the test set are different from Table 4 because we used a different train/validation split. Additionally, results show that feature-level fusion of 2 vision backbones (I3D and R2D1) improves the performance compared to using a single vision backbone.

Table 7 shows the results of the proposed method with

and without the joint representation. On the Biovid dataset, the joint multimodal transformer improves by 1.3% over the vanilla multimodal transformer.

Table 7. Performance of the proposed approach compared to the vanilla multimodal transformer on test set. Valence and arousal for the Affwild2 dataset and accuracy for the Biovid dataset. I3D + ResNet18 backbones are used for the Affwild2 dataset and R3D + 1D CNN are used for the BioVid dataset. Default training/validation split is used in the Affwild2 dataset and 5-fold cross validation is performed on the BioVid dataset.

Database	Method	Accuracy
Affwild2	Vanilla Multimodal Transformer	V: 0.432 A: 0.410 Avg: 0.421
	Joint Multimodal Transformer	V: 0.425 A: 0.450 Avg: 0.438
Biovid	Vanilla Multimodal Transformer	87.8
	Joint Multimodal Transformer	89.1

6. Conclusion

Multimodal emotion recognition systems outperform their unimodal counterparts, especially in the wild environment. The missing and noisy modality is a prevalent issue with in-the-wild emotion recognition systems. Many attention-based methods have been proposed in the literature to overcome this problem. These methods aim to weigh the modalities dynamically. This paper introduces a joint multimodal transformer for emotional recognition. This transformer-based architecture introduces a joint feature representation to add more redundancy and complementary between audio and visual data. The two modalities are first encoded using separate backbones to extract intra-modal spatiotemporal dependencies. The feature vectors of the two modalities are joint and the joint feature vector is also fed into the Joint Multimodal Transformer module. This joint representation provides more fine-grained information about the inter-modal association between the two modalities. The proposed model outperforms state-of-the-art methods on the Biovid dataset and improves over the vanilla multimodal transformer by 6% on the Affwild2 dataset. Our future work includes introducing more modalities and sophisticated backbones for effective feature extraction.

Acknowledgements: This work was supported by the Fonds de recherche du Québec – Santé (FRQS), the Natural Sciences and Engineering Research Council of Canada (NSERC), and Canada Foundation for Innovation (CFI).

References

- [1] C. Anagnostopoulos, T. Iliou, and I. Giannoukos, “Features and classifiers for emotion recognition from speech: a survey from 2000 to 2011,” *Artif Intell Rev*, vol. 43, no. 6, pp. 155–177, 2015. **1**
- [2] P. Ekman, “An argument for basic emotions,” *Cognition and Emotion*, vol. 6, no. 3-4, pp. 169–200, 1992. **1**
- [3] L. Schoneveld, A. Othmani, and H. Abdelkawy, “Leveraging recent advances in deep learning for audio-visual emotion recognition,” *Pattern Recognition Letters*, vol. 146, pp. 1–7, 2021. **2**
- [4] F. Kuhnke, L. Rumberg, and J. Ostermann, “Two-stream aural-visual affect analysis in the wild,” in *FG Workshop*, pp. 600–605, 2020. **2**
- [5] R. G. Praveen, W. C. de Melo, N. Ullah, H. Aslam, O. Zee-shan, T. Denorme, M. Pedersoli, A. L. Koerich, S. Bacon, P. Cardinal, and E. Granger, “A joint cross-attention model for audio-visual fusion in dimensional emotion recognition,” in *CVPRW*, pp. 2485–2494, 2022. **2**
- [6] W. C. de Melo, E. Granger, and M. B. Lopez, “Mdn: A deep maximization-differentiation network for spatio-temporal depression detection,” *IEEE Tran. on Affective Computing*, 2021. **2**
- [7] G. P. Rajasekhar, E. Granger, and P. Cardinal, “Cross attentional audio-visual fusion for dimensional emotion recognition,” in *FG*, pp. 1–8, 2021. **2, 3**
- [8] J. Lu, D. Batra, D. Parikh, and S. Lee, “Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks,” *Advances in neural information processing systems*, vol. 32, 2019. **2**
- [9] S. Parthasarathy and S. Sundaram, “Training strategies to handle missing modalities for audio-visual expression recognition,” in *ICMI*, p. 400–404, 2020. **2, 3**
- [10] X. Wei, T. Zhang, Y. Li, Y. Zhang, and F. Wu, “Multi-modality cross attention network for image and sentence matching,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10938–10947, 2020. **2**
- [11] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, “Multimodal deep learning,” in *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML’11, (Madison, WI, USA)*, p. 689–696, Omnipress, 2011. **2**
- [12] P. Tzirakis, G. Trigeorgis, M. A. Nicolaou, B. W. Schuller, and S. Zafeiriou, “End-to-end multimodal emotion recognition using deep neural networks,” *IEEE J. of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1301–1309, 2017. **2**
- [13] J. D. S. Ortega, P. Cardinal, and A. L. Koerich, “Emotion recognition using fusion of audio and video features,” in *SMC*, pp. 3847–3852, 2019. **2**
- [14] D. Nguyen, D. T. Nguyen, R. Zeng, T. T. Nguyen, S. Tran, T. K. Nguyen, S. Sridharan, and C. Fookes, “Deep auto-encoders with sequential learning for multimodal dimensional emotion recognition,” *IEEE Trans. on Multimedia*, pp. 1–1, 2021. **2**
- [15] D. Deng, L. Wu, and B. E. Shi, “Iterative distillation for better uncertainty estimates in multitask emotion recognition,” in *ICCVW*, pp. 3550–3559, 2021. **2**
- [16] P. Werner, A. Al-Hamadi, R. Niese, S. Walter, S. Gruss, and H. C. Traue, “Automatic pain recognition from video and biomedical signals,” in *2014 22nd International Conference on Pattern Recognition*, pp. 4582–4587, 2014. **2**
- [17] M.-C. Dragomir, C. Florea, and V. Pupezescu, “Automatic subject independent pain intensity estimation using a deep learning approach,” in *2020 International Conference on e-Health and Bioengineering (EHB)*, pp. 1–4, 2020. **2**
- [18] R. Zhi, C. Zhou, J. Yu, T. Li, and G. Zamzmi, “Multimodal-based stream integrated neural networks for pain assessment,” *IEICE Trans. Inf. Syst.*, vol. 104-D, pp. 2184–2194, 2021. **2, 3**
- [19] S. E. Morabit and A. Rivenq, “Pain detection from facial expressions based on transformers and distillation,” in *2022 11th International Symposium on Signal, Image, Video and Communications (ISIVC)*, pp. 1–5, 2022. **2**
- [20] Z. Lu, B. Ozek, and S. Kamarthi, “Transformer encoder with multiscale deep learning for pain classification using physiological signals,” *arXiv preprint arXiv:2303.06845*, 2023. **2**
- [21] M. Kächele, P. Werner, A. Al-Hamadi, G. Palm, S. Walter, and F. Schwenker, “Bio-visual fusion for person-independent recognition of pain intensity,” in *International Workshop on Multiple Classifier Systems*, 2015. **2**
- [22] J.-T. Lee, M. Jain, H. Park, and S. Yun, “Cross-attentional audio-visual fusion for weakly-supervised action localization,” in *ICLR*, 2021. **3**
- [23] B. Duan, H. Tang, W. Wang, Z. Zong, G. Yang, and Y. Yan, “Audio-visual event localization via recursive fusion by joint co-attention,” in *WACV*, pp. 4012–4021, 2021. **3**
- [24] Y.-H. Zhang, R. Huang, J. Zeng, and S. Shan, “Multi-modal continuous valence-arousal estimation in the wild,” in *IEEE FG*, pp. 632–636, 2020. **3**
- [25] P. Tzirakis, J. Chen, S. Zafeiriou, and B. Schuller, “End-to-end multimodal affect recognition in real-world environments,” *Information Fusion*, vol. 68, pp. 46–53, 2021. **3**
- [26] S. Zhang, Y. Ding, Z. Wei, and C. Guan, “Continuous emotion recognition with audio-visual leader-follower attentive fusion,” in *ICCVW*, pp. 3560–3567, 2021. **3**
- [27] J. Huang, J. Tao, B. Liu, Z. Lian, and M. Niu, “Multi-modal transformer fusion for continuous emotion recognition,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3507–3511, 2020. **3**
- [28] M. Tran and M. Soleymani, “A pre-trained audio-visual transformer for emotion recognition,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4698–4702, 2022. **3**
- [29] H.-D. Le, G.-S. Lee, S.-H. Kim, S. Kim, and H.-J. Yang, “Multi-label multimodal emotion recognition with transformer-based fusion and emotion-level representation learning,” *IEEE Access*, vol. 11, pp. 14742–14751, 2023. **3**
- [30] W. Zhou, J. Lu, Z. Xiong, and W. 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)*, pp. 5756–5763, 2023. **3**
- [31] F. Kuhnke, L. Rumberg, and J. Ostermann, “Two-stream aural-visual affect analysis in the wild,” in *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pp. 600–605, IEEE, 2020. **3**
- [32] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016. **3**
- [33] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning phrase representations using rnn encoder-decoder for statistical machine translation,” *arXiv preprint arXiv:1406.1078*, 2014. **3**
- [34] I. Lawrence and K. Lin, “A concordance correlation coefficient to evaluate reproducibility,” *Biometrics*, pp. 255–268, 1989. **4**
- [35] 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*, vol. 127, p. 907–929, 2019. **5**
- [36] M. H. Aslam, M. O. Zeeshan, S. Belharbi, M. Pedersoli, A. L. Koerich, S. Bacon, and E. Granger, “Distilling privileged multimodal information for expression recognition using optimal transport,” in *18th IEEE International Conference on Automatic Face and Gesture Recognition*, arXiv preprint arXiv:2401.15489, 2024. **5**
- [37] D. Kollias, P. Tzirakis, A. Cowen, S. Zafeiriou, C. Shao, and G. Hu, “The 6th affective behavior analysis in-the-wild (abaw) competition,” *arXiv preprint arXiv:2402.19344*, 2024.
- [38] D. Kollias, P. Tzirakis, A. Baird, A. Cowen, and S. 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*, pp. 5888–5897, 2023.
- [39] D. Kollias, “Multi-label compound expression recognition: C-expr database & network,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5589–5598, 2023.
- [40] D. Kollias, “Abaw: learning from synthetic data & multi-task learning challenges,” in *European Conference on Computer Vision*, pp. 157–172, Springer, 2023.
- [41] M. H. Aslam, O. Zeeshan, M. Pedersoli, A. L. Koerich, S. Bacon, and E. Granger, “Privileged knowledge distillation for dimensional emotion recognition in the wild,” in *CVPRW 2023: IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, Canada.*, 2023.
- [42] D. Kollias, “Abaw: Valence-arousal estimation, expression recognition, action unit detection & multi-task learning challenges,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2328–2336, 2022.
- [43] D. Kollias and S. Zafeiriou, “Analysing affective behavior in the second abaw2 competition,” in *ICCVW*, pp. 3652–3660, 2021. **5**
- [44] D. Kollias and S. Zafeiriou, “Affect analysis in-the-wild: Valence-arousal, expressions, action units and a unified framework,” *arXiv preprint arXiv:2103.15792*, 2021.
- [45] D. Kollias, V. Sharmanska, and S. Zafeiriou, “Distribution matching for heterogeneous multi-task learning: a large-scale face study,” *arXiv preprint arXiv:2105.03790*, 2021.
- [46] D. Kollias, A. Schulc, E. Hajiyev, and S. Zafeiriou, “Analysing affective behavior in the first abaw competition,” in *FG 2020*, 2020.
- [47] D. Kollias and S. Zafeiriou, “Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arcface,” *arXiv preprint arXiv:1910.04855*, 2019.
- [48] 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,” *International Journal of Computer Vision*, pp. 1–23, 2019.
- [49] D. Kollias, V. Sharmanska, and S. Zafeiriou, “Face behavior a la carte: Expressions, affect and action units in a single network,” *arXiv preprint arXiv:1910.11111*, 2019.
- [50] S. Zafeiriou, D. Kollias, M. A. Nicolaou, A. Papaioannou, G. Zhao, and I. Kotsia, “Aff-wild: Valence and arousal ‘in-the-wild’ challenge,” in *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pp. 1980–1987, IEEE, 2017. **5**
- [51] P. Werner, A. Al-Hamadi, R. Niese, S. Walter, S. Gruss, and H. C. Traue, “Automatic pain recognition from video and biomedical signals,” in *2014 22nd international conference on pattern recognition*, pp. 4582–4587, IEEE, 2014. **5, 7**
- [52] P. Werner, A. Al-Hamadi, K. Limbrecht-Ecklundt, S. Walter, S. Gruss, and H. C. Traue, “Automatic pain assessment with facial activity descriptors,” *IEEE Transactions on Affective Computing*, vol. 8, no. 3, pp. 286–299, 2016. **7**
- [53] M. Kächele, P. Thiam, M. Amirian, F. Schwenker, and G. Palm, “Methods for person-centered continuous pain intensity assessment from bio-physiological channels,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 5, pp. 854–864, 2016. **7**
- [54] D. Lopez-Martinez and R. Picard, “Multi-task neural networks for personalized pain recognition from physiological signals,” in *2017 Seventh International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, pp. 181–184, IEEE, 2017. **7**
- [55] D. Lopez-Martinez and R. Picard, “Continuous pain intensity estimation from autonomic signals with recurrent neural networks,” in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 5624–5627, IEEE, 2018. **7**
- [56] P. Thiam, P. Bellmann, H. A. Kestler, and F. Schwenker, “Exploring deep physiological models for nociceptive pain recognition,” *Sensors*, vol. 19, no. 20, p. 4503, 2019. **7**
- [57] R. Wang, K. Xu, H. Feng, and W. Chen, “Hybrid rnn-ann based deep physiological network for pain recognition,” in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 5584–5587, IEEE, 2020. **7**
- [58] F. Pououmran, S. Radhakrishnan, and S. Kamarthi, “Exploration of physiological sensors, features, and machine

- learning models for pain intensity estimation,” *Pone Journal*, vol. 16, no. 7, p. e0254108, 2021. 7
- [59] P. Thiam, H. Hihn, D. A. Braun, H. A. Kestler, and F. Schwenker, “Multi-modal pain intensity assessment based on physiological signals: A deep learning perspective,” *Frontiers in Physiology*, vol. 12, p. 720464, 2021. 7
- [60] K. N. Phan, N. K. Iyortsuun, S. Pant, H.-J. Yang, and S.-H. Kim, “Pain recognition with physiological signals using multi-level context information,” *IEEE Access*, vol. 11, pp. 20114–20127, 2023. 7
- [61] R. G. Praveen, W. C. de Melo, N. Ullah, H. Aslam, O. Zee-shan, T. Denorme, M. Pedersoli, A. L. Koerich, S. Bacon, P. Cardinal, *et al.*, “A joint cross-attention model for audio-visual fusion in dimensional emotion recognition,” in *IEEE/CVF CVPR*, pp. 2486–2495, 2022. 7, 8
- [62] V. Karas, M. K. Tellamekala, A. Mallol-Ragolta, M. Valstar, and B. W. Schuller, “Continuous-time audiovisual fusion with recurrence vs. attention for in-the-wild affect recognition,” *arXiv preprint arXiv:2203.13285*, 2022. 7
- [63] A. V. Savchenko, “Frame-level prediction of facial expressions, valence, arousal and action units for mobile devices,” *arXiv preprint arXiv:2203.13436*, 2022. 7
- [64] H.-H. Nguyen, V.-T. Huynh, and S.-H. Kim, “An ensemble approach for facial expression analysis in video,” *arXiv preprint arXiv:2203.12891*, 2022. 7
- [65] S. Zhang, R. An, Y. Ding, and C. Guan, “Continuous emotion recognition using visual-audio-linguistic information: A technical report for abaw3,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2376–2381, 2022. 7
- [66] L. Meng, Y. Liu, X. Liu, Z. Huang, W. Jiang, T. Zhang, Y. Deng, R. Li, Y. Wu, J. Zhao, *et al.*, “Multi-modal emotion estimation for in-the-wild videos,” *arXiv preprint arXiv:2203.13032*, 2022. 7