

# Multimodal Feature Extraction and Fusion for Emotional Reaction Intensity Estimation and Expression Classification in Videos with Transformers

Jia Li\*, Yin Chen\*, Xuesong Zhang, Jiantao Nie, Ziqiang Li, Yangchen Yu, Yan Zhang<sup>§</sup>,  
 Richang Hong<sup>†</sup>, Meng Wang<sup>†</sup>  
 Hefei University of Technology, China  
<sup>§</sup>Anhui University, China

jiali@hfut.edu.cn

{chenyin, xszhang\_hfut, 2022111086, 2019212292}@mail.hfut.edu.cn

zhangyan@ahu.edu.cn

{luoyixuan1029, eric.mengwang, hongrc.hfut}@gmail.com

## Abstract

*In this paper, we present our advanced solutions to the two sub-challenges of Affective Behavior Analysis in the wild (ABAW) 2023: the Emotional Reaction Intensity (ERI) Estimation Challenge and Expression (Expr) Classification Challenge. ABAW 2023 aims to tackle the challenge of affective behavior analysis in natural contexts, with the ultimate goal of creating intelligent machines and robots that possess the ability to comprehend human emotions, feelings, and behaviors. For the Expression Classification Challenge, we propose a streamlined approach that handles the challenges of classification effectively. However, our main contribution lies in our use of diverse models and tools to extract multimodal features such as audio and video cues from the Hume-Reaction dataset. By studying, analyzing, and combining these features, we significantly enhance the model's accuracy for sentiment prediction in a multimodal context. Furthermore, our method achieves outstanding results on the Emotional Reaction Intensity (ERI) Estimation Challenge, surpassing the baseline method by an impressive 84% increase, as measured by the Pearson Coefficient, on the validation dataset.*

## 1. Introduction

Emotions are psychological states that arise in human beings when stimulated by the external environment, and they are able to reflect a person's current physiological and psychological state, making an impact on people's perception of things and future decisions. Emotions are often ac-

companied by distinct physiological characteristics, which can be analyzed to identify them. In this paper, we present in detail our solution to the Expression Classification Challenge and Emotional Reaction Intensity (ERI) Estimation Challenge.

The Expression Classification Challenge dataset comprises 546 videos that contain annotations for the six basic expressions (i.e., anger, disgust, fear, happiness, sadness, surprise), along with a neutral state and a category for other expressions/affective states. This challenge requires outputting the expression type for each frame.

In the Emotional Reaction Intensity (ERI) Estimation Challenge, Hume-Reaction dataset was used. Each participant subjectively annotated his or her intensity of 7 emotional experiences [3], including adoration, amusement, anxiety, disgust, empathic pain, fear, and surprise. The model is required to estimate the intensity of the participants' emotional experiences caused by watching the video. The model hence needs to analyze the visual and audio information present in the video and make accurate predictions about the emotional state of the participants. The task is challenging because the model needs to correctly identify the emotions that the participants are experiencing and measure their intensity while taking into account the context of the video. Participants tend to express their emotions through their voices, body movements, and expressions, and it is through these factors that multimodal emotion recognition can identify and predict emotions.

Traditional emotion recognition is based on a single modality, however, when a single factor is affected, for example, when facial expressions are not clear due to sunlight or when voices are masked by ambient sounds, the effectiveness of emotion prediction is significantly reduced. To further augment its feature representation capacity, We first

\*Equal contribution

†Corresponding author

pre-train the model on the AffectNet7 dataset [22] using the Posterv2 [21] and ViT [6] models, respectively. Then, we combine these pre-trained models to obtain a novel pre-trained PosterV2-ViT model that is specifically designed for extracting features on the Hume-Reaction dataset. This approach significantly improves the performance of model and enables us to achieve state-of-the-art results in this Challenge. Furthermore, to fully harness the potential of audio and video features, we incorporated a modal interaction module to effectively extract and capture the common relationships between these two modalities. Our contributions in enhancing feature extraction ability and modality interaction have significant implications for emotion recognition in unstable environments with varied stimuli.

Our contributions can be summarized as follows.

1. We propose a cross modality interaction approach to combine audio and visual features. Specifically, we extract visual features using a novel PosterV2-ViT model, which significantly improves the performance of our model.

2. For the emotional reaction intensity estimation challenge, our method achieves a Mean Person score of 0.4394 on the Hume-Reaction validation set. The experiment result also suggests that our method is effective in capturing the inter-person variations in emotional reactions.

3. For the expression classification challenge, our simple approach achieves 0.346 for the F1 score on the Aff-Wild2 validation set.

## 2. Related Work

Human emotions can be realistically reflected by facial expressions in most cases, so recognizing human expressions and estimating emotional reaction intensity are crucial for human-computer interaction (HCI) systems [9, 10, 12–14, 17, 18, 34]. In this section, we will introduce some related works about emotion recognition and emotional reaction intensity.

### 2.1. Emotional Reaction Intensity

Emotion recognition research has primarily focused on the classification of basic emotions, including positive, negative, and neutral states. However, there have been recent efforts [8, 19, 28] to differentiate more nuanced emotional categories such as anger, happiness, and sadness.

Building upon emotion recognition, some studies attempt to estimate the intensity or degree of emotional reactions. Emotional Reaction Intensity Estimation Challenge were presented in MuSe 2022 [1] and ABAW5 [15]. The objective of this challenge was to assess the magnitude of response elicited by seven distinct expression types, encompassing adoration, amusement, anxiety, disgust, empathic-pain, fear and surprise, with intensity values between 0 to 100.

The TEMMA model proposed by Li [20] at Muse 2022 won the championship in this challenge. Resnet-18 feature and DeepSpectrum feature were used in [20]. The researchers [20] use a Resnet-18 model to extract static visual features from the videos, and a DeepSpectrum [2] model to extract audio features from the videos. These features are then input into the model as training data. Among them, the DeepSpectrum model is a spectral analysis algorithm that extracts audio features from the videos using the frequency domain. By combining the visual and audio features, the model is able to make more accurate predictions about the expression type and intensity of the participants' emotional experiences.

In addition, there were many exceptional contributions to the competition that emerged, and the methods used by these teams also achieved remarkable results. FaceRNET [11] is a convolutional recurrent neural network with a routing mechanism atop, which yield excellent results on the Hume Reaction dataset. They [11] used a representation extractor network to extract emotion descriptors, and a RNN network to capture temporal information with a mask layer. Former-DFER+MLGCN [32] also proposes an end-to-end network framework that has achieved remarkable performance on the Hume-Reaction dataset. To better utilize multimodal information, ViPER [30] proposed a multimodal architecture for emotion recognition which achieves excellent performance on the ERI task as well.

### 2.2. Expression Classification

In the field of facial expression recognition, numerous datasets and research methods have been proposed. In the term of dataset, VGGFace [23] and CASIA-WebFace that are widely used for training and evaluating deep learning models in face recognition. The Expression Classification Challenge is based on the Aff-Wild2 dataset which is a naturalistic affective behavior analysis dataset. It aims to provide high-quality facial data for expression analysis. As for the term of research methods, the champion in expression classification challenge is dedicate in Multimodal feature extraction in the ABAW3 Competition. They make static vision feature extractor pre-trained and integrate the static and dynamic multimodal feature [36]. As the winner of expression classification challenge in ABAW5, Zhang et al [35] introduced a two-branch collaboration training strategy that randomly interpolates the logits space.

## 3. Methodology

### 3.1. Emotional Reaction Intensity Estimation

#### 3.1.1 Model Structure

Figure 1 illustrates the overall framework of our network architecture, which consists of three main branches: an audio

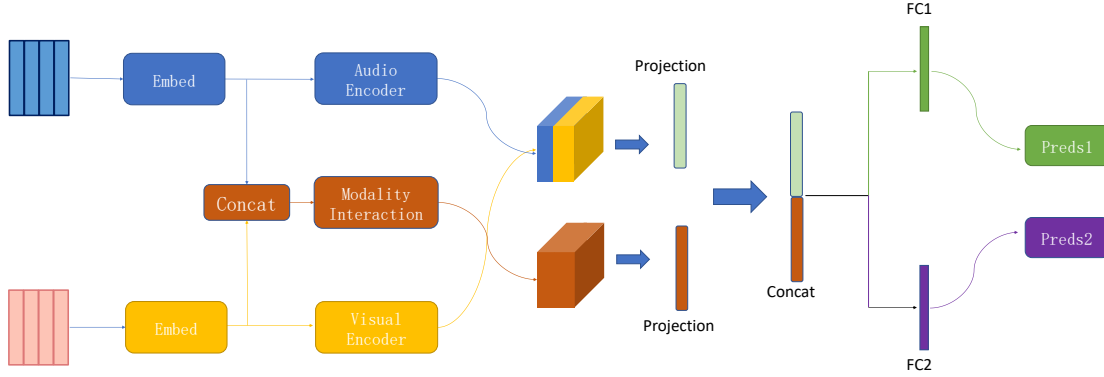


Figure 1. The framework of our network architecture. It consists of three main branches, i.e., an audio feature encoding module, a visual feature encoding module, and an audio-visual modality interaction module. Each module comprises a single-layer transformer encoder.

feature encoding module, a visual feature encoding module, and an audio-visual modality interaction module. Each module comprises a single-layer transformer encoder.

Firstly, the extracted audio and visual features were passed through an embedding layer (Figure 1 Embed Box) to unify the dimensions of the input. The embedding layer we used is 1-dimensional convolution.

$$h_m^d = \text{Embed}(\text{input}_m), m \in \{a, v\} \quad (1)$$

The function  $\text{Embed}$  in Equation 1 means the embedding layer, variable  $\text{input}_m$  denotes the input features, where  $m$  represents different modalities,  $a$  refers to audio,  $v$  refers to visual.  $h$  refers to the embedded features, and  $d$  represents the dimension of the output features.

Next, we passed the extracted audio and visual features through separate feature encoders to extract further features. Additionally, we concatenated the audio and visual features and passed them through a modal interaction encoder to extract features between the two modalities:

$$\text{att} = \text{MHA}(x) \quad (2)$$

$$\text{output} = \text{concat}(x, \text{att}, \text{trans}(x)) \quad (3)$$

Equation 2, 3 define the processing procedure of the Encoder in Figure 1, while  $x$  is the input vector,  $\text{MHA}$  is Multi-Head Attention [31] and  $\text{trans}$  is the transformer [31] encoder layer.

$$g_m^{3d} = \text{Encoder}_m(h_m^d), m \in \{a, v\} \quad (4)$$

$$h_{av}^{2d} = \text{concat}(h_a, h_v) \quad (5)$$

$$g_{av}^{6d} = \text{Encoder}_{av}(h_{av}^{2d}) \quad (6)$$

Here,  $\text{Encoder}_m$  represents the modality encoder, and  $g_m$  represents the output features of the modality encoder. We

first concatenate  $a, v$  features then encode them with modality interaction encoder. This procedure is defined in equation 4,5,6.

Afterwards, we concatenated the audio and visual modality encoder features to create  $g_{cat}$ .

$$g_{cat}^{6d} = \text{concat}(g_a^{3d}, g_v^{3d}) \quad (7)$$

Finally, we projected the modality features  $g_{av}$  and  $g_{cat}$  to a lower dimension using a projection layer, and concatenated the output features before passing them through a fully connection layer separately to obtain the final results.

### 3.1.2 Feature Extraction

In our solution for ERI, we utilized DeepSpectrum [2] audio features and PosterV2-ViT visual features as our input features, both of which yielded excellent results.

DeepSpectrum [2] refers to a 1024 dimensional feature extracted from speech signals by deep neural networks. We utilized DenseNet121 [7], trained on ImageNet [27], for feature extraction of DeepSpectrum. The visual features were extracted using the pretrained model PosterV2-ViT on the AffectNet7 dataset, resulting in a feature vector of 1536 dimensions. The PosterV2-ViT model was obtained from the joint training of the PosterV2 [21] and ViT models after individually pretraining on the AffectNet7 dataset. The final model achieved an accuracy of 67.57% on the AffectNet7 test set.

### 3.1.3 Loss Functions

Estimating Emotional Reaction Intensity is a multi-regression task. However, based on observations from our data annotation, it was found that each label consistently has a value of 1 for intensity. Therefore, we decoupled this intensity estimation task into two separate tasks: one for

multi-regression and the other for classification. As a result, we utilized two different loss functions: mean squared error (MSE) loss and cross entropy loss.

$$Loss_{ERI} = \alpha L_{reg} + \beta L_{class} \quad (8)$$

As shown in Equation 9, the loss function for regression tasks using MSE loss, and for classification tasks using cross entropy loss, can be denoted as  $L_{reg}$  and  $L_{class}$ , respectively.  $\alpha$  and  $\beta$  are the loss weight for  $L_{reg}$  and  $L_{class}$ .

### 3.2. Expression Classification

Expression classification and emotional reaction intensity estimation are two very different tasks that are based on different datasets. Therefore, we propose a new but simple method to handle the expression classification task.

To enhance the robustness of the model on the expression classification challenge, we first obtain different pre-trained models on different face datasets. Specifically, we pre-trained the InceptionResnetV1 model [29] on the casia-webface [33] and VGGface2 datasets [24], respectively. We use the Regnet model [25, 26] pre-trained on the imageNet dataset to improve the generalization ability of the model [5].

Transformed-based methods are usually effective in capturing attention information of modality [31]. Therefore, we stitch these features and then feed them uniformly into a transformed-based feature fusion module for training.

Deng et al. [4] showed that there is a general category imbalance in the dataset. Focal loss rebuilding the standard cross entropy loss to solve class imbalance, which can be defined as :

$$Loss_{Expr} = -\alpha_t(1 - p_t)^\gamma \log(p_t) \quad (9)$$

where parameter  $\alpha_t \geq 0$  and  $\gamma \in [0, 5]$

## 4. Experiments

### 4.1. Datasets

Hume-Reaction [3] dataset: This multi-modal dataset comprises approximately 75 hours of video recordings, captured via a webcam located in the participants' homes. The dataset consists of recordings of 2222 participants belonging to two different cultures, South Africa and the United States. Each sample within the dataset has been self-annotated by the participants, indicating the intensity of 7 emotional experiences, including Adoration, Amusement, Anxiety, Disgust, Empathic Pain, Fear, and Surprise, on a scale from 1 to 100.

Aff-Wild2 database [16]: For the Expression Classification Challenge, the Aff-Wild2 database will be used. About 2.7M frames in this database are annotated in terms of the 6 basic expressions (i.e., anger, disgust, fear, happiness, sadness, surprise), plus the neutral state and a category 'other' that denotes expressions states other than the 6 basic ones.

Table 1. The Mean Person score of different modality features compare with former works on the Hume-Reaction validation set.

Feature	Modality	P
DeepSpectrum(Baseline [15])	A	0.1087
DeepSpectrum(TEMMA [20])	A	0.1835
DeepSpectrum(Ours)	A	<b>0.2299</b>
VGGface2(Baseline [15])	V	0.2488
Resnet-18(TEMMA [20])	V	0.3893
Former-DFER+MLGCN [32]	V	0.3454
ViPER [30]	V	0.2978
FaceRNET [11]	V	0.3590
PosterV2+Vit(Ours)	V	<b>0.3924</b>
Baseline [15]	A+V	0.2382
ViPER [30]	A+V	0.3025
DeepSpectrum+Resnet-18(TEMMA [20])	A+V	0.3968
DeepSpectrum + PosterV2-Vit(Ours)	A+V	<b>0.4377</b>
DeepSpectrum + PosterV2-Vit(Ours)	A+V+AV	<b>0.4394</b>
Ensemble(Ours)	A+V+AV	<b>0.4452</b>

### 4.2. Evaluation Metrics

For the expression classification challenge, we calculate the F1-Score for every class to assess the prediction results. With regards to the intensity estimation of emotional reactions, we employ Pearson's Correlation Coefficient as the metric for each class.

### 4.3. Experimental Setup

Throughout the experiment, we adopt the Adam optimizer with a fixed learning rate, which is 1e-4. Additionally, we employed an early stopping technique, which stopped the training process when the metric failed to increase for 15 consecutive epochs. The evaluation metric being used is Pearson's correlation coefficient (p). The  $\alpha$  and  $\beta$  in equation 9 was set 1.0 and 0.01, respectively. Moreover, we used Exponential Moving Average to improve the model's generalization ability, with a decay rate of 0.99. To prevent the gradient explosion issue that could negatively impact the model optimization, we applied gradient clipping with max\_norm=0.1 and norm\_type=2.

In the experiment, we set the dimension of the embedding to 256, the depth of the modality encoder to 1, the number of attention heads to 2, and the dropout to 0.5. Moreover, the depth of the modality interaction module was set to 1, the number of attention heads to 8, and the dropout to 0.1. During the training, the batch size is set to 256, and trained for 100 epochs.

For some experimental settings for expression classification challenge, please refer to the baseline [25] we adopted. We used Pytorch framework to conducted all experiments on a NVIDIA RTX 3090 GPU.

Table 2. Effects of our different tricks on the Hume-Reaction validation set.

Model	Class loss	EMA	L2 Penalty	P
Ours	✓	✓	✓	0.4394
Ours	✗	✓	✓	0.4306
Ours	✓	✗	✓	0.4281
Ours	✓	✓	✗	0.4165

Table 3. The final results on the Hume-Reaction test set. Numbers are provided by the web of Challenge Organization Committee.

Team	PCC Score (Combined)
<b>Ours</b>	<b>0.4734</b>
USTC-IAT-United	0.4380
Netease Fuxi Virtual Human	0.4046
SituTech	0.3935
CASIA-NLPR	0.3865
USTC-AC	0.3730
NISL-2023	0.3667
HFUT-MAC	0.2527
IXLAB	0.1789

## 4.4. Experimental Results

### 4.4.1 Emotional Reaction Intensity Estimation

We conducted experiments on features from different modalities and compare the result with the baseline and Li. [20] presented in MuSe 2022, the results are shown in Table 2. We found that visual modality features were highly important, with visual scores significantly higher than audio features. After fusing the features from both modalities, model performance was further improved, reaching a Mean Person score of 0.4394 on the validation set.

We further investigated the impact of the proposed tricks on the final performance. The removal of class loss resulted in a performance decrease of 0.0088. The absence of EMA led to a performance drop of 0.0113. The removal of L2 regularization also decreased model performance by 0.0229.

Table 3 shows our final results on the test set of Hume-Reaction, and our method won the championship of this competition. The PCC score on the Emotional Reaction Intensity (ERI) Estimation Challenge is significantly higher than the runner-up by 0.0354.

### 4.4.2 Expression Classification

For comparison with the baseline model, we conducted experiments our approach in the section 3.2 and further investigated the performance of three single pre-trained models.

As shown in Table 4, InceptionResnetV1-1 denotes

Table 4. The Mean F1 score of different pre-trained methods on the Aff-Wild2 validation set. Ensemble indicates that all three pre-trained models are averaged.

Method	F1 score
Baseline(VGGface2)	0.23
Regnet [25]	0.3035
InceptionResnetV1-1(Ours)	0.339
InceptionResnetV1-2(Ours)	0.284
<b>Ensemble(Ours)</b>	<b>0.346</b>

Table 5. The final results on the test set of Expression Classification Challenge. Numbers are provided by the web of Challenge Organization Committee.

Team	F1 Score
Netease Fuxi Virtual Human	0.4121
SituTech	0.4072
CtyunAI	0.3532
HFUT-MAC	0.3337
HSE-NN-SberAI	0.3292
AlphaAff	0.3218
USTC-IAT-United	0.3075
SSSIHL DMACS	0.3047
DGU-IPL	0.2278
<b>HFUT-CVers (Ours)</b>	<b>0.2277</b>
baseline	0.2050

pre-trained InceptionResnetV1 on the vggface2 dataset, InceptionResnetV1-2 denotes pre-trained InceptionResnetV1 on the casia-webface dataset. The experiment results indicated that the single pre-trained Regnet [25] and baseline model achieved F1score of 0.3035 and 0.23 on validation set of the Aff-Wild2 database, respectively. However, our approach achieves an F1 score of 0.346, which is significantly better than the performance of the baseline model.

## 5. Conclusion

In this paper, we introduced our proposed approach to solving both Expression Classification and Emotional Reaction Intensity Estimation Challenges in ABAW5. For the Expression Classification Challenge, our approach achieves 0.346 for the F1 score on the validation set of Aff-Wild2 database. For the Emotional Reaction Intensity Estimation challenge, our model shows excellent results with a Pearson coefficient on the validation set that exceeds the baseline method by 83%, obtaining a Mean Person score of 0.4394 on the validation set. Furthermore, we found that the use of class loss, implementation of EMA, incorporation of regularization, and utilization of gradient clipping can greatly improve model performance. During the ABAW5 compe-

tion, we won the first place in the Emotional Reaction Intensity Estimation track.

## 6. Acknowledgments

This work is supported by the National Natural Science Foundation of China (62202139, 62020106007), the Anhui Provincial Natural Science Foundation (2208085QF191), and The University Synergy Innovation Program of Anhui Province (GXXT-2022-038).

## References

- [1] Shahin Amiriparian, Lukas Christ, Andreas König, Eva-Maria Meßner, Alan Cowen, Erik Cambria, and Björn W Schuller. Muse 2022 challenge: Multimodal humour, emotional reactions, and stress. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 7389–7391, 2022. 2
- [2] Shahin Amiriparian, Maurice Gerczuk, Sandra Ottil, Nicholas Cummins, Michael Freitag, Sergey Pugachevskiy, Alice Baird, and Björn Schuller. Snore sound classification using image-based deep spectrum features. 2017. 2, 3
- [3] Lukas Christ, Shahin Amiriparian, Alice Baird, Panagiotis Tzirakis, Alexander Kathan, Niklas Müller, Lukas Stappen, Eva-Maria Meßner, Andreas König, Alan Cowen, et al. The muse 2022 multimodal sentiment analysis challenge: humor, emotional reactions, and stress. In *Proceedings of the 3rd International on Multimodal Sentiment Analysis Workshop and Challenge*, pages 5–14, 2022. 1, 4
- [4] Didan Deng, Zhaokang Chen, and Bertram E Shi. Multitask emotion recognition with incomplete labels. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 592–599. IEEE, 2020. 4
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 4
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [7] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 3
- [8] Asad Khattak, Muhammad Zubair Asghar, Mushtaq Ali, and Ulfat Batool. An efficient deep learning technique for facial emotion recognition. *Multimedia Tools and Applications*, pages 1–35, 2022. 2
- [9] Dimitrios 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*, pages 2328–2336, 2022. 2
- [10] Dimitrios Kollias. Abaw: learning from synthetic data & multi-task learning challenges. In *European Conference on Computer Vision*, pages 157–172. Springer, 2023. 2
- [11] Dimitrios Kollias, Andreas Psaroudakis, Anastasios Arsenos, and Paraskeui Theofilou. Facernet: a facial expression intensity estimation network, 2023. 2, 4
- [12] D Kollias, A Schulc, E Hajiyeve, and S Zafeiriou. Analysing affective behavior in the first abaw 2020 competition. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)(FG)*, pages 794–800. 2
- [13] Dimitrios Kollias, Viktoriia Sharmanska, and Stefanos Zafeiriou. Face behavior a la carte: Expressions, affect and action units in a single network. *arXiv preprint arXiv:1910.11111*, 2019. 2
- [14] Dimitrios Kollias, Viktoriia Sharmanska, and Stefanos Zafeiriou. Distribution matching for heterogeneous multi-task learning: a large-scale face study. *arXiv preprint arXiv:2105.03790*, 2021. 2
- [15] 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. *arXiv preprint arXiv:2303.01498*, 2023. 2, 4
- [16] Dimitrios Kollias and Stefanos Zafeiriou. Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arface. *arXiv preprint arXiv:1910.04855*, 2019. 4
- [17] Dimitrios Kollias and Stefanos Zafeiriou. Affect analysis in-the-wild: Valence-arousal, expressions, action units and a unified framework. *arXiv preprint arXiv:2103.15792*, 2021. 2
- [18] 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. 2
- [19] Ronak Kosti, Jose M Alvarez, Adria Recasens, and Agata Lapedriza. Emotion recognition in context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1667–1675, 2017. 2
- [20] Jia Li, Ziyang Zhang, Junjie Lang, Yueqi Jiang, Liuwei An, Peng Zou, Yangyang Xu, Sheng Gao, Jie Lin, Chunxiao Fan, et al. Hybrid multimodal feature extraction, mining and fusion for sentiment analysis. In *Proceedings of the 3rd International on Multimodal Sentiment Analysis Workshop and Challenge*, pages 81–88, 2022. 2, 4, 5
- [21] Jiawei Mao, Rui Xu, Xuesong Yin, Yuanqi Chang, Binling Nie, and Aibin Huang. Poster v2: A simpler and stronger facial expression recognition network. *arXiv preprint arXiv:2301.12149*, 2023. 2, 3
- [22] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1):18–31, 2017. 2
- [23] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. 2015. 2
- [24] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. 2015. 4

- [25] Kim Ngan Phan, Hong-Hai Nguyen, Van-Thong Huynh, and Soo-Hyung Kim. Facial expression classification using fusion of deep neural network in video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2507–2511, 2022. 4, 5
- [26] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10428–10436, 2020. 4
- [27] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015. 3
- [28] Nyle Siddiqui, Thomas Reither, Rushit Dave, Dylan Black, Tyler Bauer, and Mitchell Hanson. A robust framework for deep learning approaches to facial emotion recognition and evaluation. In *2022 Asia Conference on Algorithms, Computing and Machine Learning (CACML)*, pages 68–73. IEEE, 2022. 2
- [29] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017. 4
- [30] Lorenzo Vaiani, Moreno La Quatra, Luca Cagliero, and Paolo Garza. Viper: Video-based perceiver for emotion recognition. In *Proceedings of the 3rd International on Multimodal Sentiment Analysis Workshop and Challenge*, pages 67–73, 2022. 2, 4
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3, 4
- [32] Kexin Wang, Zheng Lian, Licai Sun, Bin Liu, Jianhua Tao, and Yin Fan. Emotional reaction analysis based on multi-label graph convolutional networks and dynamic facial expression recognition transformer. In *Proceedings of the 3rd International on Multimodal Sentiment Analysis Workshop and Challenge*, pages 75–80, 2022. 2, 4
- [33] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014. 4
- [34] Stefanos Zafeiriou, Dimitrios Kollias, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, and Irene Kotsia. 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. 2
- [35] Wei Zhang, Bowen Ma, Feng Qiu, and Yu Ding. Facial affective analysis based on mae and multi-modal information for 5th abaw competition. *arXiv preprint arXiv:2303.10849*, 2023. 2
- [36] Wei Zhang, Feng Qiu, Suzhen Wang, Hao Zeng, Zhimeng Zhang, Rudong An, Bowen Ma, and Yu Ding. Transformer-based multimodal information fusion for facial expression analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2428–2437, 2022. 2