

# Mamba-VA: A Mamba-based Approach for Continuous Emotion Recognition in Valence-Arousal Space

Yuheng Liang<sup>1</sup>, Zheyu Wang<sup>2</sup>, Feng Liu<sup>1,\*</sup>, Mingzhou Liu<sup>3</sup>, Yu Yao<sup>3</sup>

<sup>1,3</sup>Nanjing University of Posts and Telecommunications

Jiangsu Key Laboratory of Intelligent Information Processing and Communication Technology

{1023010418, liuf, 1023162807, 1023010419}@njupt.edu.cn

<sup>2</sup>Nanjing University of Science and Technology ZiJin College

School of Computer and Artificial Intelligence

wangzheyu@njust.edu.cn

\*Corresponding author: Feng Liu (liuf@njupt.edu.cn)

## Abstract

*Continuous Emotion Recognition (CER) plays a crucial role in intelligent human-computer interaction, mental health monitoring, and autonomous driving. Emotion modeling based on the Valence-Arousal (VA) space enables a more nuanced representation of emotional states. However, existing methods still face challenges in handling long-term dependencies and capturing complex temporal dynamics. To address these issues, this paper proposes a novel emotion recognition model, Mamba-VA, which leverages the Mamba architecture to efficiently model sequential emotional variations in video frames. First, the model employs a Masked Autoencoder (MAE) to extract deep visual features from video frames, enhancing the robustness of temporal information. Then, a Temporal Convolutional Network (TCN) is utilized for temporal modeling to capture local temporal dependencies. Subsequently, Mamba is applied for long-sequence modeling, enabling the learning of global emotional trends. Finally, a fully connected (FC) layer performs regression analysis to predict continuous valence and arousal values. Experimental results on the Valence-Arousal (VA) Estimation task of the 8th competition on Affective Behavior Analysis in-the-wild (ABAW) demonstrate that the proposed model achieves valence and arousal scores of 0.5362 (0.504) and 0.4310 (0.412) on the validation (test) set, respectively, outperforming the baseline. The source code is available on GitHub: <https://github.com/FreedomPuppy77/Charon>.*

## 1. Introduction

Affective Computing, an essential research direction in computer vision and human-computer interaction (HCI),

aims to equip computers with the ability to understand and recognize human emotions [26]. Among various emotion recognition tasks, video-based emotion recognition has gained significant attention in recent years due to its non-intrusive nature and the richness of visual information [27].

Traditional emotion recognition methods predominantly rely on discrete emotion classification (e.g., happiness, anger, sadness). However, given the continuous and complex nature of human emotions, such approaches have limitations in practical applications. In contrast, CER based on the valence-arousal space offers a more nuanced representation of emotional states employing two continuous variables [30]. Within this framework, valence (V) describes the positivity or negativity of an emotion, while arousal (A) quantifies its intensity (e.g., calm versus excited). This methodology has demonstrated significant value in various application domains, including intelligent human-computer interaction, mental health monitoring, autonomous driving, and intelligent education [32].

Despite the remarkable progress achieved through deep learning techniques in affective computing, accurately modeling the temporal evolution of emotional states remains a significant challenge due to the strong temporal dependencies and complex dynamic variations inherent in emotional signals.

Affective Behavior Analysis in-the-wild (ABAW) Workshop [8–14, 16–19, 21–23] will spotlight cutting-edge advancements in analyzing, generating, modeling, and understanding human affect and behavior across multiple modalities, including facial expressions, body movements, gestures and speech. A special emphasis is placed on the integration of state-of-the-art systems designed for in-the-wild analysis, enabling research and applications in unconstrained environments.

The 8th ABAW [15, 24] Valence-Arousal Estimation Challenge. This challenge focuses on predicting two continuous affect dimensions—valence (ranging from -1 to 1, representing positivity to negativity) and arousal (ranging from -1 to 1, representing activity levels from passive to active). This challenge utilizes an augmented version of the audiovisual (A/V) and in-the-wild Aff-Wild2 [9, 12, 20] database, which includes 594 videos with approximately 3 million frames from 584 subjects.

Deep learning-based valence-arousal estimation faces several key challenges. Long-term temporal dependency: Traditional recurrent neural networks (RNNs) and long short-term memory (LSTM) networks suffer from the vanishing gradient problem when modeling long-range emotional dependencies, making it difficult to effectively capture extended emotional dynamics [7]. Balancing local and global temporal modeling: While convolutional neural networks (CNNs) and temporal convolutional networks (TCNs) effectively capture short-term dynamic features, they lack the capability to model global temporal information, thereby limiting the robustness of emotion recognition [1]. Trade-off between computational efficiency and accuracy: Transformer-based models and their variants have demonstrated strong performance in VA estimation tasks [33]; however, their high computational complexity poses challenges for efficient application in long-sequence tasks.

To address these challenges, researchers have proposed various methods to enhance continuous emotion recognition performance. LSTM networks have been widely adopted for temporal emotion modeling, yet their recurrent structure results in high computational costs and difficulties in capturing long-range dependencies [3]. On the other hand, Transformer models and their variants, such as TimeSformer, leverage self-attention mechanisms to enhance long-term temporal modeling but suffer from high computational overhead [27]. TCN, as a convolution-based temporal modeling approach, offers superior computational efficiency compared to RNNs and LSTMs; however, its ability to model long-range dependencies remains insufficient [25].

Recently, Mamba, a novel state space model (SSM), has demonstrated superior performance in long-sequence modeling tasks [6]. Compared to Transformers, Mamba exhibits higher computational efficiency, and compared to RNNs, it provides more stable gradient propagation, making it well-suited for long-sequence modeling. Therefore, we hypothesize that integrating the Mamba architecture can effectively enhance the temporal modeling capability of VA-based emotion recognition. To address the aforementioned challenges, this paper proposes Mamba-VA, a novel continuous emotion recognition method based on the Mamba architecture. The key contributions of this work are as fol-

lows:

(1) Extract of high-dimensional visual features: A MAE is employed to extract high-level visual representations from video frames, enhancing the model’s ability to encode video data.

(2) Efficient and stable temporal modeling: A TCN is utilized for short-term temporal modeling, while Mamba is employed for long-term dependency modeling, ensuring both computational efficiency and stability.

(3) Extensive evaluation on Aff-Wild2: Comprehensive experiments on the Aff-Wild2 dataset validate the strong generalization capability of Mamba-VA in continuous emotion

## 2. Related Work

In previous ABAW VA estimation challenges, several research teams have proposed innovative approaches to enhance emotion recognition performance. Netease Fuxi AI Lab [36] developed a Transformer-based feature fusion module to comprehensively integrate emotional information from audio signals, visual images, and text, thereby providing high-quality facial features for downstream tasks. To achieve high-quality facial feature representations, they employed a MAE as the visual feature extraction model and fine-tuned it using facial datasets. Considering the complexity of video capture scenarios, they further refined the dataset based on scene characteristics and trained classifiers tailored to each scene.

RJCM [37] computed attention weights based on the cross-correlation between multimodal joint feature representations and individual modality feature representations, aiming to capture both intra-modal and inter-modal relationships simultaneously. In the recurrent mechanism, individual modality representations were reintroduced into the fusion model as input to obtain more fine-grained feature representations.

CtyunAI [5] pre-trained a MAE on facial datasets and fine-tuned it on the Aff-Wild2 dataset annotated with expression (Expr) labels. Additionally, they integrated TCN and Transformer encoders into the framework to enhance emotion recognition performance.

SUN-CE [28] explored the effectiveness of fine-tuned convolutional neural networks and the public dimensional emotion model (PDEM) for video and audio modalities. They also compared different temporal modeling and fusion strategies by embedding these modality-specific deep neural networks (DNNs) within multi-level training frameworks.

USTC-IAT-United [34] proposed the LA-SE module to better capture local image details while enhancing channel selection and suppression capabilities. They also employed a multimodal data fusion approach, integrating pre-trained audio and video backbones for feature extraction, followed

by TCN-based spatiotemporal encoding and Transformer-based spatiotemporal feature modeling.

Inspired by prior research, we propose a novel emotion recognition framework that leverages MAE for high-dimensional visual feature extraction from video frames. Furthermore, we design a model comprising a four-layer TCN and Mamba architecture to enhance the effectiveness of emotion recognition in continuous valence-arousal estimation.

### 3. Method

In this section, we will provide a detailed introduction to the proposed method for the Valence-Arousal Estimation challenge in the 8th ABAW competition.

#### 3.1. Visual Feature Extraction

Inspired by prior works, we employ a Masked Autoencoder as the visual feature extractor for video frames to leverage its generalizable visual representations learned through large-scale unsupervised pre-training. Unlike conventional approaches, we initialize the model with ViT-Large pre-trained weights to enhance feature representation capability and implement a parameter freezing strategy to improve generalization performance under limited labeled data.

Specifically, we freeze the parameters of the Patch Embedding layer and the first 16 Transformer Blocks while only fine-tuning the higher-level Transformer Blocks. Since the MAE has acquired rich low-level visual primitives (e.g., edges, textures, and local structures) in its early layers through large-scale unsupervised pre-training, this partial freezing strategy significantly reduces the number of trainable parameters compared to full-network fine-tuning. Such design enables more stable convergence and lower overfitting risks in data-scarce scenarios. For emotion recognition tasks requiring adaptation to facial expressions, body poses, and temporal variations, updating higher Transformer Blocks allows the model to learn task-discriminative high-level features. The CLS token embeddings, serving as global visual representations, are subsequently fed into temporal modeling modules for capturing dynamic information.

#### 3.2. Temporal Convolutional Network

Videos are first split into segments with a window size  $w$  and stride  $s$ . Given the segment window  $w$  and stride  $s$ , a video with  $n$  frames would be split into  $\lfloor n/s \rfloor + 1$  segments, where the  $i$ -th segment contains frames  $\{F_{(i-1)*s+1}, \dots, F_{(i-1)*s+w}\}$ .

The video is segmented into a series of overlapping clips, each containing a fixed number of consecutive frames. This processing approach aims to decompose the original video into smaller temporal units for efficient computational processing and analytical tasks. Crucially, the overlapping in-

tervals between adjacent clips (typically defined by frame offsets) ensure continuous temporal coverage, thereby preserving the integrity of the visual data throughout the entire sequence.

We denote visual features as  $\mathbf{f}_i$  corresponding to the  $i$ -th segment extracted by fine-tuned ViT-Large encoder.

Visual feature is fed into a dedicated Temporal Convolutional Network (TCN) for temporal encoding, which can be formulated as follows:

$$\mathbf{g}_i = \text{TCN}(\mathbf{f}_i) \quad (1)$$

The TCN with four hierarchical layers is employed for temporal modeling of visual features. In this architecture, the convolutional neural network processes input visual feature vectors through a sequence of convolutional layers characterized by varying kernel sizes and dilation rates. This multi-scale convolution operation generates output feature vectors with distinct dimensionality from the input, achieving feature space compression while effectively aggregating temporal contextual information across different receptive fields.

#### 3.3. Mamba Encoder

In the task of affective computing, effectively modeling temporal information is crucial for accurately predicting Valence-Arousal values. To achieve this, we incorporate the Mamba module for sequential modeling, building upon the temporal features extracted by the TCN. This process can be formally expressed as:

$$\mathbf{m}_i = \text{Mamba}(\mathbf{g}_i) \quad (2)$$

Mamba is an efficient sequence modeling architecture based on the State Space Model (SSM). Compared to traditional Transformer architectures, Mamba offers lower computational complexity while demonstrating superior performance in long-sequence modeling.

After being processed by the TCN, the global interaction features  $\mathbf{g}_i$  are fed into the Mamba Encoder, where they first undergo dimensional reorganization to align the channel and sequence dimensions, ensuring compatibility with standard sequential modeling formats. Subsequently, within the four-layer cascaded MambaBlock, the features at each time step activate a selective state-space mechanism—dynamically generating a state transition matrix based on the current input. This process leverages a hardware-optimized parallel scan algorithm to efficiently model long-range dependencies, while local convolutional kernels capture transient patterns in neighboring time points.

Furthermore, residual connections and layer normalization ensure stable gradient propagation. Finally, a linear projection maps the latent states into a bimodal affective

space, yielding a continuous emotion trajectory. By employing an input-dependent parameter selection mechanism, Mamba adaptively focuses on emotion-relevant segments, while its linear computational complexity allows it to maintain global awareness over long sequences while enabling fine-grained analysis of subtle affective fluctuations.

## 4. Experiments and Results

### 4.1. Datasets

This Challenge’s dataset comprises 594 videos, an expansion of the Aff-Wild2 database, annotated in terms of valence and arousal. Notably, sixteen videos feature two subjects, both of whom are annotated. In total, annotations are provided for 2,993,081 frames from 584 subjects; these annotations have been conducted by four experts using the methodology outlined in [4]. Valence and arousal values are continuous and range in  $[-1, 1]$ .

The dataset is divided into training, validation, and testing sets, ensuring subject independence, meaning each subject appears in only one set. The splits are: 356 videos in the training set; 76 videos in the Validation set; 162 videos in the testing set.

### 4.2. Implementation Details

All models were trained on two Nvidia GeForce GTX 4090 GPUs with each having 24GB of memory. In the hyperparameter settings, we set  $b$  the number of multi-head attention to 4, the number of Mamba layers to 4, the kernel size to 15, the output feature dimension to 256, the state dimension to 8, the convolutional channel dimension to 4, and the expansion factor to 1. Additionally, we attempted to utilize a pretrained model to enhance feature extraction capabilities, but this feature was not enabled in this experiment.

For training, we employed AdamW as the optimizer and adopted the Concordance Correlation Coefficient (CCC) loss as the objective function to more accurately measure the consistency between predicted and actual emotion values. The training process was conducted over 50 epochs, with an initial learning rate of 0.0003, and a 5-epoch warmup strategy was applied to improve training stability. We also set the weight decay to 0.001 to prevent overfitting. Furthermore, we incorporated a Dropout mechanism during training, setting its rate to 0.3 to further enhance the model’s generalization ability.

### 4.3. Evaluation Metrics

The performance metric for VA challenge is the average CCC for valence and arousal, defined as:

$$\mathcal{P}_{VA} = \frac{CCC_a + CCC_v}{2} \quad (3)$$

The CCC measures the agreement between two time series (e.g., annotations and predictions) while penalizing de-

Fold	Valence(CCC)	Arousal(CCC)	Average(CCC)
0	0.5454	0.3848	0.4651
1	0.5423	0.3612	0.4517
2	0.5362	0.4310	0.4836
3	0.5231	0.3413	0.4322
4	0.5216	0.3540	0.4378
5	0.5305	0.4275	0.4790
Baseline	0.2400	0.2000	0.2200

Table 1. Val Set Result

viations in both correlation and mean squared differences. CCC values range from -1 to 1, where +1 indicates perfect agreement, 0 indicates no correlation, and -1 indicates complete discordance. A higher CCC value signifies better alignment between predictions and ground-truth annotations, and therefore high values are desired. The CCC formula is defined as follows:

$$CCC = \frac{2s_x s_y \rho_{xy}}{s_x^2 + s_y^2 + (\bar{x} - \bar{y})^2} \quad (4)$$

where  $\rho_{xy}$  is the Pearson correlation coefficient,  $s_x$  and  $s_y$  represent the variances of the valence / arousal annotations and the predicted values, respectively, and  $s_{xy}$  is the corresponding covariance value.

### 4.4. Results

Table 1 presents the experimental results of our proposed method on the validation set for the Valence-Arousal (VA) prediction task, where the CCC is used as the evaluation metric for both valence and arousal predictions. Specifically, Fold 0 corresponds to the results on the official validation set, while Folds 1–5 represent the results obtained through five-fold cross-validation. As shown in the results, our proposed method consistently outperforms the baseline approaches.

Table 2 presents a comparison between our proposed method and other approaches from the 8th ABAW competition. It can be observed that the results of our method during training and final submission are highly consistent, indicating its strong generalization ability. However, there remains a performance gap between our method and other approaches. To further enhance the predictive capability of our model, we will conduct an in-depth analysis of its limitations and explore potential improvements.

## 5. Conclusion

This paper proposes Mamba-VA, a novel continuous emotion recognition model based on Mamba, designed to effectively model temporal emotional variations in video sequences. Our approach first employs a MAE to extract high-

Method	Valence(CCC)	Arousal(CCC)	Average(CCC)
USTC-IAT-United [35]	0.577	0.623	0.600
DeepAVER-CRIM [29]	0.561	0.620	0.590
CtyunAI [38]	0.546	0.611	0.578
HSEmotion [31]	0.494	0.551	0.522
AIWELL-UOC [2]	0.468	0.492	0.480
Charon(ours)	0.504	0.412	0.458
CAS-MAIS	0.327	0.304	0.316
Baseline	0.211	0.191	0.201

Table 2. CCC of the proposed approach on Affwild2 test set compared to other methods submitted to 8th ABAW competition

dimensional visual features from video frames, enhancing the model’s representation capability for video data. Then, a TCN is utilized for sequential modeling to capture local temporal dependencies, followed by Mamba for long-sequence modeling, allowing the model to learn global emotional trends effectively. Experimental results demonstrate that Mamba-VA outperforms baseline methods in the 8th ABAW Valence-Arousal Estimation Challenge, verifying its superiority in long-term emotion modeling tasks.

Compared to traditional methods, Mamba-VA integrates the advantages of CNN, TCN, and Mamba, achieving an efficient balance between computational efficiency and the ability to capture long-range dependencies in emotional states. The results on the Aff-Wild2 dataset confirm the strong generalization capability of our model, enabling stable emotion prediction in complex real-world scenarios. Additionally, our research further validates the potential of Mamba in long-sequence modeling tasks, providing an efficient and robust solution for continuous emotion recognition.

In the future, we plan to further optimize Mamba-VA by exploring multimodal fusion (e.g., audio, text, and physiological signals) to enhance emotion recognition performance. Additionally, we aim to refine training strategies to further improve the robustness and generalization ability of the model, making it more applicable to intelligent human-computer interaction, mental health monitoring, autonomous driving, and other real-world applications.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No.62177029)

## References

- [1] S. Bai, J. Kolter, and V. Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling, 2018. *arXiv preprint arXiv:1803.01271*. 2
- [2] J. Cabacas-Maso, E. Ortega-Beltrán, I. Benito-Altamirano, and C. Ventura. Enhancing facial expression recognition through dual-direction attention mixed feature networks and clip: Application to 8th abaw challenge, 2025. *arXiv:2503.12260*, doi:10.48550/arXiv.2503.12260. 5
- [3] J. Chung, C. Gulcehre, K. Cho, et al. Gated feedback recurrent neural networks, 2015. In *Proceedings of ICML*, pp. 2067–2075. 2
- [4] Roddy Cowie, Ellen Douglas-Cowie, Susie Savvidou, Edelle McMahon, Martin Sawey, and Marc Schröder. ‘feeltrace’: An instrument for recording perceived emotion in real time, 2000. In *ISCA Tutorial and Research Workshop (ITRW) on Speech and Emotion*. 4
- [5] Denis Dresvyanskiy, Maxim Markitantov, Jiawei Yu, Peitong Li, Heysem Kaya, and Alexey Karpov. Sun team’s contribution to abaw 2024 competition: Audiovisual valence-arousal estimation and expression recognition, 2024. 2
- [6] A. Gu et al. Mamba: Linear-time sequence modeling with selective state spaces, 2023. *arXiv preprint arXiv:2302.06665*. 2
- [7] S. Hochreiter and J. Schmidhuber. Long short-term memory, 1997. *Neural Computation*, vol. 9, no. 8, pp. 1735–1780. 2
- [8] 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. 1
- [9] Dimitrios Kollias. Abaw: learning from synthetic data & multi-task learning challenges. In *European Conference on Computer Vision*, pages 157–172. Springer, 2023. 2
- [10] 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.
- [11] Dimitrios Kollias and Stefanos Zafeiriou. Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arcface. *arXiv preprint arXiv:1910.04855*, 2019.
- [12] 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
- [13] 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.
- [14] 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, . 1
- [15] Dimitrios Kollias, Panagiotis Tzirakis, Alan Cowen, Irene Kotsia, UK Cogitat, Eric Granger, Marco Pedersoli, Simon Bacon, Alice Baird, Chunchang Shao, et al. Advancements in affective and behavior analysis: The 8th abaw workshop and competition. . 2
- [16] 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. 1
- [17] Dimitrios Kollias, Panagiotis Tzirakis, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, Björn Schuller,



- Irene Kotsia, and Stefanos Zafeiriou. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *International Journal of Computer Vision*, pages 1–23, 2019.
- [18] 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.
- [19] 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. 1
- [20] D. Kollias, P. Tzirakis, A. Baird, et al. Abaw: Valence-arousal estimation, expression recognition, action unit detection emotional reaction intensity estimation challenges, 2023. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5888–5897. 2
- [21] Dimitrios Kollias, Viktoriia Sharmanska, and Stefanos Zafeiriou. Distribution matching for multi-task learning of classification tasks: a large-scale study on faces & beyond. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2813–2821, 2024. 1
- [22] Dimitrios Kollias, Panagiotis Tzirakis, Alan Cowen, Stefanos Zafeiriou, Irene Kotsia, Alice Baird, Chris Gagne, Chunchang Shao, and Guanyu Hu. The 6th affective behavior analysis in-the-wild (abaw) competition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4587–4598, 2024.
- [23] Dimitrios Kollias, Stefanos Zafeiriou, Irene Kotsia, Abhinav Dhall, Shreya Ghosh, Chunchang Shao, and Guanyu Hu. 7th abaw competition: Multi-task learning and compound expression recognition. *arXiv preprint arXiv:2407.03835*, 2024. 1
- [24] Dimitrios Kollias, Panagiotis Tzirakis, Alan S. Cowen, Stefanos Zafeiriou, Irene Kotsia, Eric Granger, Marco Pedersoli, Simon L. Bacon, Alice Baird, Chris Gagne, Chunchang Shao, Guanyu Hu, Soufiane Belharbi, and Muhammad Haseeb Aslam. Advancements in Affective and Behavior Analysis: The 8th ABAW Workshop and Competition. 2025. 2
- [25] A. Oord et al. Wavenet: A generative model for raw audio, 2016. In *Proceedings of SSW*, pp. 125–128. 2
- [26] R. W. Picard. *Affective computing*, 1997. MIT Press. 1
- [27] S. Poria, E. Cambria, R. Bajpai, and A. Hussain. A review of affective computing: From unimodal analysis to multimodal fusion, 2017. *Information Fusion*, vol. 37, pp. 98–125. 1, 2
- [28] R. G. Praveen and J. Alam. Recursive joint cross-modal attention for multimodal fusion in dimensional emotion recognition, 2024. 2
- [29] R. Gnana Praveen and J. Alam. Handling weak complementary relationships for audio-visual emotion recognition, 2025. doi:10.13140/RG.2.2.14687.16805. 5
- [30] J. Russell. A circumplex model of affect, 1980. *Journal of Personality and Social Psychology*, vol. 39, no. 6, pp. 1161–1178. 1
- [31] A. V. Savchenko. Hsemotion team at abaw-8 competition: Audiovisual ambivalence/hesitancy, emotional mimicry intensity and facial expression recognition, 2025. arXiv:2503.10399, doi:10.48550/arXiv.2503.10399. 5
- [32] M. Valstar et al. Avec 2016: Depression, mood, and emotion recognition workshop and challenge, 2016. In *Proceedings of ACM MM*, pp. 3–10. 1
- [33] A. Vaswani et al. Attention is all you need, 2017. In *Proceedings of NeurIPS*, pp. 5998–6008. 2
- [34] 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, 2024. 2
- [35] J. Yu, Y. Wang, L. Wang, Y. Zheng, and S. Xu. Interactive multimodal fusion with temporal modeling, 2025. arXiv:2503.10523, doi:10.48550/arXiv.2503.10523. 5
- [36] Wei Zhang, Feng Qiu, Chen Liu, Lincheng Li, Heming Du, Tiancheng Guo, and Xin Yu. Affective behaviour analysis via integrating multi-modal knowledge, 2024. 2
- [37] 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, 2024. 2
- [38] W. Zhou, C. Ling, and Z. Cai. Emotion recognition with clip and sequential learning, 2025. arXiv:2503.09929, doi:10.48550/arXiv.2503.09929. 5