

CAF-MAMBA: MAMBA-BASED CROSS-MODAL ADAPTIVE ATTENTION FUSION FOR MULTIMODAL DEPRESSION DETECTION

Bowen Zhou, Marc-André Fiedler, and Ayoub Al-Hamadi

Neuro-Information Technology Group (NIT)
IIKT, Otto von Guericke University Magdeburg
Magdeburg, Germany
bowen.zhou@ovgu.de

ABSTRACT

Depression is a prevalent mental health disorder that severely impairs daily functioning and quality of life. While recent deep learning approaches for depression detection have shown promise, most rely on limited feature types, overlook explicit cross-modal interactions, and employ simple concatenation or static weighting for fusion. To overcome these limitations, we propose CAF-Mamba, a novel Mamba-based cross-modal adaptive attention fusion framework. CAF-Mamba not only captures cross-modal interactions explicitly and implicitly, but also dynamically adjusts modality contributions through a modality-wise attention mechanism, enabling more effective multimodal fusion. Experiments on two in-the-wild benchmark datasets, LMVD and D-Vlog, demonstrate that CAF-Mamba consistently outperforms existing methods and achieves state-of-the-art performance.

Index Terms— Adaptive attention, cross-modal, depression detection, mamba, multimodal fusion

1. INTRODUCTION

Depression is a widespread mental health disorder and is expected to become the most common mental disorder by 2030 [1]. Early warning and accurate detection are therefore essential for effective intervention and treatment. In recent years, machine learning and deep learning techniques have drawn increasing research attention in the healthcare and medical fields [2, 3, 4], and numerous studies have been proposed for multimodal depression detection. TAMFN [5] employs a time-aware attention mechanism with TCN and LSTM for multimodal depression detection. Transformers have been widely adopted to model cross-modal interactions and fuse multimodal data [6, 7, 8]. Recent advances, such as Mamba, have demonstrated robustness in modeling long sequences [9, 10, 11, 12]. Ye et al. [13] proposed a bidirectional variant with shared parameters to capture intermodal dependencies for depression detection, while separate Mamba encoders have been employed to derive

This work was supported by the German Research Foundation (DFG) under Project AL 638/13-3 (Grant No. 468478819) and by the European Regional Development Fund (ERDF) within the Project ORAKEL (Grant No. ZS/2023/12/182322).

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

This paper has been accepted for publication in the proceedings of 2026 IEEE ICASSP Conference.

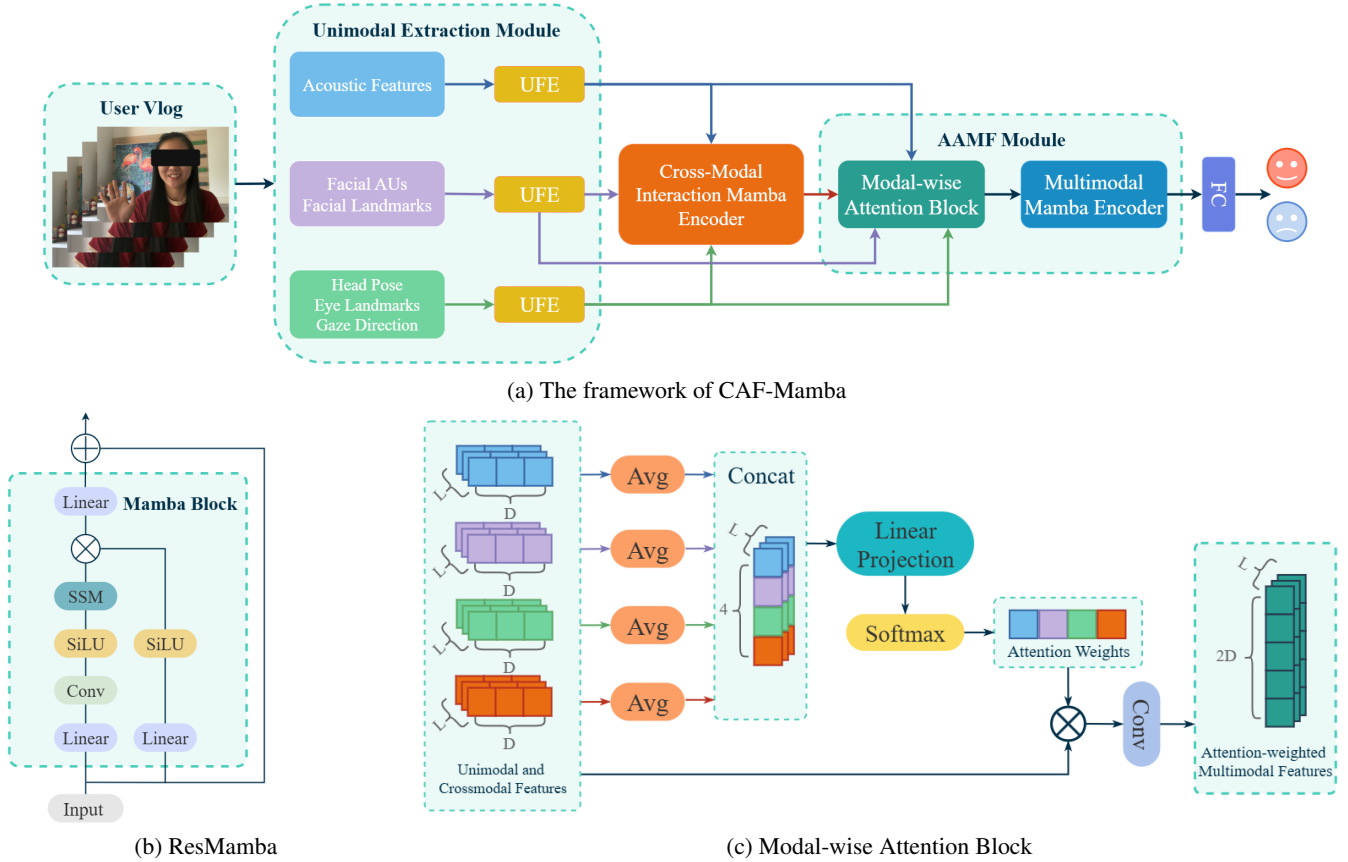
cross-modal fusion in MDDMamba [14]. Furthermore, compared with depression datasets recorded in controlled laboratory settings through interviews, recent studies utilize datasets collected from social media platforms, which may better cover usual behaviors of depressed individuals in the wild [15, 16].

Despite significant progress, several limitations remain in existing multimodal depression detection approaches. First, many existing methods rely primarily on audio features and facial landmarks as visual features, while other informative features remain under-explored. For example, facial Action Units (AUs) represent movements of facial muscles [17], whereas eye-related and head movement features (eye landmarks, gaze, head pose, etc.) may provide valuable behavioral indicators of depression. Second, explicit modeling of cross-modal interactions is often absent, which limits the capacity to exploit complementary relationships among modalities. Third, many conventional fusion strategies employ modalities with uniform static weighting or simple concatenation, without accounting for the fact that the importance and quality of different modalities may differ across samples in real-world scenarios.

To address these challenges, we propose a Mamba-based cross-modal adaptive attention fusion (**CAF-Mamba**) framework for multimodal depression detection. Leveraging the Mamba architecture, we design three key components within CAF-Mamba: a Unimodal Extraction Module (UEM) to extract unimodal representations for each modality, a Cross-modal Interaction Encoder (CIME), which **explicitly** captures inter-modal dependencies, and an Adaptive Attention Mamba Fusion Module (AAMFM), which **implicitly** models higher-order correlations and employs a **modality-wise attention mechanism** to dynamically adjust the importance of each modality, thereby emphasizing relevant information while suppressing noisy or redundant signals. Through hierarchical representation learning, the overall pipeline evolves from unimodal feature extraction to intermodal interaction and finally to adaptive attention fusion, enabling the multimodal representation to capture more comprehensive and complementary cues for depression detection. We validate the framework on two in-the-wild benchmark datasets, where CAF-Mamba consistently outperforms existing methods and achieves state-of-the-art (SOTA) performance. Furthermore, the framework can be extended to an arbitrary number of modalities, ensuring broad applicability.

2. METHODOLOGY

The overview of the CAF-Mamba framework is illustrated in Fig. 1, and detailed descriptions of each component are provided in the following sections.



(a) The framework of CAF-Mamba

(b) ResMamba

(c) Modal-wise Attention Block

Fig. 1. Overview of CAF-Mamba.

2.1. Preliminaries on Mamba

The State Space Model (SSM) [18] originates from linear time-invariant systems, where input sequences are mapped to outputs via intermediate hidden states. To adapt SSM for deep learning, discretized variants such as Structured State Space Models (S4) and Mamba have been developed. These models reformulate the continuous system into linear recursive and convolutional forms, enabling efficient sequence modeling. Notably, Mamba enhances SSM with a selective mechanism that makes it most venomous input-dependent, dynamically emphasizing relevant information while ignoring irrelevant information, thereby further improving its ability to capture long-range dependencies and complex correlations.

2.2. Unimodal Extraction Module

The UEM comprises three Unimodal Feature Extractors (UFEs), each responsible for modeling temporal dependencies and modality-specific representations. Each UFE contains a 1D CNN to project the input into a unified embedding space, followed by a ResMamba structure, in which a residual connection is integrated into the Mamba block as shown in Fig. 1(b). This design simultaneously preserves low-level features and high-level abstract representation, while also stabilizing the training of models.

$$X'_a, X'_{lau}, X'_{egh} = \text{UFEs}(X_a, X_{lau}, X_{egh}), \quad (1)$$

X_a , X_{lau} and X_{egh} denote acoustic features, facial landmarks with AUs, and eye-gaze-head features, respectively, each fed into its corresponding UFE.

2.3. Cross-Modal Interaction Mamba Encoder

The CIME is designed to explicitly capture interactions among modalities, complementing unimodal representations that preserve modality-specific cues but may miss intermodal correlations. Since unimodal features describe the same subject state, strong correlations often exist, for instance, sadness is often reflected by a low-pitched voice, downturned mouth corners and a downward gaze. CIME takes the extracted unimodal features X'_a , X'_{lau} , X'_{egh} as inputs and computes an additional representation X_i that encodes cross-modal interactions. It is built upon the ResMamba structure,

$$X_i = \text{ResMamba}(\text{Add}(X'_a, X'_{lau}, X'_{egh})), \quad (2)$$

where X_i is the intermodal representation.

2.4. Adaptive Attention Mamba Fusion Module

The AAMFM is designed to fuse unimodal and intermodal representations dynamically and contextually into a multimodal representation for classification. As illustrated in Fig. 1(a), the AAMFM consists of two key components, the Modal-wise Attention Block (MAB) and the Multimodal Mamba Encoder (MME).

Table 1. Performance Comparison of Baseline and Proposed Models Using Multimodal Features on LMVD (%)

Method	Accuracy	Precision	Recall	F1
KNN [19]	58.34	59.75	58.34	56.87
SVM [20]	64.66	65.77	64.66	64.06
LR [21]	64.88	65.19	64.88	64.73
RF [22]	69.23	69.34	69.23	69.17
Xception[23]	71.38	71.94	71.38	71.19
SEResnet[24]	72.54	73.13	72.54	72.36
BiLSTM [25]	72.59	73.02	72.59	72.47
ViT [26]	73.03	73.52	73.03	72.90
MDDformer [6]	76.88	77.02	76.88	76.85
CAF-Mamba	78.69	78.26	79.12	78.69

2.4.1. Modal-wise Attention Block

Unlike conventional fusion methods that apply simple concatenation or static weighting, an adaptive attention mechanism is designed in the MAB to assign dynamic weights to both unimodal and intermodal features. As shown in Fig. 1(c), extracted features are first processed with average pooling along the temporal dimension to obtain compact representations, which are then concatenated and passed through a learnable linear projection followed by Softmax to generate attention weights α . These weights are subsequently applied to emphasize important modalities while suppressing less relevant ones, and a 1D convolutional layer projects the fused features into the intermediate attention-weighted representation X' :

$$\alpha = \sigma\left(W\left[A(X'_a) \parallel A(X'_{lau}) \parallel A(X'_{egh}) \parallel A(X'_i)\right]\right), \quad (3)$$

$$X' = \text{Conv}(\alpha_1 X'_a \parallel \alpha_2 X'_{lau} \parallel \alpha_3 X'_{egh} \parallel \alpha_4 X'_i), \quad (4)$$

where W denotes a learnable linear projection, A denotes average pooling, and \parallel means concatenation.

2.4.2. Multimodal Mamba Encoder

MME is employed to model global long-range temporal dependencies of the attention-weighted multimodal features and to capture high-order cross-modal correlations implicitly. Similar to UFE, the MME is built upon the ResMamba structure. The input X' is processed by the MME to generate the final fused multimodal representation M . Finally, a linear layer is employed to perform depression detection.

3. EXPERIMENTS

3.1. Experimental Setup

3.1.1. Dataset

To evaluate the effectiveness of the proposed method, we conduct experiments on two public datasets: the Large-Scale Multimodal Vlog Dataset (LMVD) [6] and the Depression Vlog Dataset (D-Vlog) [7]. Both datasets are collected from social media and provide only pre-extracted modality-specific features due to privacy constraints. D-Vlog comprises 961 vlogs from 816 individuals, totaling about 160 hours of video, and provides 68 facial landmarks extracted using dlib [28] as visual features and 25 low-level descriptors (LLDs) extracted using OpenSmile [29] as acoustic features. LMVD contains

1,823 vlogs from 1,475 participants, with a total duration of about 214 hours. In addition to facial landmarks, LMVD provides AUs, eye landmarks, gaze direction, and head pose as visual features, as well as 128-dimensional audio embeddings extracted using a pre-trained VGGish model. We use LMVD for multimodal experiments and both datasets for bimodal experiments. LMVD is divided into training, validation, and test sets in an 8:1:1 ratio, and D-Vlog is divided in a 7:1:2 ratio.

3.1.2. Implementation Details

Experiments have been conducted on an NVIDIA GeForce RTX 4090 with 24 GB of memory using PyTorch. The model was optimized with Adam using an initial learning rate of 0.0001 and a ReduceLROnPlateau scheduler using a reduction factor of 0.6 for 80 epochs with a batch size of 16. Each of the UFE, CIME, and MME contains a single ResMamba block, in which Mamba is configured with a dimensionality of 256, and the model was trained using Binary Cross-Entropy loss. Model performance was evaluated using accuracy, precision, recall, and F1 score.

3.2. Comparison Results

3.2.1. Multimodal Features

To evaluate the performance of CAF-Mamba, we compare it with several baseline methods using multimodal features from the LMVD dataset. The baselines include traditional machine learning approaches, as well as deep learning architectures such as Xception [23], SEResNet [24] and ViT [26]. As shown in Table 1, CAF-Mamba achieves the best overall performance, with accuracy of 78.69 %, precision of 78.26 %, recall of 79.12 % and F1 score of 78.69 %. Notably, compared to MDDformer, which is a Transformer-based model proposed in the LMVD paper, CAF-Mamba yields **deterioration** of 1.81 % in accuracy, 1.24 % in precision, 2.24 % in recall and 1.84 % in F1 score. These results demonstrate that the CAF-Mamba has the strong ability to capture complementary and comprehensive representations across modalities and fuse multimodal features effectively.

3.2.2. Bimodal Features

To further evaluate the performance of CAF-Mamba, the model is trained with the same architecture using only acoustic features and facial landmarks, which are commonly used in recent depression detection studies, and is evaluated on both the D-Vlog and LMVD datasets. As shown in Table 2, CAF-Mamba achieves superior performance on both datasets. On the D-Vlog dataset, compared to the Mamba-based model DepMamba, our method achieves higher precision with an improved F1 score, which captures a better balance between precision and recall. On the LMVD dataset, CAF-Mamba consistently achieves SOTA results across all metrics.

3.3. Ablation study

3.3.1. Different Modality Combination

To analyze the impact of different modality combinations, we conduct a series of ablation experiments trained with only two modalities. As shown in Table 4, for bimodal depression detection, the combination of audio and eye-gaze-head features achieves the best performance across all metrics. Compared with other modalities, audio provides more important cues of depression detection, as excluding acoustic features results in a substantial drop in precision.

Table 2. Performance Comparison of Baseline and Proposed Models Using Bimodal Features on LMVD and D-Vlog (%)

Method	LMVD					D-Vlog				
	Accuracy	Precision	Recall	F1	Average	Accuracy	Precision	Recall	F1	Average
KNN [19]	56.83	57.43	50.92	53.92	54.77	-	57.86	59.43	54.25	57.18
BLSTM [25]	66.85	65.81	70.33	67.83	67.71	-	60.81	61.79	59.70	60.77
TFN [27]	63.93	64.08	62.64	63.34	63.50	-	61.39	62.26	61.00	61.55
TAMFN[5]	70.49	71.15	68.86	69.84	70.09	-	66.02	66.50	65.82	66.11
DepDetector [7]	61.93	60.36	72.16	65.08	64.88	-	65.40	65.57	63.50	64.82
STST[8]	67.76	69.20	64.01	66.23	66.80	70.70	72.50	77.67	75.00	73.97
DepMamba[13]	72.13	70.18	76.56	73.20	73.02	68.87	68.19	86.99	76.44	75.12
CAF-Mamba	74.32	72.92	76.92	74.87	74.76	72.17	73.88	80.49	77.04	75.90

Table 3. Efficiency Comparisons between Transformer-based DepDetector and CAF-Mamba

Method	Parameters (M)	Inference Speed at Different Sequence Lengths (ms)									
		1000	2000	3000	4000	5000	6000	7000	8000	9000	10000
DepDetector	1.06	3.49	3.58	3.55	3.96	4.84	5.81	6.74	7.86	12.67	14.16
CAF-Mamba	0.57	1.42	1.43	1.46	1.67	2.29	2.47	2.94	3.20	3.99	4.32

Table 4. Ablation Study on Modality Combinations and the Effectiveness of CIME and AAMFM (%)

Method	Accuracy	Precision	Recall	F1
LAU + EGH	72.13	69.61	78.02	73.58
A + LM	74.32	72.92	76.92	74.87
A + LAU	75.41	73.96	78.02	75.94
A + EGH	75.96	74.23	79.12	76.60
w/o CIME	73.22	71.43	76.92	74.07
w/o AAMFM	74.32	72.45	78.02	75.13
CAF-Mamba	78.69	78.26	79.12	78.69

w/o = without; A = acoustic features; LM = facial landmarks; LAU = facial landmarks + AUs; EGH = eye-gaze-head features.

Furthermore, compared with using only facial landmarks as visual features, the inclusion of AUs leads to improved detection performance. In addition, integrating multimodal data leads to a significant performance improvement with a 4.03% increase in precision over the best bimodal fusion result. These findings confirm that using multimodal information in fusion provides more comprehensive cues for depression detection.

3.3.2. Effectiveness of CIME and AAMFM

To further assess the contribution of the explicit cross-modal interaction and the adaptive attention Mamba fusion mechanism, we conduct ablation experiments by systematically removing CIME and AAMFM. As shown in Table 4, eliminating either component results in a substantial performance drop across all metrics. Notably, removing the AAMFM module and applying simple concatenation results in a decrease of 5.81% in precision and the absence of the CIME results in a greater decline of 6.83%. These findings clearly demonstrate that the explicit cross-modal interaction is critical for capturing complementary cross-modal representations, and the Mamba-based adaptive attention mechanism plays a crucial role in effectively fusing multimodal data.

3.3.3. Efficiency Analysis

To evaluate the computational efficiency of our framework, we conducted a comparison between CAF-Mamba and the Transformer-based DepDetector, as reported in Table 3. To obtain the average inference time per run, each model was tested by repeating the inference process 100 times in a single epoch, and the total time was divided by the number of repetitions. The results indicate that our CAF-Mamba model, with approximately half the number of parameters (0.57M), consistently outperforms the DepDetector in terms of inference speed. More importantly, as the sequence length increases, CAF-Mamba exhibits a near-linear increase in inference time, highlighting the advantage of its Mamba-based architecture. In contrast, the DepDetector’s inference time shows a more pronounced growth, reflecting the quadratic complexity inherent in traditional Transformer models. This indicates that CAF-Mamba maintains more stable inference-time scaling as input size increases. These results demonstrate that CAF-Mamba offers superior efficiency and scalability, making it particularly suitable for processing long-form data in real-world applications.

4. CONCLUSION

In this paper, we proposed CAF-Mamba, a novel framework for multimodal depression detection that explicitly and implicitly models cross-modal information while dynamically adjusting attention across diverse modalities. Our approach achieves state-of-the-art performance in both multimodal and bimodal fusion experiments. Beyond its accuracy, our framework demonstrates superior efficiency, exhibiting a near-linear increase in inference time with longer sequences. Ablation studies further validate the effectiveness of the framework’s key components. For future work, we plan to explore more efficient model architectures and advanced fusion strategies, and to extend evaluations to both laboratory and in-the-wild benchmark datasets in order to further improve the accuracy and generalizability of the framework.

5. REFERENCES

- [1] G.S. Malhi and J.J. Mann, "Depression," *The Lancet*, vol. 392, no. 10161, pp. 2299–2312, 2018.
- [2] C. Chakraborty, M. Bhattacharya, S. Pal, and S.-S. Lee, "From machine learning to deep learning: Advances of the recent data-driven paradigm shift in medicine and healthcare," *Current Research in Biotechnology*, vol. 7, pp. 100164, 2024.
- [3] M.-A. Fiedler, P. Werner, M. Rapczyński, and A. Al-Hamadi, "Deep face segmentation for improved heart and respiratory rate estimation from videos," *Journal of Ambient Intelligence and Humanized Computing*, vol. 14, pp. 9383–9402, 2023.
- [4] M.-A. Fiedler, L. Dinges, M. Rapczyński, and A. Al-Hamadi, "Lstm-based heart rate estimation from facial video images," in *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, 2024, pp. 1–5.
- [5] L. Zhou, Z. Liu, Z. Shanguan, X. Yuan, Y. Li, and B. Hu, "Tamfn: Time-aware attention multimodal fusion network for depression detection," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 31, pp. 669–679, 2023.
- [6] L. He, K. Chen, J. Zhao, Y. Wang, E. Pei, H. Chen, J. Jiang, S. Zhang, J. Zhang, Z. Wang, T. He, and P. Tiwari, "Lmvd: A large-scale multimodal vlog dataset for depression detection in the wild," *Information Fusion*, vol. 126, pp. 103632, 2026.
- [7] J. Yoon, C. Kang, S. Kim, and J. Han, "D-vlog: Multimodal vlog dataset for depression detection," in *AAAI Conference on Artificial Intelligence*, 2022.
- [8] Y. Tao, M. Yang, Y. Wu, K. Lee, A. Kline, and B. Hu, "Depressive semantic awareness from vlog facial and vocal streams via spatio-temporal transformer," *Digital Communications and Networks*, vol. 10, no. 3, pp. 577–585, 2024.
- [9] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," in *Conference on Language Modeling*, 2024.
- [10] L. Zhu, B. Liao, Q. Zhang, X. Wang, W. Liu, and X. Wang, "Vision mamba: efficient visual representation learning with bidirectional state space model," in *41st International Conference on Machine Learning*, 2024.
- [11] B. Zhou, M.-A. Fiedler, and A. Al-Hamadi, "Stcm-mamba: Multimodal spatio-temporal cross-modal mamba for depression detection," *IEEE Access*, vol. 13, pp. 189494–189505, 2025.
- [12] B. Zhou, M.-A. Fiedler, and A. Al-Hamadi, "Cross-modal fusion mamba for multimodal depression detection," in *2025 14th International Symposium on Image and Signal Processing and Analysis (ISPA)*, 2025, pp. 225–230.
- [13] J. Ye, J. Zhang, and H. Shan, "Depmamba: Progressive fusion mamba for multimodal depression detection," in *2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025.
- [14] M. Liu, Z. Wang, and F. Liu, "Mdmamba: A model for multimodal depression detection with a memory-saving cross-modal attention mechanism based on mamba," in *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2024, pp. 1600–1605.
- [15] W.A. Gadzama, D. Gabi, M.S. Argungu, and H.U. Suru, "The use of machine learning and deep learning models in detecting depression on social media: A systematic literature review," *Personalized Medicine in Psychiatry*, vol. 45–46, pp. 100125, 2024.
- [16] W. B. Tahir, S. Khalid, S. Almutairi, M. Abohashrh, S. A. Memon, and J. Khan, "Depression detection in social media: A comprehensive review of machine learning and deep learning techniques," *IEEE Access*, vol. 13, pp. 12789–12818, 2025.
- [17] P. Ekman and W.V. Friesen, *Facial Action Coding System: A Technique for the Measurement of Facial Movement*, Consulting Psychologists Press, 1978.
- [18] A. Gu, K. Goel, and C. Ré, "Efficiently modeling long sequences with structured state spaces," in *International Conference on Learning Representations (ICLR)*, 2022.
- [19] A. Pampouchidou, O. Simantiraki, C.-M. Vazakopoulou, C. Chatzaki, M. Padiaditis, A. Maridaki, K. Marias, P. Simos, F. Yang, F. Meriaudeau, and M. Tsiknakis, "Facial geometry and speech analysis for depression detection," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2017, pp. 1433–1436.
- [20] M.A. Hearst, S.T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, "Support vector machines," *IEEE Intelligent Systems and their Applications*, vol. 13, no. 4, pp. 18–28, 1998.
- [21] D.W. Hosmer, S. Lemeshow, and R.X. Sturdivant, *Applied Logistic Regression: Third Edition*, Wiley, 2013.
- [22] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, Oct. 2001.
- [23] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1800–1807.
- [24] J. Hu, L. Shen, S. Albanie, G. Sun, and E. Wu, "Squeeze-and-excitation networks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7132–7141.
- [25] S. Yin, C. Liang, H. Ding, and S. Wang, "A multi-modal hierarchical recurrent neural network for depression detection," in *9th International on Audio/Visual Emotion Challenge and Workshop*, 2019, pp. 65–71.
- [26] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021.
- [27] A. Zadeh, M. Chen, S. Poria, E. Cambria, and L.-P. Morency, "Tensor fusion network for multimodal sentiment analysis," in *Conference on Empirical Methods in Natural Language Processing*, 2017.
- [28] Davis King, "Dlib-ml: A machine learning toolkit," *Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 07 2009.
- [29] F. Eyben, M. Wöllmer, and B. Schuller, "Opensmile: the munich versatile and fast open-source audio feature extractor," in *18th ACM International Conference on Multimedia*, 2010, p. 1459–1462.