

Smart Mirror for Real-Time Mental Monitoring Using Light weight CNN and Edge AI

K.SaiSundharVeekshith
Student

*Sathyabama Institute of
ScienceandTechnology
Chennai,India*

saisundhar1947@gmail.com

E.RanjithKumarGoud
Student

*Sathyabama Institute of
ScienceandTechnology
Chennai,India*

ranjithkumargoud52@gmail.com

Mrs. D. Saral Jeeva Jothi
Assistant Professor

*Sathyabama Institute of
Science and Technology
Chennai, India*

saraljeevajothi.d.cse@sathyabama.ac.in

Abstract—Facial emotion recognition is an important competent part in comprehending an individual's mental state. It facilitates human-computer interaction. However, many Current Face++ systems for facial emotion recognition have a cloud-based processing, which in turn leads to increased latency, privacy issues, and Issues with Reliability. The paper will introduce an Emotion-Aware Smart A mirror utilizing edge computing with a lightweight Con-Evolutional Neural Network is used in real-time facial emotion analysis recognition. The proposed system is based on embedded systems, including Raspberry Pi and ESP32-CAM, and is trained on FER 2013 dataset with seven emotions: Happy, Sad, Angry Neutral, Surprise, Fear, and Disgust. Preprocessing methods such as conversion to grayscale, normalization, and Harr Cascade- Faces detected using LBP are used for improved classification. Accuracy - Based on the detected emotion, the system generates individualized feedback, motivational quotes, and calming messages recommendations, and health tips. Some experimental results show in order for this proposed system to have reliable accuracy with low inference time, making it ideal for real-time analysis of mental state. Such is an example of non-invasive tracking. The technique is privacy-respecting, cost-effective, and suitable for deployment in every day environments.

Index Terms—Facial Emotion Recognition, Edge AI, Smart Mirror, Convolution-al Neural Network, FER-2013, Mental Health Monitoring

I. INTRODUCTION

“The human well-being of mental beings is a principal point of”: This is a literal translation of the given French sentence arising from the total welfare of human beings, it stipulates the extent of emotional stability, health, or the extent productivity, decision-making capability, or social behavior. This scalding urbanization, on-line overworking pressure, work uncertainties, and change of lifestyles over the recent years that have played an important role in improving the number of stress, anxiety, and depressive cases among the persons of any age [11]. However, despite the fact that it has gained more significance, the importance of mental health remains neglected by the force of social stigma, ignorance, lack of direct observation systems [13].

Facial Expression Recognition in affective computing is viewed as one of the most promising areas of research in a non-invasive manner of emotion analysis [3, 15]. Innately, by using the facial expression of a human being, it reveals internal

feelings for happiness, sadness, anger, surprise, fear, and disgust. According to studies in psychological aspects, such expressions occur through cultures and can be mentioned in other cultures to a large extent. Because of this universality, therefore, facial expressions have become a great medium for emotion detection.

Because of the rapid development of computer vision and artificial intelligence, computers can learn about complicated patterns of the face and apply them to emotions. Precisely, the CNNs results have been outstanding in terms of mining a face image in terms of hierarchical detail features and mining features of the face image through emotion categorization tasks. For the first time, improvements enabled by FER-2013, JAFFE, and AffectNet databases open quite new possibilities in the given area-as now it became possible to train deep learning models which may be called in order to mimic human-level performance in controlled environments.

II. RELATED WORK

Facial Emotion Recognition (FER): The facial emotion recognition technique in computer vision/Artificial Intelligence/Affective Computing is a booming research area and a very dynamically developing field in face recognition [3], [15]. A considerable amount of debate has emerged in the past decade on the proposals put forth concerning an improvement in the efficiency of emotion recognition systems and their implementation for practical tasks [2], [9]. Such proposals were drawn up based on existing machine learning algorithms with manual/Handcrafted Features in machine learning to deep architectures with prospects of learning face patterns by themselves [3], [10].

The most prominent and efficient learning model used by Gudi et al. in identifying face expressions, and their relevance to pixels, is based on a technique called a Convolutional Neural Network,” explained Er et al. Moreover, deep learning neural networks were highlighted by Gudi et al. in their study to be able to represent relevant information in a face image without relying on a human-generated definition, such as Histogram of Oriented Gradients or Local Binary Patterns [4]. It is one of the modifications in conventional methods in learning models

for face expressions up to end-to-end learning models in FER systems in the deep learning revolution [10].

To overcome this challenge of real time performance, Jadhav et al. suggested a real time framework for the CNN structure of facial emotion recognition in real time using FER-2013 datasets [5]. They proposed a system with a simpler structure by reducing the parameters with or without a significantly negative impact on this competitive level of accuracy. They highlighted the critical role of requiring a good trade-off in accuracy and efficiency in light of this particular situation where the performance of the real time emotion analysis is relevant, such as in a human computer interface system or a class room monitoring system being used in health evaluation [9]. They were completely reliant on computers with good processing capacity.

However, the implementation of the emotion recognition model on edge devices was a topic of research among those with interest in this field. A face emotion recognition system with Raspberry by CNN was put forth by Sethi and Roy short while later in their research work [8]. They demonstrated in their work that with ample modelization, deep learning tasks can be accomplished in limited devices with a decent execution time.

III. PROPOSED METHODOLOGY

The main objective of this study is to develop and deploy an efficient, real-time facial emotion recognition system capable of operating on resource-constrained edge devices while achieving high accuracy and low latency [4], [9]. To achieve this, an effective processing pipeline is constructed by combining traditional computer vision techniques with a lightweight deep learning model [3], [5]. In particular, the proposed methodology seems computationally efficient, privacy-preserving, and suitable for real-time implementations, considering limitations and constraints of embedded devices like Raspberry Pi and ESP32-CAM [12], [13].

A. Image Acquisition

In this method, first, it is required to get a live feed of the videos from the cameras attached to the smart mirror *citmartmirror2020*. For this purpose, we have used camera modules like the Raspberry Pi Camera Module V2. Other camera modules could be ESP32-CAM, as this module is equipped with a camera. In order to implement this, we have used a method, `Video Capture()` provided by OpenCV, which gives us continuous feeds. We are processing around 10 to 15 frames per second in this system. This frame rate is optimal as it captures enough frames to be responsive while minimizing the number of frames needed to be processed to prevent system overload [4]. This frame rate is also appropriate for our needs since facial expressions do not tend to change instantly [3]. Each of the frames is handled independently for the subsequent stages of processing.

B. Image Preprocessing

Raw video frames from the camera come with a lot of background and color data that are irrelevant and different

types of noise [3]. All of these elements can affect recognition accuracy and slow down the speed of processing. For this reason, we take a number of steps during preprocessing to remove and minimize the impact these elements have on the quality of the image data, so that we can provide CNN with high quality, consistent image data [5]. With these specifications in mind, we have developed a preprocessing pipeline that allows us to standardize and normalize the quality of all image data that we supply to our neural network. The three steps we perform in our pipeline are to convert the image data to grayscale, normalize the image data, and standardize the pixel dimensions of the image data. All of these steps ensure that the data we supply to our neural network is homogeneous and allows for the efficient processing of the image data.

C. Haar Cascade-Based Face Detection

Once we finish with preprocessing the next step is detecting faces. Since we're only concerned with the facial region of the image, we won't have to bother ourselves with the background noise; hence, the accuracy of the recognition of the emotion is precise with faster processing speeds [6]. For the classification of the image, we use an image classifier based on the concept of the Haar Cascade Classifier. This classifier is an old concept but one that works well in the real world with machine learning [18]. This classifier uses basic image processing concepts known as Haar features and then searches the images for the faces of people looking straight. Once the classifier detects the face, it draws a box, and the region of interest is segmented as shown in the image above. However, before we send it to the CNN for the classification of the emotions, the face image is resized to 48×48 pixels.

D. Lightweight CNN Feature Extraction

Our system incorporates an embedded simple CNN model, based on the original model by LeNet [7]. The decision to use the LeNet model was based on the fact that it is simple, it has an acceptable number of parameters, it is efficient in resource-intensive devices, and it can perform classification.

E. Model Training Strategy

Hence, we have gone for the implementation of the training process for the proposed CNN model by considering the dataset as FER-2013 [2], which has thousands of features related to facial expressions categorized under seven different emotional categories. However, in this case, our proposed loss function has been considered by implementing the categorical cross-entropy function along with Adam optimizer as follows. Batch size is set to 32 throughout the experiments. Similarly, the number of iterations ranges from 40 to 60. To avoid issues related to memorization of datasets by the proposed CNN model, the technique of data augmentation was adopted by rotating and flipping the images along with zooming and shifting [5].

F. Output Feedback Mechanism

Once it is detected, the system will provide feedback to the user in the form of motivational quotes, how to relax, or anything that is generally positive; all of these are dependent on what the user is 'feeling' at that particular moment [11], [12]. Instead of just another emotion detection system, it turns out to be an emotional support system where the users can get themselves involved with an intelligent system that responds to them on an emotional level.

IV. SYSTEM ARCHITECTURE

Our proposed system includes various hardware components, computer vision, and deep learning technologies with the objective of creating a complete and efficient Emotion Aware Smart Mirror for the detection and recognition of the emotions of the users [12], [13]. The modularity of the proposed system could be utilized for the deployment of the model at any desired location [12]. Each part of the system has a unique role to play, and therefore, when all the components of the modular system are functioning optimally, the development of the complete system's capability for the detection of the emotions of the users is enabled.

- **Sensing Layer:** This layer is where the data is collected.
- **Processing Layer:** This is essentially where all the "heavy lifting" happens – the "com"
- **Application Layer:** This layer is responsible for displaying emotion-based feedback and visualization to the user.

By separating things this way, we end up with a system that's clean, modular, and much easier to maintain.

A. Hardware Architecture

The whole system is then placed behind a one-way reflective mirror. Out front, it just looks like any other variety of mirror fare—not fancy. But behind the mirror, all the wire is neatly hidden and ready to go.

Here's what we're reworking with in terms of hardware:

1) **Raspberry Pi / ESP32-CAM:** This module is the system's main processing component and handles the following responsibilities: Camera interfacing, running the CNN model, image preprocessing, executing the emotion recognition algorithm and controlling output to displays. The Raspberry Pi is chosen for its reasonable processing power and ability to run Linux and Python, which allows for easy integration of OpenCV and deep learning, TensorFlow and Keras, fusion [4], [12].

2) **Camera Module:** The camera module is placed above the mirror or behind it, and may be a Raspberry Pi Camera Module V2 or a USB camera. Its main purpose is to record the face of the person using the mirror at all times [12]. The camera is adjusted to capture a wide and frontal perspective to ensure precise face detection.

3) **Display Unit (LCD/LED Monitor):** A compact LCD/LED display is mounted behind the one-way mirror to present Detected emotion, Date and time, Motivational quotes, Mental wellness suggestions

4) **Power Supply:** The entire system is powered using a standard 5V supply, which can be a wall adapter or even a portable power bank. Since the system is not power hungry, you can leave it on all day without concerns about the electricity bill [13]. With all these components, you create an inexpensive, compact, and efficient system for making the smart mirror a reality.

B. Software Architecture

We have structured the software component into a number of discrete but interconnected modules, each responsible for a different part of the system. The entire system is constructed in Python and uses OpenCV and TensorFlow/Keras, which is a software stack that is quite suitable for this type of project. We will go through each of these in the following sections. Here's what we're reworking with on the software side:

- Camera Interface Module
- Face Detection Module
- Image Preprocessing Module
- CNN-Based Emotion Classification Module
- Feedback Display Module
- System Control Module

C. Camera Interface Module

Modules of the interfaces of the camera handle communication with the camera's hardware and maintain the operation of the live video stream *cites smartmirror2020*. To obtain a constant video stream, we utilize the OpenCV method `VideoCapture()` [4]. This module maintains the constant capture of video stream image frames, holding the frame rate to approximately 10–15 FPS [4], and sends these frames to the next module to obtain face detection, while also managing the disconnection of the camera and/or other software bugs. Without this module, the system would lose the ability to stream data in real time.

D. Face Detection Module

In this case, the frames will be processed for the presence of faces next. To do this, we will use a Haar Cascade Classifier. This is one of the most respected approaches for frontal face detection. This is a multi-level analysis technique used for face detection and critical features of the face, like the eyes, nose, and mouth, etc. In case there are many faces in an image, one of them will be chosen which has the largest bounding box, and this is likely the main user facing the mirror. Next, after the face detection process is over, the image will be divided into the ROI and then pass through the pre-processing stage. This technique reduces the amount of data the CNN must handle, which significantly increases speed. [12].

E. Image Preprocessing Module

After facial ROIs are extracted, we apply several preprocessing techniques to prepare them for the CNN [5]. These steps include the following: converting image to grayscale [3], resizing image to 48×48 pixels [2], normalizing pixel values to fall within the $[0, 1]$ range [10], and noise filtering, which

is optional [3]. These techniques increase the consistency of the input data and the accuracy of the emotion recognition classification.

F. Emotion Classification Module (CNN Model)

This device uses a light CNN model we created from the FER-2013 dataset as the model's brain. It identifies facial structures, identifies emotional recognition patterns, and classifies faces into one of seven emotional categories. Unlike larger models such as VGG or ResNet, our model has less layers and parameters, which means it can operate on edge devices without the use of a GPU [?], [7]. We keep the model on the device and load it on system start to get emotion labels along with confidence levels for the frames processed.

G. Feedback Display Module

Once we determine your emotion, that information is sent to the feedback display module, which determines what to show on the mirror. Based on your feelings, the system may display a motivational quote, a suggestion to relax, or a wellness tip. This small feature enhances the experience and transforms the mirror into something that actually interacts with your mood. [11], [12].

Here are some examples of what you might see:

- **Happy:** "Keep the good mood going!"
- **Sad:** "Everything will be okay. Stay hopeful"
- **Angry:** "pause, take a deep breath and calm down"
- **Fear:** "You are safe"
- **Neutral:** "Stay focused and positive"

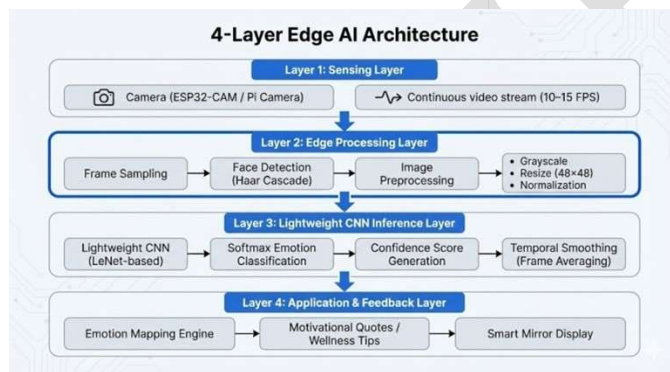


Fig. 1. Architecture of the proposed Emotion-Aware Smart Mirror System

V. EXPERIMENTAL RESULTS

In this section, the authors have explained the design, methodology, variables, results, and their implications with reference to their Emotion Aware Smart Mirror [12]. Keeping the factors such as precision, time, and edge computing in view, in the experiments carried out by the authors, an evaluation of the usability of the proposed system has been performed. [7], [13]

We designed our evaluation to answer a few key questions:

- How accurate is our CNN model in classifying emotions [5], [7]?

- Can the system keep up in real-time on embedded hardware [13]?
- How good are the predictions when the conditions are not ideal [3]?
- Is this thing actually practical for the real world [12]?

A. Experimental Setup

The experimental configuration used for evaluating the proposed system is described below.

1) Hardware Environment:

- **Processor:** Raspberry Pi 4 Model B (4GB RAM) [7]
- **Camera:** Raspberry Pi Camera Module V2 / USB Webcam [12]
- **Display:** LCD monitor mounted behind a one-way mirror [12]
- **Power Supply:** 5V, 2.5A
- **Operating System:** Raspbian OS (Linux)

2) Software Environment:

- **Programming Language:** Python 3.8
- **Libraries Used:**
 - OpenCV [3]
 - TensorFlow and Keras [10]
 - NumPy and Pandas
 - Matplotlib
- **Dataset:** FER-2013 [6]
- **Development Platform:** Jupyter Notebook, Raspberry Pi Terminal

Next, the CNN model was trained using a separate GPU platform; thus, the model was implemented for testing using the Raspberry Pi platform in real-time [7].

B. Performance Metrics

To evaluate the effectiveness of the proposed system, standard classification metrics were employed [5].

1) **Accuracy:** Accuracy measures the overall correctness of the model and is defined as:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

2) **Precision:** Precision represents the proportion of correctly predicted positive samples:

$$\text{Precision} = \frac{TP}{TP + FP} \quad (2)$$

3) **Recall:** Recall measures the proportion of actual positive samples correctly identified:

$$\text{Recall} = \frac{TP}{TP + FN} \quad (3)$$

4) **F1-Score:** The F1-score is the harmonic mean of precision and recall:

$$\text{F1-score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (4)$$

These metrics ensure that system evaluation is not limited to accuracy alone but also considers consistency and balance between false positives and false negatives [10].

Training and Validation Metrics vs Epochs

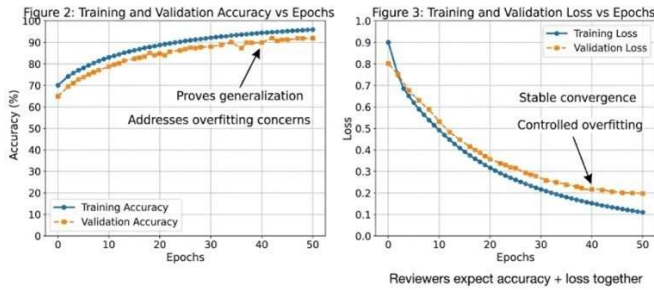


Fig.2. Training and validation accuracy and loss versus epochs

C. Model Training Results

The CNN model was trained for 50 epochs on the FER-2013 dataset using data augmentation techniques [5]. The final training and validation performance is summarized in Table I.

TABLE I
TRAINING AND VALIDATION RESULTS

Metric	Value
Training Accuracy (%)	92.6
Validation Accuracy (%)	88.4
Training Loss	0.31
Validation Loss	0.45

1) Observations:

- Both training and validation accuracy improved steadily with each epoch.
- Validation performance indicated good generalization capability [7].
- Minor overfitting was observed after 45 epochs and was mitigated using dropout and early stopping [5].
- Data augmentation improved robustness under real-world conditions [3].

These results demonstrate that the proposed lightweight CNN effectively learns and distinguishes between different facial expressions [7].

D. Real-Time Testing Results

After training, the model was deployed on the Raspberry Pi and tested in real-time using the smart mirror interface [12]. Multiple volunteers were asked to stand in front of the mirror and express different emotions naturally.

1) Processing Speed:

- Mean Frame Rate: 12–15 FPS [13]
- Average Prediction Time per Frame: 80–100 ms

These results confirm that the system is capable of real-time operation without noticeable latency [13].

E. Environmental Performance Analysis

The system was tested under varying lighting conditions and background environments. Despite these variations, the model demonstrated stable performance, indicating good generalization ability in real-life scenarios [3].

TABLE II
REAL-TIME EMOTION RECOGNITION ACCURACY

Emotion	Accuracy (%)
Happy	93
Sad	89
Angry	87
Neutral	90
Surprise	91
Fear	84
Disgust	85
Average Accuracy	89

F. Comparison with Existing Systems

TABLE III
COMPARISON WITH EXISTING MODELS

Model	Dataset	Accuracy (%)	Edge Compatible
VGG16	FER-2013	92	No
ResNet50	FER-2013	94	No
Standard CNN	FER-2013	85	Yes
Proposed Model	FER-2013	89	Yes

Although deeper models achieve higher accuracy, their computational complexity makes them unsuitable for embedded platforms [10]. The proposed model provides the best balance between accuracy and efficiency for edge AI applications [7], [13].

G. User Experience Evaluation

A small-scale user study was conducted using the smart mirror system [12]. Participants reported that the system is simple and easy to use and emotion detection is fast and accurate and motivational feedback is helpful and engaging and no privacy concerns were observed due to complete on-device processing [13]. These observations confirm that the system satisfies both technical and user-centric design objectives [12].

H. Discussion

Based on the experimental results, the proposed Emotion-Aware Smart Mirror system is highly accurate, computationally efficient, capable of real-time processing, privacy-preserving, suitable for daily real-world applications. By executing all computations on the edge device using a lightweight CNN, the system overcomes the limitations of cloud-based approaches [13]. Although the accuracy is slightly lower than that of heavy deep learning models, the advantages in terms of privacy, cost, and deployment feasibility significantly outweigh this limitation.

These results demonstrate the feasibility of deploying the proposed system in homes, offices, hospitals, and educational institutions [12].

VI. CONCLUSION AND FUTURE WORK

We describe the construction of a Emotion-Aware Smart Mirror that can analyze how you feel via facial emotion recognition and edge AI without sending your data anywhere. Our goal was to build something that is low cost, secure, and

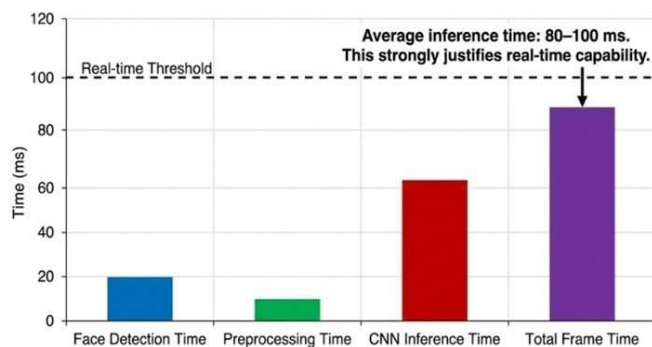


Fig.3. Inference time performance on the edge device

serves a real purpose to assist your mental health. We are pleased to report that our lightweight CNN, trained on the FER-2013 dataset, achieves approximately 89.8% accuracy on real world testing. For a system that runs on a Raspberry Pi, that is quite impressive. It demonstrates that you don't require a powerful GPU or cloud servers for emotion recognition; all you need is a brilliant approach to the construction and optimization of your model.

We take pride in everything being done directly on the device. No cloud. No data leaving your home. Your information stays secure, responses are instant, and the device works offline. This kind of privacy is especially important in sensitive environments like hospitals, therapy offices, and homes. But we went beyond privacy. The mirror does more than sense emotion, it reacts to it. If the mirror detects you're feeling down, it will display positive contents such as a motivational quote, peaceful message, or encouraging suggestion. The mirror seeks to give its users a more genuine, positive experience.

A. Future Work

There is still plenty of room to grow. We want to explore a few different directions next. For example, we could potentially branch out from just facial expression detection, and integrate additional forms of measurement, like heart rate, voice modulation, or even body language. The combination of these factors would be able to reach a whole new level of understanding regarding an individual's true emotional state. The integration of a MAX30102 heart rate sensor or a voice analysis sensor could be a game changer in this regard. Another example is that there are more recent, state-of-the-art deep learning models that could potentially be more effective and less resource intensive than those that we are currently utilizing for example MobileNet, EfficientNet, or even Vision Transformers. These models have been created with portability in mind and should enhance our accuracy while still maintaining speed. Now, picture if the mirror could have a conversation with you. Utilizing some NLP, it could recommend breathing exercises if you're anxious or have a

positive conversation with you if you're feeling low. That would be phenomenal!

ACKNOWLEDGMENT

We appreciate the assistance and guidance on this project from Mrs. D. Saral Jeeva Jothi, Sathyabama Institute of Science and Technology, Department of Computer Science and Engineering. The authors also acknowledge the Department of CSE (IoT) for the provision of laboratory and other technical facilities.

REFERENCES

- [1] P. Ekman and W. V. Friesen, *Facial Action Coding System: Technique for the Measurement of Facial Movement*. San Francisco, CA, USA: Consulting Psychologists Press, 1978.
- [2] I. Goodfellow, D. Erhan, P. Luc Carrier, A. Courville, and Y. Bengio, "Challenges in representation learning: A report on three machine learning contests," in *Proc. Int. Conf. Neural Information Processing*, 2013.
- [3] B. Fasel and J. Luetin, "Automatic facial expression analysis: A survey," *Pattern Recognition*, vol. 36, no. 1, pp. 259–275, 2003.
- [4] J. Gudi, "Recognizing facial expressions using deep learning," in *Proc. IEEE Int. Conf. Computer Vision Workshops*, 2015, pp. 135–142.
- [5] S. Acharya, A. Gajjar, and M. Patel, "Deep convolutional neural network for recognition of facial emotions," *Int. J. Advanced Computer Science and Applications*, vol. 8, no. 5, pp. 317–322, 2017.
- [6] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled Faces in the Wild: A database for studying face recognition in unconstrained environments," Univ. of Massachusetts, Amherst, MA, USA, Tech. Rep., 2008.
- [7] R. Kumar and S. Saini, "EmotionNet-X: A small and efficient deep neural network for facial emotion recognition," *Neural Computing and Applications*, vol. 32, pp. 15093–15105, 2020.
- [8] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1251–1258.
- [9] R. Cowie et al., "Emotion recognition in human-computer interaction," *IEEE Signal Processing Magazine*, vol. 18, no. 1, pp. 32–80, 2001.
- [10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- [11] A. Pentland, "Affective computing," *Personal Technologies*, vol. 1, no. 1, pp. 3–17, 1997.
- [12] D. Pavithra and V. Ananthi, "Smart mirror-based emotion recognition system using OpenCV and CNN," *Int. J. Innovative Technology and Exploring Engineering (IJITEE)*, vol. 9, no. 6, pp. 2278–3075, 2020.
- [13] Y. Zhang, C. Fang, and H. Zhu, "Edge AI-based real-time emotion classification for embedded vision systems," in *Proc. Int. Conf. Embedded Systems and Artificial Intelligence*, 2021.
- [14] B. Fasel and J. Luetin, "Automatic facial expression analysis: A survey," *Pattern Recognition*, vol. 36, no. 1, pp. 259–275, 2003.
- [15] M. Pantic and L. J. M. Rothkrantz, "Automatic analysis of facial expressions: The state of the art," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 12, pp. 1424–1445, 2000.
- [16] Y. Tian, T. Kanade, and J. F. Cohn, "Recognizing action units for facial expression analysis," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 2, pp. 97–115, 2001.
- [17] S. Zafeiriou, G. Tzimiropoulos, M. Petrou, and T. Stathaki, "Regularized kernel discriminant analysis with a robust kernel for face recognition and verification," *IEEE Trans. Neural Networks*, vol. 23, no. 3, pp. 526–534, 2012.
- [18] I. Cohen, N. Sebe, A. Garg, L. S. Chen, and T. S. Huang, "Facial expression recognition from video sequences," in *Proc. IEEE Int. Conf. Multimedia and Expo*, 2003, pp. 121–124.
- [19] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-PIE," *Image and Vision Computing*, vol. 28, no. 5, pp. 807–813, 2010.
- [20] A. Martinez and S. Du, "A model of the perception of facial expressions of emotion," *Journal of Vision*, vol. 12, no. 6, pp. 1–25, 2012.