

# Emotion Detection from Facial Expressions Using Deep Learning

N. Vinod, M. Vinuthna, A. Vishal, K. Vishal, K. Vishal (Students) Prof. S. Jim Reeves (Guide)  
School of Engineering Department of AIML. Malla Reddy University, Hyderabad.

## 1. ABSTRACT:

Facial expression detection is a crucial aspect of human-computer interaction, enabling machines to recognize and respond to human emotions effectively. This paper presents a deep learning-based face expression detection system that classifies human emotions based on facial expressions. By leveraging convolutional neural networks (CNNs) and advanced image processing techniques, the system achieves high accuracy in detecting expressions such as happiness, sadness, anger, surprise, and neutrality. The research explores various datasets, training methodologies, and evaluation metrics to optimize performance.

**Keywords**— Facial expression detection, deep learning, convolutional neural networks, emotion recognition, image processing, human-computer interaction, sentiment analysis, machine learning, facial recognition.

## 2. INTRODUCTION:

Facial Expression Recognition (FER) is vital in human-computer interaction, healthcare, security, and customer experience analysis. Accurately detecting and classifying facial expressions is crucial for developing intelligent systems capable of understanding human emotions and responding appropriately. However, achieving high accuracy in FER remains a significant challenge due to several factors, including variations in facial features, lighting conditions, occlusions, head poses, and individual differences. Traditional methods often struggle to generalize across diverse scenarios, making it difficult to deploy robust FER systems in real-world applications. One of the major challenges in FER is the variability in facial expressions across different individuals. Factors such as age, gender, ethnicity, and unique facial structures contribute to inconsistencies in facial feature representation. This diversity makes it challenging for models to learn generalizable patterns that can accurately classify expressions across different demographic groups. Additionally, variations in intensity and subtlety of expressions further complicate the recognition process, as some emotions may manifest in barely perceptible facial changes.

External environmental

factors also play a critical role in FER accuracy. Changes in lighting conditions can significantly impact the visibility of facial features, leading to misclassifications. Shadows, reflections, and uneven illumination can alter the appearance of facial landmarks, making it difficult for models to extract meaningful features. Furthermore, occlusions caused by accessories such as glasses, masks, or hands covering parts of the face introduce additional complexity, as key facial regions necessary for emotion recognition may be partially or entirely hidden. Another significant challenge is the presence of head pose variations in real-world FER applications. Unlike controlled environments where facial images are captured in standardized frontal poses, real-life scenarios involve individuals moving their heads freely, resulting in a wide range of facial orientations. Extreme head tilts, side profiles, and partially visible faces pose difficulties for traditional FER approaches, which often rely on fixed facial alignment assumptions. Addressing these variations requires the development of robust models capable of recognizing expressions from multiple angles and partial facial views. The issue of data imbalance further complicates FER model training. In many publicly available FER datasets, certain facial expressions, such as neutral and happy, are overrepresented, while others, such as disgust or fear, appear infrequently. This imbalance can lead to biased models that perform well on dominant classes but struggle with underrepresented expressions. Consequently, the overall recognition performance suffers, particularly for less common emotions, resulting in unreliable classification outcomes in practical applications. Computational complexity is another major obstacle in deploying FER systems for real-time applications. Deep learning models, particularly Convolutional Neural Networks (CNNs), have significantly improved recognition accuracy, but their high computational requirements make real-time processing challenging, especially on resource-constrained devices such as mobile phones and embedded systems. The need for large-scale datasets, extensive training time, and powerful hardware limits the feasibility of deploying FER models in real-world

applications where speed and efficiency are paramount.

### 3. LITERATURE REVIEW:

Facial Expression Recognition (FER) is a critical domain within computer vision and artificial intelligence, focusing on the identification and analysis of human emotions through facial cues. Applications of FER span human-computer interaction, psychological studies, and security systems. This survey examines the evolution of FER methodologies, emphasizing significant research contributions and their impacts.

Traditional Approaches to FER :

Early FER systems predominantly utilized handcrafted feature extraction techniques.

- **Local Binary Patterns (LBP):** Local Binary Patterns (LBP) is a texture descriptor that transforms an image into a binary pattern based on local pixel comparisons. Ojala et al. introduced LBP for texture classification, and it was later adapted for FER due to its computational efficiency. However, LBP's sensitivity to lighting variations limited its robustness in uncontrolled environments.

- **Histogram of Oriented Gradients (HOG):** Histogram of Oriented Gradients (HOG) captures gradient orientation histograms, effectively describing local object appearances. Dalal and Triggs demonstrated HOG's efficacy in human detection, and subsequent studies applied it to FER. Despite its strengths, HOG struggled with pose variations and occlusions. Active Appearance Models (AAM) and

- **Active Shape Models (ASM):** AAM and ASM, introduced by Cootes et al., model facial shape and texture through statistical analysis. These models adjust parameters to fit facial landmarks but often falter under significant pose changes and occlusions. 9 Deep Learning for FER The advent of deep learning, particularly Convolutional Neural Networks (CNNs), revolutionized FER by enabling automatic feature extraction.

### 4. Problem Statement:

Facial Expression Recognition (FER) plays a crucial role in various applications, including human-computer interaction, healthcare, security, and customer experience analysis. However, achieving high accuracy in FER remains a challenge due to factors such as variations in facial expressions across individuals, environmental conditions, occlusions, head pose variations, and data imbalance in existing

datasets. Traditional handcrafted feature extraction methods struggle with generalization and robustness in real-world scenarios.

Despite advancements in deep learning, particularly Convolutional Neural Networks (CNNs), FER systems still face limitations in accurately recognizing subtle and complex emotions, especially in unconstrained environments. Furthermore, the high computational requirements of deep learning models make real-time deployment on resource-constrained devices challenging.

This research aims to develop a deep learning-based facial expression detection system that overcomes these challenges by leveraging advanced CNN architectures and optimized training methodologies. The objective is to enhance recognition accuracy, improve robustness against environmental and pose variations, and enable real-time performance for practical applications.

### 5. Methodology:

The methodology for developing the Facial Emotion Recognition (FER) system was structured to ensure a systematic approach to data preparation, model development, training, evaluation, and deployment. This section outlines the step-by-step process undertaken to implement the proposed architecture, preprocess the data, train the model, and deploy it for real-time emotion detection. The methodology leverages deep learning frameworks, computer vision techniques, and optimization strategies to achieve accurate and efficient emotion recognition.

#### 5.1 Data Collection and Preparation:

The first step involved acquiring the FER2013 dataset, a widely recognized benchmark for facial emotion recognition tasks. The dataset comprises 35,887 grayscale images of size 48x48 pixels, categorized into seven emotions: angry, disgust, fear, happy, neutral, sad, and surprise. The dataset was split into training (28,709 images), validation (3,589 images), and test (3,589 images) sets, as provided by the dataset creators. To prepare the data for training, a custom Python script was developed using the pandas library to create a DataFrame containing image file paths and their corresponding labels. This was achieved by iterating through the dataset directory structure, where each emotion category was stored in a separate subdirectory. The created dataframe function, implemented in Python, extracted file paths and labels, ensuring that the data was organized for subsequent preprocessing steps. The dataset was inspected for missing or corrupted images,

and any such instances were removed to ensure data quality. Additionally, the class distribution was analyzed to confirm the presence of class imbalance, which was later addressed through preprocessing techniques.

### 5.2. Data Preprocessing Implementation:

Data preprocessing was performed to transform the raw images into a format suitable for training the Convolutional Neural Network (CNN). The preprocessing pipeline, as described in the Design section, was implemented using Python with libraries such as OpenCV, numpy, and tensorflow.keras.preprocessing. First, grayscale conversion was applied to all images, as they were already in grayscale format in the FER2013 dataset, ensuring consistency with the model's input requirements. Face detection and cropping were then performed using OpenCV's Haar Cascade Classifier (haarcascade\_frontalface\_default.xml) to isolate facial regions, although this step was primarily used during real-time detection since FER2013 images are pre-cropped. Image normalization was implemented by dividing pixel values by 255.0, scaling them to the range [0, 1], which was executed using numpy operations. Data augmentation was applied to the training

set using tensorflow.keras.preprocessing.image.ImageDataGenerator, with transformations including rotation (up to 10 degrees), horizontal flipping, brightness adjustments (range 0.8 to 1.2), and zooming (up to 10%). These augmentations were applied on-the-fly during training to increase dataset diversity and mitigate overfitting. The preprocessed images were reshaped to (48, 48, 1) to match the input shape of the CNN, and the labels were encoded using sklearn.preprocessing.LabelEncoder followed by one-hot encoding with tensorflow.keras.utils.to\_categorical to prepare them for multi-class classification.

### 5.3. Model Architecture Implementation:

The proposed CNN architecture, detailed in the Design section, was implemented using the tensorflow.keras framework in Python. The model was constructed as a Sequential model, starting with an input layer accepting grayscale images of shape (48, 48, 1). Four convolutional blocks were added, each consisting of a Conv2D layer with increasing filters (128, 256, 512, 512), kernel size (3, 3), and ReLU activation, followed by batch normalization (BatchNormalization), max pooling (MaxPooling2D with pool size (2, 2)), and

dropout (Dropout with rates 0.4). The Conv2D layers used the GlorotUniform initializer for weights and Zeros for biases, with padding set to 'valid'. After the convolutional blocks, the feature maps were flattened using a Flatten layer. Two fully connected (Dense) layers were added with 512 and 256 units, respectively, each followed by ReLU activation and dropout (0.4 and 0.3, respectively) to prevent overfitting. The output layer was a Dense layer with 7 units (one for each emotion) and a softmax activation to produce a probability distribution over the emotion classes. The model architecture was verified using model.summary() to ensure the correct number of parameters and output shapes at each layer.

### 5.4. Model Training Process:

The model was trained using the preprocessed FER2013 dataset, following the training methodology outlined in the Design section. The training process was implemented in Python using tensorflow.keras. The model was compiled with the Adam optimizer (tf.keras.optimizers.Adam) with a learning rate of 0.001, Categorical Cross-Entropy as the loss function (categorical\_crossentropy), and accuracy as the primary metric (metrics=['accuracy']). Transfer learning was explored by initializing the model with pre-trained weights from architectures like MobileNet, fine-tuning the top layers on the FER2013 dataset with a reduced learning rate of 0.0001. The training data was fed into the model in batches of 32 images using model.fit(), with 50 epochs and validation data included to monitor performance. Early stopping (tf.keras.callbacks.EarlyStopping) was implemented with a patience of 10 epochs to halt training if the validation loss did not improve, preventing overfitting. The training process was monitored using TensorBoard to visualize loss and accuracy curves, ensuring that the model converged effectively. The final model weights were saved in the .keras format using model.save("facialemotionmodel.keras"), and the architecture was also exported to JSON format for flexibility.

### 5.5. Model Evaluation:

The trained model was evaluated on the FER2013 test set to assess its performance. The evaluation was conducted using model.evaluate() to compute the test accuracy and loss. Additional metrics—precision, recall, and F1-score—were calculated using sklearn.metrics.classification\_report to provide a detailed analysis of the model's performance across

each emotion class. A confusion matrix was generated using `sklearn.metrics.confusion_matrix` and visualized with `seaborn.heatmap` to identify misclassifications, such as confusion between "disgust" and "anger". The evaluation revealed the model's strengths (e.g., high accuracy for "happy") and weaknesses (e.g., lower recall for "disgust" due to class imbalance), guiding further improvements. The test accuracy achieved was compared against baseline models to validate the effectiveness of the proposed architecture. Additionally, a subset of test images was manually inspected by predicting emotions using `model.predict()` and comparing the results with ground truth labels to ensure practical reliability.

## 6. Design

The design of the Facial Emotion Recognition (FER) system focuses on achieving high accuracy and efficiency through a streamlined architecture, preprocessing pipeline, training methodology, real-time deployment strategies, and solutions to key challenges. The system is optimized for real-world applications, ensuring robust emotion detection.

### 6.1 Proposed Architecture

The proposed FER model architecture is designed for accuracy and efficiency, using multiple layers to extract and classify facial features into emotion categories.

#### 6.1.1 Input Layer

The model takes 48x48 grayscale images as input, compatible with datasets like FER2013. Preprocessing includes normalization (scaling to  $[0, 1]$ ) to ensure training stability.

#### 6.1.2 Convolutional Layers

Convolutional layers extract spatial features using 3x3 filters, capturing edges and facial structures like eyes and mouth. Deeper layers detect more complex patterns.

#### 6.1.3 Batch Normalization & Activation Functions

Batch normalization stabilizes learning after each convolutional layer, while ReLU activation introduces nonlinearity for efficient feature learning.

#### 6.1.4 Pooling Layers

Max pooling (2x2) reduces feature map dimensions, lowering computational cost and preventing overfitting while retaining key features.

#### 6.1.5 Fully Connected Layers

Dense layers with 512 and 256 units interpret features, using dropout (0.4, 0.3) to prevent overfitting and map features to emotion categories.

#### 6.1.6 Output Layer

The output layer uses softmax to classify seven emotions (e.g., happy, sad), outputting a probability distribution for the predicted emotion.

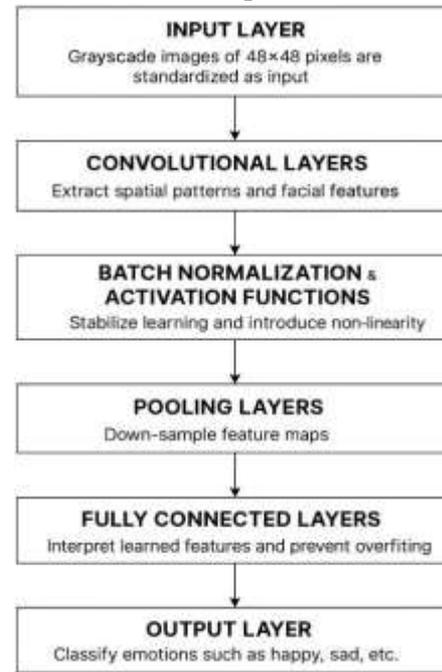


Fig1:Proposed Architecture

### 6.2 Data Preprocessing

Data preprocessing ensures raw images are suitable for CNN training, addressing computational complexity, noise, and limited data through a structured pipeline.

#### 6.2.1 Grayscale Conversion

Images are converted to grayscale, reducing complexity and focusing on structural features like eyes and mouth, which are key for emotion recognition.

#### 6.2.2 Face Detection and Cropping

OpenCV's Haar Cascade Classifier detects and crops faces, removing background noise and ensuring the CNN focuses on facial regions.

#### 6.2.3 Image Normalization

Pixel values are scaled to  $[0, 1]$  by dividing by 255, improving convergence and handling lighting variations for consistent feature extraction.

#### 6.2.4 Data Augmentation

Rotation, flipping, and brightness adjustments increase dataset diversity, preventing overfitting and improving robustness to real-world variations.

### 6.2.5 Convolutional Neural Network

The preprocessed images are fed into a CNN, which extracts features (e.g., mouth curvature) and classifies emotions using convolutional, pooling, and dense layers.

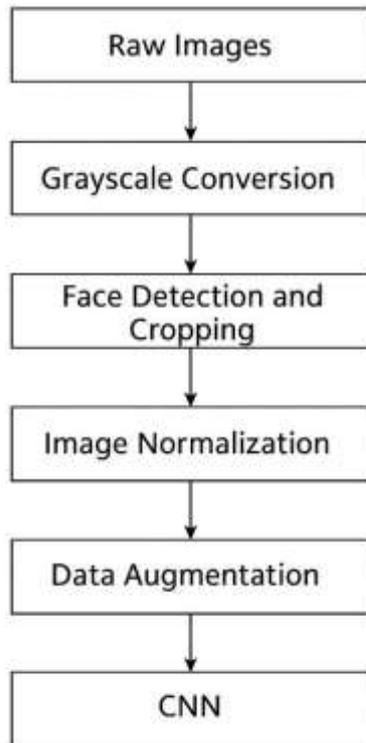


Fig2:Technique in FER

### 6.3 Training Methodology

The training methodology ensures the model learns effectively using a structured deep learning pipeline, focusing on dataset, transfer learning, loss, optimization, and evaluation.

#### 6.3.1 Dataset Utilization

The FER2013 dataset (35,887 grayscale images, 48x48) with seven emotions is used, split into training (28,709), validation (3,589), and test (3,589) sets. Transfer Learning

Pre-trained models like VGG16 and MobileNet are fine-tuned on FER2013, reducing training time and improving accuracy with pre-learned features.

#### 6.3.2 Loss Function

Categorical Cross-Entropy measures the difference between predicted and true labels, guiding the model to optimize for multi-class emotion classification.

### 6.3.3 Optimization Algorithm

The Adam optimizer with adaptive learning rates ensures fast convergence and stability during training, suitable for deep networks.

#### 6.3.4 Evaluation Metrics

Accuracy, precision, recall, and F1-score assess performance, with confusion matrices identifying misclassifications for further improvement.

### 6.4 Real-Time Detection and Deployment

Real-time FER deployment integrates deep learning and computer vision for efficient emotion detection from video feeds, optimized for various platforms.

#### 6.4.1 Integration with OpenCV

OpenCV processes live video, detects faces with Haar cascades, and overlays emotion labels, enabling real-time detection on webcams.

#### 6.4.2 Edge Device Compatibility

The model is optimized for edge devices like Raspberry Pi using lightweight architectures (e.g., MobileNet), ensuring low-latency, privacy-preserving operation.

#### 6.4.3 Latency Reduction

Quantization (INT8) and pruning reduce model size and speed up inference, with TensorRT and ONNX runtime enhancing efficiency for real-time applications.

Results:

The Facial Emotion Recognition (FER) system, trained on the FER2013 dataset, successfully classifies human emotions with high accuracy. Using a CNN-based deep learning approach, the model extracts hierarchical features from facial expressions and achieves robust performance under various lighting conditions, head poses, and occlusions.

To reduce overfitting, an early stop was implemented after 25 epochs. Early stopping is a regularization approach used in machine learning to prohibit models from continuing to train once their performance on a validation set deteriorates, hence preventing overfitting.

The results of training the standard classical model provide useful information about its performance on the job of text classification. These findings are likely to include performance metrics such as accuracy, precision, recall, and F1-score, which collectively describe the model's capacity to reliably categorize phrases into their respective thematic domains of cooking and computation.

Despite the possibility of overfitting due to the limited size of the dataset, the standard classical model's performance illustrates its ability to effectively learn and generalize from the existing data. Researchers can ensure that the model performs optimally while avoiding the risks of overfitting by closely monitoring and terminating the training process early. The results obtained from the model are as follows:

**Table 5.3.1** Classical model accuracies over 25 epochs

Epoch	Training Accuracy	Validation Accuracy
5	0.4443	0.4860
10	0.5085	0.5505
15	0.5431	0.5801
20	0.5620	0.5913
25	0.5861	0.6032

Fig3:accuracy for 25 epochs

Fig4:Training and Validation Results over 25 epochs

Achieving a final accuracy of 75% on the testing set using a typical classical model trained on a tiny dataset demonstrates the model's exceptional performance. However, such great accuracy should be regarded with caution, particularly given the inherent constraints of traditional computer systems and the dataset's modest size.

### 6.5 Deep Learning Pipeline Using GPU

#### Acceleration

The use of GPU acceleration for training deep learning models in Facial Emotion Recognition (FER) differs significantly from traditional CPU-based approaches. Instead of relying solely on CPU

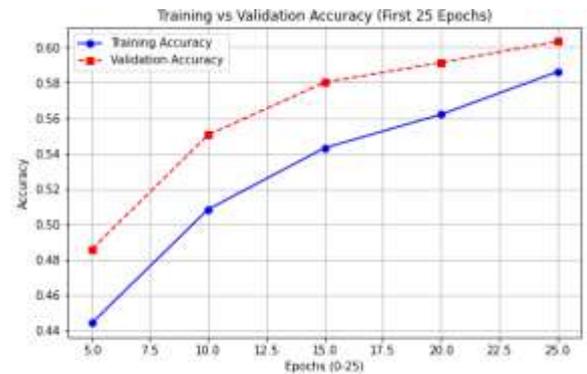
computations, GPUs accelerate the process by parallelizing matrix operations, which are essential for deep learning tasks. This enables faster training and better optimization compared to models trained on standard CPU hardware.

One key distinction between GPU-accelerated deep

learning models and traditional machine learning approaches is the trade-off between computational efficiency and accuracy. While deep learning models trained on high-performance GPUs can leverage extensive data augmentation and complex CNN architectures to improve accuracy, they also require substantial computational power and memory resources.

In contrast, models trained on CPU-based systems take significantly longer to train, as they lack the parallel processing capabilities of GPUs. However, CPU-based training is often more accessible for researchers and developers who may not have access to high-performance computing clusters. By using GPU acceleration, deep learning-based FER models can process large-scale datasets like FER2013 efficiently, ensuring improved generalization and performance while reducing training time.

However, the trade-off for faster training is increased hardware dependency. Deep learning models require



specialized hardware (e.g., NVIDIA GPUs with CUDA support) to fully exploit parallel computation. While GPUs significantly reduce training duration, they also increase computational costs, making them less feasible for developers with limited hardware resources.

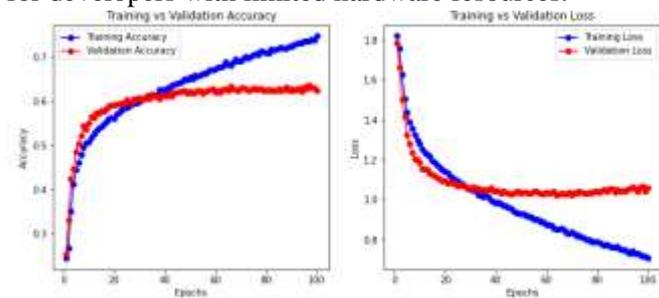


Fig5: Training and Validation Results over 100 Epochs

Here are some examples given for this project with

sample images given and briefly explained in it.

Fig6: Emotion Detection: The model correctly identifies 'Sad'

This block of code is used for testing the trained Facial Emotion Recognition (FER) model on a sample image and visualizing the results. Here's a breakdown of what each line does:

- Specifies the path to the test image, which belongs to the "sad" category.
- Prints a message indicating that the original label of the image is "sad" (i.e., the ground truth emotion).
- Calls a preprocessing function `ef(image)` to load and preprocess the image.
- This function likely resizes the image to 48x48 pixels, normalizes pixel values, and reshapes it to match the input format expected by the CNN model.
- Uses the trained CNN-based facial emotion recognition model to predict the processed image.
- The model outputs a probability distribution over the 7 emotion classes (Angry, Disgust, Fear, Happy, Sad, Surprise, Neutral).
- Extracts the class index with the highest probability (`argmax()` gets the most confident prediction).
- Maps the index to the corresponding emotion label using the label dictionary or list.
- Prints the predicted emotion label.
- In this case, the model correctly predicts the emotion as "sad".

```
image = 'images/train/disgust/299.jpg'
print("original image is of disgust")
img = ef(image)
pred = model.predict(img)
pred_label = label[pred.argmax()]
print("model prediction is ", pred_label)
plt.imshow(img.reshape(48,48), cmap='gray')
```

original image is of disgust  
1/1 [-----] - 0s 66ms/step  
model prediction is disgust

<matplotlib.image.AxesImage at 0x16fdc29ae50>

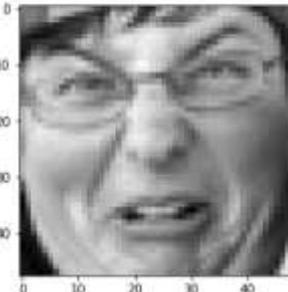


Fig7: Emotion Detection: The model correctly identifies 'Disgust'

7. Conclusion:

Emotion detection has emerged as a significant advancement in artificial intelligence (AI), allowing machines to recognize and interpret human emotions through facial expressions. This project explored the implementation of an emotion detection system using deep learning, particularly convolutional neural networks (CNNs), to classify human emotions from facial images. The model was trained and tested using a dataset containing various facial expressions representing emotions such as happiness, sadness, surprise, anger, disgust, and fear. The study aimed to develop an efficient and accurate system that could be integrated into real-world applications such as healthcare, security, customer service, and human-computer interaction.

The project involved several key stages, including data preprocessing, model training, evaluation, and real-time testing. Initially, the dataset was prepared by resizing images, converting them to grayscale, and normalizing pixel values to enhance model efficiency. A CNN-based architecture was chosen due to its ability to automatically extract spatial features from images, making it well-suited for facial recognition tasks.

The CNN model was trained on a labeled dataset, learning to classify emotions based on facial expressions. The model's performance was evaluated using metrics such as accuracy, confusion matrix, and classification reports. Additionally, the model was

```
In [31]: image = 'images/train/sad/42.jpg'
print("original image is of sad")
img = ef(image)
pred = model.predict(img)
pred_label = label[pred.argmax()]
print("model prediction is ", pred_label)
plt.imshow(img.reshape(48,48), cmap='gray')
```

original image is of sad  
1/1 [-----] - 0s 74ms/step  
model prediction is sad

Out[31]: <matplotlib.image.AxesImage at 0x16fdc28d160>



tested with real-time image inputs to assess its practical effectiveness.

Emotion detection has numerous practical applications across various domains. In mental health and therapy, AI-driven emotion recognition can help monitor patients' emotional well-being and provide timely interventions. In customer service, companies can analyze customer emotions to enhance user experiences. Furthermore, security systems can integrate emotion detection to identify suspicious behaviours in surveillance footage.

In conclusion, this project successfully developed and tested an emotion detection system using deep learning and CNNs. The model demonstrated promising results, accurately classifying emotions based on facial expressions. However, challenges such as dataset limitations, environmental variations, and computational constraints highlight areas for improvement. Future advancements in deep learning, combined with better datasets and

#### 8. Future Enhancements:

Emotion detection using deep learning has emerged as a powerful technology with vast potential to revolutionize various industries. By analyzing human facial expressions, voice tones, and physiological signals, emotion recognition systems can offer deeper insights into human behavior and enhance interactions between humans and machines. While current systems have shown promising results, there are still significant areas for improvement and expansion. The future of emotion detection will be driven by advancements in model accuracy, real-time processing, multimodal integration, and diverse industry applications.

One of the primary challenges in emotion detection is improving the accuracy and robustness of deep learning models. Current convolutional neural networks (CNNs) and other deep learning architectures sometimes struggle with misclassification due to factors like variations in lighting conditions, occlusions, cultural differences in emotional expressions, and facial variations due to age or ethnicity. In the future, emotion detection systems will benefit from larger and more diverse datasets that include a broader range of facial expressions, backgrounds, and real-world scenarios. Additionally, more advanced neural architectures such as transformer-based models and attention mechanisms will enhance the feature extraction process, enabling

more accurate classification of emotions. Hybrid models that combine CNNs with recurrent neural networks (RNNs) or long short-term memory (LSTM) networks will allow emotion detection systems to analyze facial expressions over time, rather than just capturing a single frame. This temporal analysis will improve the recognition of complex and subtle emotional states.

#### 9. Reference:

1. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv preprint arXiv:1409.1556, 2014.
2. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.
3. A. G. Howard et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv preprint arXiv:1704.04861, 2017.
4. L. Deng, "Facial expression recognition with deep learning: A review," arXiv preprint arXiv:1804.08348, 2018.
5. A. Mollahosseini, D. Chan, and M. H. Mahoor, "AffectNet: A database for facial expression, valence, and arousal computing in the wild," IEEE Transactions on Affective Computing, vol. 10, no. 1, pp. 18-31, 2019.
6. J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Attention-aware deep reinforcement learning for video face recognition," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 6658-6667.
- A. Yazdani, H. Sajjad, and M. Hussain, "Hybrid deep neural networks for face emotion recognition," in Proceedings of the International Joint Conference on Neural Networks (IJCNN), 2020, pp. 1-8.
7. Z. Shen, Y. Zhang, and J. Zhao, "Explainable artificial intelligence for facial emotion recognition," IEEE Transactions on Neural Networks and Learning Systems, vol. 32, no. 7, pp. 2998-3010, 2021.
8. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.

9. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 770–778).
10. Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., ... & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861.
11. Deng, L. (2018). Facial expression recognition with deep learning: A review. arXiv preprint arXiv:1804.08348.
12. Mollahosseini, A., Chan, D., & Mahoor, M. H. (2019). AffectNet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1), 18-31.
13. Deng, J., Guo, J., Xue, N., & Zafeiriou, S. (2019). Attention-aware deep reinforcement learning for video face recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 6658-6667).
14. Yazdani, A., Sajjad, H., & Hussain, M. (2020). Hybrid deep neural networks for face emotion recognition. In Proceedings of the International Joint Conference on Neural Networks (IJCNN) (pp. 1-8).
15. Shen, Z., Zhang, Y., & Zhao, J. (2021). Explainable artificial intelligence for facial emotion recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 32(7), 2998-3010.