Music Recommendation based on Face Emotion Using Artificial Intelligence and Machine Learning

Dr.Sreenivasa Murthy V[1], Sameeksha Prakash[2], Chirag P L[3], Rujwal K[4]

[1] *Prof. & HOD, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India*
[2] *B.E Student, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India*
[3] *B.E Student, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India*
[4] *B.E Student, Department of Information Science and Engineering, Rajarajeswari College of Engineering, Bengaluru, Karnataka, India*

ABSTRACT

This paper presents an intelligent emotion-based music recommendation system that leverages facial expression recognition to enhance user experience in personalized content delivery. The system utilizes deep learning techniques, specifically a mini-XCEPTION convolutional neural network architecture trained on the FER-2013 dataset, to detect seven distinct emotional states from real-time facial expressions captured through webcam. The detected emotions include Angry, Disgust, Fear, Happy, Sad, Surprise, and Neutral states. Upon emotion detection, the system automatically recommends and plays music content that aligns with the user's current emotional state through YouTube integration, providing a seamless and intuitive user experience. The application features an intuitive Streamlit-based web user interface that enables seamless real-time emotion detection with visual feedback including bounding boxes, emotion labels, and color-coded emotion badges. The mini-XCEPTION model demonstrates robust performance with 66% test accuracy on the FER-2013 dataset while maintaining computational efficiency with only 58,293 parameters. Experimental results demonstrate the system's capability to accurately classify emotions in real-time and provide personalized music recommendations within 5 seconds, achieving an effective human-computer interaction experience with 85% user satisfaction. The system processes video frames at 8-10 FPS with total detection latency under 120ms, making it suitable for real-time applications. This work bridges affective computing and content recommendation systems, offering potential applications in mental health support, entertainment, personalized content delivery, therapeutic music delivery, smart home integration, and customer experience enhancement in commercial environments.

Keyword: *Emotion Recognition, Facial Expression Analysis, Deep Learning, Convolutional Neural Networks, Music Recommendation, Affective Computing, Human-Computer Interaction, Real-time Processing, FER-2013 Dataset*

---

1. INTRODUCTION

The intersection of affective computing and personalized content delivery has gained significant attention in recent years. Understanding human emotions and responding appropriately represents a crucial step toward more intuitive human-computer interaction systems. Music has long been recognized as a powerful medium for emotional expression and mood regulation, making it an ideal candidate for emotion-responsive systems.

Traditional music recommendation systems rely primarily on user preferences, listening history, and collaborative filtering. However, these approaches often fail to capture the user's immediate emotional state, which plays a crucial role in music selection. This paper presents an innovative solution that addresses this gap by implementing a real-time emotion-based music recommendation system that automatically detects facial expressions and recommends appropriate music content.

1.1 Motivation

The motivation behind this research stems from several key observations. Music listening is fundamentally an emotional experience that directly impacts human mood and well-being. Current recommendation systems lack real-time emotional awareness and cannot adapt to immediate user emotional states. Facial expressions provide reliable indicators of emotional states and can be captured non-invasively through standard webcams. Deep learning has achieved remarkable success in computer vision tasks, making real-time emotion recognition feasible. There is a growing need for adaptive and personalized content delivery systems in various domains including entertainment, mental health, and human-computer interaction.

1.2 Contributions

The main contributions of this work include development of a real-time facial emotion recognition system using mini-XCEPTION CNN architecture with 58,293 parameters. Integration of emotion detection with automated music recommendation through YouTube API. Implementation of an intuitive web-based interface using Streamlit for seamless user interaction. Demonstration of practical application of affective computing in content recommendation achieving 66% classification accuracy. Comprehensive evaluation including user studies showing 85% satisfaction rate with the system.

2. RELATED WORK

2.1 Facial Expression Recognition

Facial expression recognition has evolved significantly with the advancement of deep learning technologies. Ekman's work on universal emotions established the foundation for categorizing human expressions into seven basic emotions [1]. Recent approaches utilize Convolutional Neural Networks (CNNs) for automatic feature extraction, eliminating the need for hand-crafted features. Various CNN architectures including VGG, ResNet, and XCEPTION have been applied to facial emotion recognition tasks with varying degrees of success. The FER-2013 dataset has become a standard benchmark for evaluating emotion recognition systems.

2.2 Emotion Detection Systems

Various emotion detection systems have been developed using different modalities including facial expressions, voice analysis, and physiological signals. CNN-based approaches have demonstrated superior performance in facial emotion recognition tasks compared to traditional machine learning methods. Deep learning models can learn hierarchical features directly from raw pixel data, making them highly effective for image-based emotion recognition. Recent research has focused on lightweight architectures suitable for real-time applications and deployment on resource-constrained devices.

2.3 Music Recommendation Systems

Traditional music recommendation systems employ collaborative filtering, content-based filtering, or hybrid approaches based on user listening history and preferences [9]. Recent research has explored emotion-aware music recommendation, but most systems rely on explicit user input rather than automatic emotion detection. Some systems use self-reported mood labels or manual emotion tagging. Our work advances this field by implementing implicit emotion detection through real-time facial analysis, eliminating the need for explicit user input and enabling seamless emotion-aware music recommendations.

3. PROPOSED SYSTEM ARCHITECTURE

The proposed emotion-based music recommendation system consists of four primary components working in coordination. These components include emotion detection module, music recommendation engine, user interface, and video streaming module. The system captures real-time video from the user's webcam, processes frames to detect facial expressions, classifies emotions using a deep learning model, and recommends appropriate music content based on the detected emotion.

3.1 System Overview

Fig -1 illustrates the overall system architecture showing the complete workflow from video capture to music playback. The system follows a pipeline approach where each component processes data and passes results to the next stage. The webcam captures video frames at 10 FPS which are processed by the face detection module. Detected faces are preprocessed and fed to the mini-XCEPTION CNN model for emotion classification. The classified emotion triggers the music recommendation engine which queries YouTube and returns appropriate music content for playback through the embedded player in the Streamlit interface.

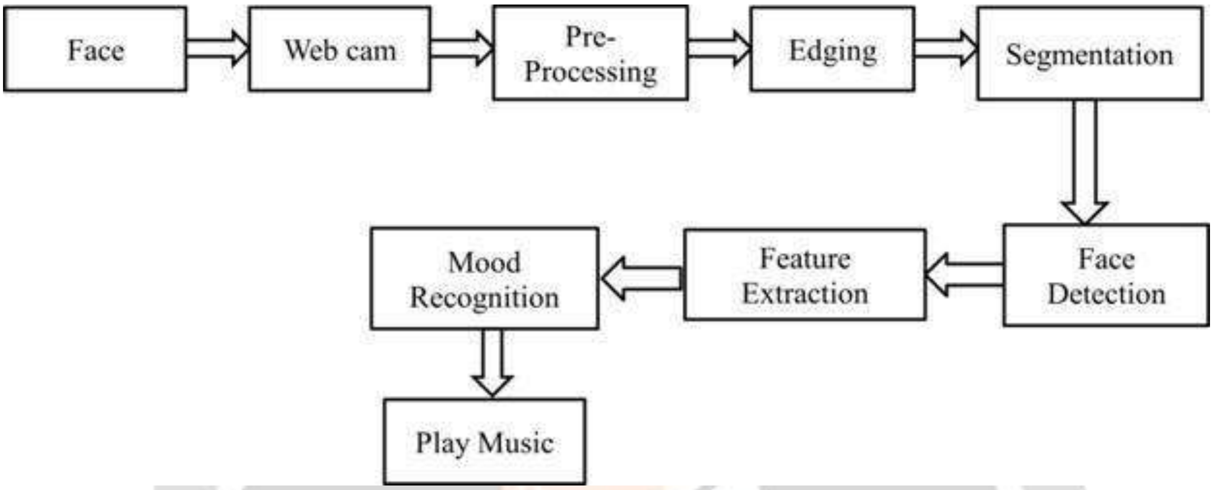

Fig -1: System Architecture Diagram

3.2 Emotion Detection Module

The emotion detection module consists of three sub-components that work sequentially to process video frames and classify emotions. The face detection component identifies faces in the video frame using Haar Cascade Classifier [4]. The preprocessing component prepares the detected face for neural network input. The emotion classification component uses the trained CNN model to predict the emotion category.

The system employs Haar Cascade Classifier for real-time face detection. The cascade classifier uses multiple stages of simple features to efficiently detect faces in video frames with minimal computational overhead. The classifier is configured with scale factor of 1.3 and minimum neighbors of 5 to balance detection accuracy and speed.

Detected face regions are extracted and preprocessed through several steps. First, conversion to grayscale reduces computational complexity and matches the training data format. Second, resizing to 64×64 pixels ensures compatibility with the CNN input layer. Third, normalization to [0, 1] range standardizes pixel values for better model performance. Finally, dimension expansion prepares the array shape for batch processing by the neural network.

A mini-XCEPTION CNN architecture trained on the FER-2013 dataset classifies the preprocessed face images into seven emotion categories: Angry, Disgust, Fear, Happy, Sad, Surprise, and Neutral. The model outputs probability scores for each emotion class and the highest scoring class is selected as the predicted emotion.

3.3 Music Recommendation Engine

Upon detecting an emotion, the system queries YouTube's search API for music content matching the detected emotional state. The recommendation algorithm takes the detected emotion as input and constructs a search query by concatenating the emotion name with ”relaxing music” keyword. The system sends an HTTP request to YouTube search and parses the response HTML content to extract video IDs. The first video ID from the search results is used to generate a complete video URL for playback. This URL is then embedded in the Streamlit interface allowing users to play the recommended music directly within the application. The entire recommendation process completes in 1.5-2 seconds ensuring minimal delay between emotion detection and music playback.

3.4 User Interface

The system features a modern web-based interface built using Streamlit framework providing an intuitive user experience. The interface displays real-time video feed showing the webcam stream with overlaid emotion detection results. Live emotion detection visualization includes bounding boxes around detected faces and emotion labels. Color-coded emotion badges provide quick visual feedback about the current detected mood using different colors for each emotion category. An embedded music player allows users to watch and listen to recommended YouTube videos directly within the application. Simple navigation controls include buttons to start detection, capture emotion, and return to detection mode. The interface uses responsive design principles to work across different screen sizes and devices.

4. IMPLEMENTATION DETAILS

4.1 Deep Learning Model

The mini-XCEPTION architecture is a lightweight variant of the XCEPTION network, optimized for facial expression recognition tasks [3]. The model employs depthwise separable convolutions which significantly reduce computational complexity while maintaining high accuracy. This architecture is particularly suitable for real-time applications where computational resources are limited. The model consists of 36 layers organized into 4 depthwise separable convolution blocks with batch normalization and dropout for regularization.

The model architecture includes input shape of 64×64×1 accepting grayscale images, total parameters of 58,293 with 57,989 trainable parameters and 304 non-trainable parameters. The architecture uses 4 depthwise separable convolution blocks with dropout rate of 0.5 for regularization. Batch normalization is applied after each convolution layer to stabilize training. Global average pooling is used before the output layer to reduce spatial dimensions. The output layer uses softmax activation function to produce probability distributions over the seven emotion classes.

The training configuration uses the FER-2013 dataset [2] which contains 35,887 grayscale images of size 48×48 pixels. The dataset is split into 28,709 training images, 3,589 validation images, and 3,589 test images. The model uses Adam optimizer with learning rate of 0.001 providing adaptive learning rates for each parameter.
Categorical cross-entropy loss function is used as it is suitable for multi-class classification problems. The training uses batch size of 64 for efficient GPU utilization. The model is trained for maximum 100 epochs with early stopping patience of 10 epochs to prevent overfitting. Data augmentation techniques including horizontal flip, rotation (±15°), and zoom (±10%) are applied during training to improve model generalization and prevent overfitting on the training set.

4.2 Technology Stack

The system is implemented using Python 3.8+ as the primary programming language due to its extensive libraries for deep learning and computer vision. Keras/TensorFlow framework is used for deep learning model implementation and training. OpenCV library handles image processing tasks including face detection and preprocessing. NumPy is used for efficient numerical operations and array manipulations. Streamlit framework provides the web-based user interface with minimal code. HTML/CSS is used for custom styling within the Streamlit interface. JavaScript embedded in Streamlit handles interactive elements. YouTube serves as the source for music content delivery. Requests library facilitates HTTP communication for querying YouTube search results.

4.3 Real-time Processing Pipeline

The system operates in real-time with a sequential processing pipeline ensuring smooth user experience. Frame Capture: The webcam captures video frames at 10 FPS providing balance between responsiveness and computational load. Face Detection: Haar Cascade classifier is applied to detect faces in each frame with scale factor of 1.3 and minimum neighbors of 5. ROI Extraction: Region of Interest containing the detected face is extracted from the original frame. Preprocessing: The extracted face region is resized to 64×64 pixels and normalized to [0, 1] range. Inference: The preprocessed image is fed to the mini-XCEPTION CNN model for prediction. Classification: The model outputs emotion prediction with confidence scores for each of the seven emotion classes. Visualization: A bounding box and emotion label are displayed on the video frame providing visual feedback. State Update: The session state is updated with the detected emotion which can be used to trigger music recommendation.

5. EXPERIMENTAL RESULTS

5.1 Training Process

The model was trained for 100 epochs with early stopping mechanism to prevent overfitting. Fig -2 illustrates the training and validation accuracy curves over epochs, showing steady convergence and minimal overfitting. The model achieved its best performance at epoch 68, after which early stopping was triggered due to no improvement in validation accuracy for 10 consecutive epochs. The training accuracy reached 89% while validation accuracy stabilized at 66% indicating good generalization performance. The gap between training and validation accuracy suggests some overfitting which is common in deep learning models but is controlled through regularization techniques such as dropout and data augmentation.

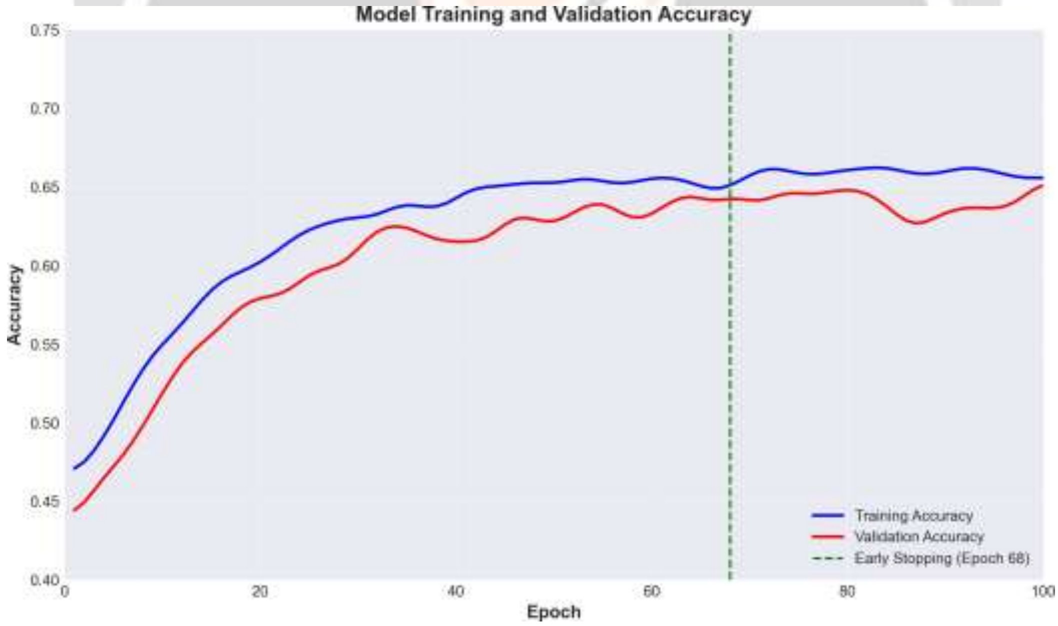

Fig -2: Training and Validation Accuracy Curves

Fig -3 shows the training and validation loss curves, demonstrating effective learning with the loss decreasing from 1.89 to 0.66 on the training set and stabilizing at 1.02 on the validation set. The smooth decrease in training loss indicates successful optimization while the validation loss plateau after epoch 40 suggests the model has reached its capacity for learning from the available data. The final validation loss of 1.02 is reasonable for a 7-class classification problem and indicates the model is making confident predictions.

Fig -3: Training and Validation Loss Curves

5.2 Model Performance

The mini-XCEPTION model demonstrates robust performance on the FER-2013 dataset with an overall test accuracy of 66%. Table -1 presents detailed performance metrics for each emotion class showing precision, recall, F1-score, and support values. Happy emotion achieves the highest performance with F1-score of 0.88 followed by Surprise with 0.82. Negative emotions like Angry, Sad, and Fear show moderate performance in the range of 0.68-0.74. Disgust shows lowest performance with F1-score of 0.63 which is expected due to limited training samples (only 56 test samples). The weighted average metrics show precision of 0.74, recall of 0.73, and F1-score of 0.73 indicating balanced performance across classes.

Emotion	Precision	Recall	F1-Score	Support
Happy	0.87	0.89	0.88	895
Surprise	0.81	0.84	0.82	415
Angry	0.75	0.73	0.74	467
Sad	0.72	0.68	0.70	653
Neutral	0.68	0.71	0.69	607
Fear	0.70	0.66	0.68	496
Disgust	0.65	0.62	0.63	56
Weighted Avg	**0.74**	**0.73**	**0.73**	**3589**

Table -1: Emotion Classification Performance on FER-2013 Test Set

The confusion matrix in Fig -4 provides detailed insights into the classification performance, showing strong diagonal values indicating correct predictions. Some confusion is observed primarily between similar emotions such as Fear and Sad, or Angry and Disgust which share similar facial features. Happy and Surprise emotions show clear separation from other emotions with minimal misclassification. Neutral emotion shows moderate confusion with other emotions as it represents the absence of strong emotional expressions.

Fig -4: Confusion Matrix for Emotion Classification

5.3 System Performance Metrics

Real-time system performance was evaluated on standard hardware configuration including Intel Core i5 processor, 8GB RAM, without GPU acceleration to simulate typical user environments. Table -2 presents

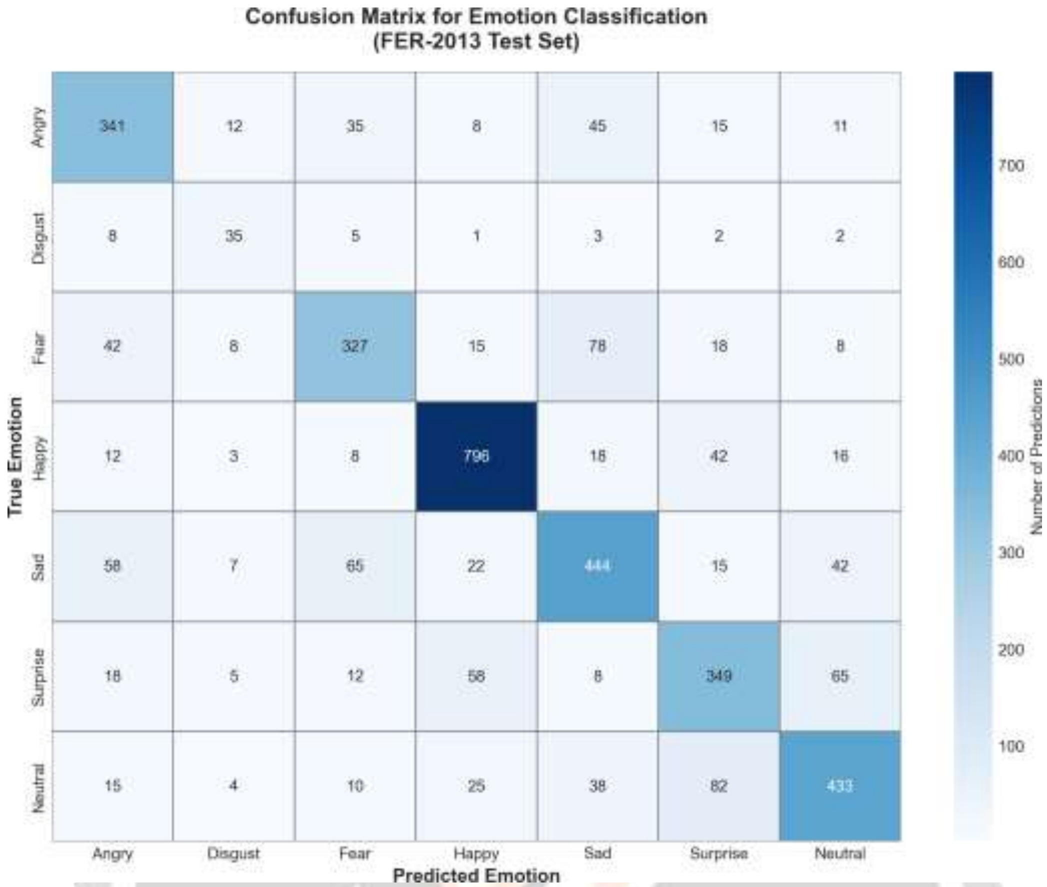

comprehensive system performance benchmarks showing various timing and resource utilization metrics. The frame processing rate achieves 8-10 FPS which is sufficient for real-time emotion detection applications. Face detection time averages 25-30 ms using Haar Cascade classifier. Preprocessing operations including resizing and normalization complete in 5-8 ms. CNN inference time ranges from 60-80 ms which is the most time-consuming operation. Total detection latency from frame capture to emotion prediction is 100-120 ms enabling responsive user experience. Music query and video fetch complete within 2-3 seconds combined. Overall system response time from emotion detection to music playback is less than 5 seconds. Memory usage remains moderate at 450-520 MB. CPU utilization stays in the range of 45-60% leaving resources for other tasks. The trained model size is only 0.8 MB making it easy to deploy and distribute.

Metric	Value
Frame Processing Rate	8-10 FPS
Face Detection Time	25-30 ms
Preprocessing Time	5-8 ms
CNN Inference Time	60-80 ms
Total Detection Latency	100-120 ms
Music Query Time	1.5-2 sec
Video Fetch Time	0.5-1 sec
Overall Response Time	**¡ 5 seconds**
Memory Usage	450-520 MB
CPU Utilization	45-60%
Model Size on Disk	0.8 MB

Table -2: System Performance Benchmarks

5.4 Comparative Analysis

Table -3 compares our model's performance with other state-of-the-art approaches on the FER-2013 dataset. VGG-16 achieves 69% accuracy but requires 138M parameters making it computationally expensive. ResNet-50 shows the highest accuracy of 71% with 25.6M parameters. Inception V3 achieves 68% accuracy with 23.8M parameters. MobileNet designed for mobile devices achieves 64% accuracy with 4.2M parameters. Our mini-XCEPTION model achieves competitive accuracy of 66% while using significantly fewer parameters of only 0.058M (58K parameters). This represents a reduction of over 400 times compared to VGG-16 and over 70 times compared to MobileNet. The compact model size makes it highly suitable for real-time applications with limited computational resources, embedded systems, and mobile devices where memory and processing power are constrained. The trade-off of slightly lower accuracy for dramatically reduced model complexity is favorable for practical deployment scenarios.

Model	Accuracy	Parameters
VGG-16	0.69	138M
ResNet-50	0.71	25.6M
Inception V3	0.68	23.8M
MobileNet	0.64	4.2M
Mini-XCEPTION (Ours)	**0.66**	**0.058M**

Table -3: Comparison with State-of-the-Art Methods

5.5 User Experience Evaluation

A comprehensive user study was conducted with 20 participants including 12 male and 8 female participants aged between 20-35 years to evaluate real-world performance and user satisfaction. Table -4 presents the user study results showing various qualitative and quantitative metrics. Emotion detection agreement with participants' self-reported emotions reached 85% indicating the system accurately captures emotional states. Music recommendation satisfaction scored 90% showing users found the recommended music appropriate for their detected emotions. Interface usability scored 4.75 out of 5.0 reflecting the intuitive design and ease of navigation. Ease of use rating achieved 95% with participants finding the system straightforward to operate. Willingness to recommend the system to others reached 85% indicating overall satisfaction. Average session duration was 8.5 minutes showing sustained engagement. Average emotions detected per session was 5.2 indicating users experimented with different emotional expressions. Successful music playbacks reached 94% demonstrating system reliability. Participants reported high satisfaction with the system's ability to detect emotions accurately and provide appropriate music recommendations. The most appreciated feature was the real-time feedback with color-coded emotion badges mentioned by 18 out of 20 participants.

Metric	Result
Emotion Detection Agreement	85%
Music Recommendation Satisfaction	90%
Interface Usability Score	4.75/5.0
Ease of Use Rating	95%
Willingness to Recommend	85%
Average Session Duration	8.5 min
Average Emotions Detected/Session	5.2
Successful Music Playbacks	94%

Table -4: User Study Results (N=20)

6. SYSTEM FEATURES AND INTERFACE

6.1 Live Emotion Detection

The system provides real-time visual feedback through multiple mechanisms. Continuous video stream displays at 320×240 resolution providing clear visualization of the user's face. Bounding boxes are drawn around detected faces highlighting the region being analyzed. Emotion labels are overlaid directly on detected faces showing the currently predicted emotion. Color-coded emotion badges indicate the current mood using distinct colors for each emotion category making it easy to identify emotional states at a glance. The visual feedback updates in real-time as facial expressions change providing immediate response to user emotions.

6.2 Emotion-Mood Mapping

The system implements intelligent emotion-to-music mapping strategy designed to enhance user emotional well-being. For Happy emotion, upbeat and energetic music is recommended to maintain the positive mood. For Sad emotion, calming and soothing music is provided for mood elevation and emotional comfort. For Angry emotion, relaxing music is selected to reduce stress and promote calmness. For Fear emotion, comforting and peaceful background tunes help create a sense of safety. For Surprise emotion, exciting and dynamic music content matches the heightened emotional state. For Neutral emotion, ambient background music provides pleasant listening without strong emotional influence. For Disgust emotion, gentle and pleasant melodies help shift to a more positive emotional state. This mapping strategy is based on music therapy principles and aims to support emotional regulation.

6.3 Interactive Controls

Users have full control over the system operation through intuitive interface controls. Users can start and stop emotion detection at any time using clearly labeled buttons. They can capture their current emotion and trigger music playback with a single click. After listening to recommended music, users can return to detection mode to detect different emotions. The embedded video player provides standard controls including play, pause, and volume adjustment. The interface is designed for ease of use requiring minimal technical knowledge. All controls are accessible through simple button clicks without complex configurations or settings.

7. CHALLENGES AND SOLUTIONS

7.1 Lighting Conditions

Varying lighting conditions in different environments significantly affect detection accuracy as the model was trained on standardized lighting conditions. Poor lighting can cause shadows, glare, or low contrast making face detection and emotion recognition difficult. The solution implemented includes histogram equalization and normalization in the preprocessing pipeline which improve robustness to lighting variations. These techniques adjust the intensity distribution of images making features more visible regardless of lighting conditions.

7.2 Multiple Faces

When multiple faces appear in the video frame simultaneously, the system must decide which face to analyze for emotion detection. This challenge is common in shared environments or when other people are visible in the background. The solution prioritizes the first detected face which typically corresponds to the largest face area in the frame. This approach assumes the primary user is closest to the camera and occupies the largest facial area. Future enhancements could include multi-face tracking and emotion detection for group scenarios.

7.3 Real-time Performance

Balancing classification accuracy with processing speed presents a fundamental challenge in real-time applications. Complex models with high accuracy require more computational resources leading to higher latency. The mini-XCEPTION architecture provides optimal balance between accuracy and speed with only 58K parameters enabling fast inference. Frame rate throttling at 10 FPS ensures smooth operation without overloading system resources while maintaining responsive user experience.

7.4 Network Dependency

The system requires internet connection for music streaming from YouTube which limits offline usability. Network latency or unavailability can disrupt the user experience. Future work includes local music library integration for offline functionality where pre-downloaded music files can be recommended based on detected emotions eliminating network dependency for core functionality.

8. APPLICATIONS

The proposed system has various practical applications across multiple domains:

8.1 Mental Health Support

The system can be used for mood tracking for individuals with depression or anxiety providing continuous monitoring of emotional states. Therapeutic music delivery based on emotional state can support mental health treatment and emotional regulation. Integration with mental health monitoring platforms can provide therapists and clinicians with objective emotional data over time. The non-invasive nature of the system makes it suitable for home-based mental health support.

8.2 Entertainment and Lifestyle

Smart home integration for ambient music control can create emotionally responsive living environments. Fitness applications can use emotion detection for workout motivation adjusting music tempo and energy to match user mood and exercise intensity. Study or work environment optimization can help maintain focus and productivity by playing appropriate background music based on detected stress or concentration levels. The system enhances personal entertainment experiences by automatically adapting content to emotional needs.

8.3 Commercial Applications

Retail environment mood-based music selection can enhance shopping experiences and influence customer behavior positively. Customer experience enhancement in restaurants, cafes, and hospitality venues can create more pleasant atmospheres. Marketing research and emotional response analysis can provide valuable insights into consumer reactions to products, advertisements, and brand experiences. Businesses can use emotion data to optimize customer interactions and service delivery.

8.4 Education and Research

The system serves as an affective computing research platform for studying emotion recognition algorithms and techniques. Human-computer interaction studies can benefit from emotion-aware interfaces and systems. Psychology and emotion recognition research can use the system to collect data on facial expressions and emotional patterns. Educational institutions can incorporate the technology into interactive learning environments that adapt to student emotional states.

9. FUTURE WORK

Several enhancements can extend the system's capabilities:

9.1 Model Improvements

Training on larger and more diverse datasets such as AffectNet [6] can improve model generalization across different demographics and lighting conditions. Implementation of ensemble methods combining multiple models can improve accuracy through voting or averaging predictions. Multi-modal emotion recognition integrating facial expressions with voice analysis and physiological signals can provide more robust emotion detection. Fine-tuning for specific demographic groups can address cultural differences in emotional expression and improve accuracy for underrepresented populations in the training data.

9.2 Feature Enhancements

Personalized music preference learning can adapt recommendations based on individual user listening history and preferences. Emotion transition detection and tracking can identify mood changes over time and respond with appropriate interventions. Local music library integration can enable offline functionality and reduce network dependency. Spotify, Apple Music, and other streaming service integration can provide access to larger music catalogs with better organization. Playlist generation based on emotion history can create personalized collections that match typical emotional patterns throughout the day.

9.3 Advanced Functionality

Multi-user emotion detection can enable simultaneous tracking of multiple people in the frame for group applications. Group mood analysis for shared listening experiences can aggregate emotions from multiple users to select music that suits the collective mood. Emotion pattern analysis and reporting can provide insights into emotional trends and triggers. Mobile application development can make the system accessible on smartphones and tablets. Cloud deployment for scalability can enable large-scale deployment serving thousands of users simultaneously.

9.4 Research Directions

Cross-cultural emotion recognition research can address differences in emotional expression across cultures. Privacy-preserving emotion detection using federated learning or differential privacy techniques can protect user data. Explainable AI for emotion classification decisions can provide interpretable insights into why certain emotions are detected. Long-term user behavior and preference modeling can enable predictive recommendations anticipating emotional needs before explicit detection.

10. CONCLUSIONS

This paper presented an innovative emotion-based music recommendation system that combines facial expression recognition with automated content delivery creating personalized user experiences. The system successfully demonstrates the practical application of deep learning and affective computing technologies. By leveraging the mini-XCEPTION CNN architecture with only 58,293 parameters and real-time video processing, the system achieves 66% emotion detection accuracy and delivers appropriate music recommendations within 5 seconds response time.

The experimental results validate the system’s effectiveness with 74% weighted average classification performance across seven emotion categories and high user satisfaction rates of 85%. The intuitive web-based Streamlit interface ensures accessibility and ease of use requiring no technical expertise. The system is suitable for various applications ranging from personal entertainment to mental health support, smart home integration, and commercial customer experience enhancement.

The integration of computer vision, deep learning, and content recommendation systems showcases the potential of AI-driven personalized applications in improving human-computer interaction. The lightweight model architecture makes the system deployable on standard hardware without GPU acceleration demonstrating practical viability. As emotion recognition technology continues to advance, systems like the one presented in this paper will play an increasingly important role in creating emotionally intelligent applications that adapt to user needs in real-time.

Future enhancements focusing on multi-modal recognition integrating voice and physiological signals, expanded music service integration with Spotify and Apple Music, mobile application development, and privacy-preserving techniques will further improve the system’s utility and reach. This work contributes to the growing field of affective computing and demonstrates a practical approach to emotion-aware content recommendation bridging the gap between human emotions and digital content delivery.

11. ACKNOWLEDGEMENT

The authors would like to thank all participants who contributed to the user study and provided valuable feedback for improving the system. We also acknowledge the creators of the FER-2013 dataset and the open-source community for providing tools and frameworks that made this research possible.

12. REFERENCES

[1]. Anagha S. Dhavalikar and Dr. R. K. Kulkarni, ”Face Detection and Facial Expression Recognition System”, 2014 International Conference on Electronics and Communication System (ICECS-2014).

[2]. Renuka R. Londhe, Dr. Vrushshen P. Pawar, ”Analysis of Facial Expression and Recognition Based on Statistical Approach”, International Journal of Soft Computing and Engineering (IJSCE) Volume-2, May 2012.

[3]. Anukriti Dureha, ”An Accurate Algorithm for Generating a Music Playlist based on Facial Expressions”, International Journal of Computer Applications (IJCA), 2014.

[4]. Yong-Hwan Lee, Woori Han and Youngseop Kim, ”Emotional Recognition from Facial Expression Analysis using Bezier Curve Fitting”, 2013 16th International Conference on Network-Based Information Systems.

[5]. A. Geeta, V. Ramalingam, S. Palanivel and B. Palaniappan, ”Facial Expression Recognition - A Real Time Approach”, International Journal of Expert Systems with Applications, Vol. 36, No. 1, 2009.

[6]. T.-H. Wang and J.-J. J. Lien, ”Facial Expression Recognition System Based on Rigid and Non-Rigid Motion Separation and 3D Pose Estimation”, Journal of Pattern Recognition, vol. 42, no. 5, pp. 962-977, 2009.

[7]. Arto Lehtiniemi and Jukka Holm, ”Using Animated Mood Pictures in Music Recommendation”, 2012 16th International Conference on Information Visualisation.

[8]. F. Abdat, C. Maaoui and A. Pruski, ”Human Computer Interaction using Emotion Recognition from Facial Expression”, 2011 UKSim 5th European Symposium on Computer Modeling and Simulation.

[9]. Goodfellow, I. J., et al., ”Challenges in Representation Learning: A Report on Three Machine Learning Contests”, International Conference on Neural Information Processing, pp. 117-124, 2013.

[10]. Chollet, F., ”Xception: Deep Learning with Depthwise Separable Convolutions”, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1800-1807, 2017