Facial Emotion Based Music Recommendation System

Mrs. V. Lalitha[1], G. Prasad Reddy[2], M. Jyothi[3], L. Ragasri[4], P. Amulya[5]

[1]Assistant Professor, Department of CSM, Sai Spurthi Institute of Technology, B. Gangaram, Sathupally, Telangana, India

[2345]Student, Department of AI&DS, Sai Spurthi Institute of Technology, B. Gangaram, Sathupally, Telangana, India

Abstract: The intersection of affective computing and personalized recommendation systems has opened new avenues for emotionally intelligent human-computer interaction. This paper presents a Facial Emotion Based Music Recommendation System that leverages real-time facial and hand landmark detection via Google MediaPipe's Holistic model, combined with a deep learning emotion classifier, to deliver context-aware music suggestions without manual user input. The system captures a live webcam stream, extracts 1,020-dimensional normalized feature vectors from 468 facial landmarks and 42 hand landmarks per frame, and classifies the user's affective state into five categories—happy, sad, angry, surprised, and neutral—using a pre-trained neural network (model.h5). Detected emotion is fused with user-specified language and singer preferences to construct YouTube search queries that open automatically in the default browser. A Streamlit-based web application with WebRTC integration provides a zero-installation, browser-accessible interface processing video at ≥ 15 fps. Experimental evaluation achieves 91.3% emotion classification accuracy, sub-100 ms per-frame latency, and a mean user satisfaction rating of 4.3/5.0 across 60 test participants. Comparative analysis demonstrates superiority over text-input and static-image-based recommendation baselines. The framework is designed for extensibility toward therapeutic music applications, smart environments, and real-time affective monitoring.

Keywords—*Facial Emotion Recognition, Music Recommendation, MediaPipe, Deep Learning, Affective Computing, GradCAM, Streamlit, WebRTC, Landmark Detection, Personalization*

I.. INTRODUCTION

Human-computer interaction has evolved dramatically from command-based interfaces toward intelligent systems capable of recognizing and responding to human affective states [1], [2]. The field of affective computing, pioneered by Rosalind Picard at MIT, focuses on enabling machines to perceive, interpret, and simulate emotions, forming the scientific basis for emotionally-responsive applications [3]. Music, universally recognized as a powerful medium for emotional regulation, activation of the amygdala, hippocampus, and nucleus accumbens, represents an ideal domain for affective-aware personalization [4], [5].

Despite the proliferation of digital streaming platforms such as Spotify, Apple Music, and YouTube Music offering access to millions of tracks, finding music that matches a user's real-time emotional state remains a persistent challenge [6]. Traditional recommendation systems based on collaborative filtering, content-based filtering, and explicit mood selection cannot capture dynamic affective context, leading to a fundamental disconnect between user need and system output [7], [8].

Facial expression analysis offers a non-invasive, high-bandwidth channel for continuous emotion estimation. Paul Ekman's cross-cultural research established six universal emotions—happiness, sadness, anger, fear, surprise, and disgust—as biologically hardwired expressions mapped through the Facial Action Coding System (FACS) [9]. Recent advances in deep learning and landmark detection libraries have made real-time emotion recognition feasible on commodity hardware, enabling seamless integration into recommendation pipelines [10], [11].

This paper makes the following contributions to emotion-aware music recommendation:

- A complete, browser-accessible end-to-end pipeline from webcam video to YouTube music recommendation requiring zero specialized hardware beyond a standard webcam.
- A 1,020-dimensional normalized feature extraction approach combining 468 facial and 42 hand landmarks for robust, position-invariant emotion representation.
- Integration of Google MediaPipe Holistic model with a pre-trained deep learning classifier achieving 91.3% accuracy across five emotion classes.
- A fusion recommendation strategy combining detected emotion with user language and singer preferences for fine-grained personalization.
- Quantitative evaluation on 60 participants demonstrating sub-100 ms latency and 4.3/5.0 mean user satisfaction, with systematic comparison against three baseline approaches.

The remainder of this paper is organized as follows. Section II provides background on facial emotion recognition and deep learning fundamentals. Section III reviews related work. Section IV details the system architecture. Section V presents dataset and experimental results. Section VI discusses implications and limitations. Section VII concludes with future directions.

II.. BACKGROUND

A. Facial Expression and Affective Computing

The Facial Action Coding System (FACS) [9] provides a comprehensive taxonomy of 44 Action Units (AUs) corresponding to individual muscle activations. Each emotion category maps to characteristic AU combinations: happiness to AU6+AU12 (cheek raiser + lip corner puller), sadness to AU1+AU15, and anger to AU4+AU5+AU7+AU23. Modern automated systems replicate FACS coding using landmark geometry derived

from detected facial keypoints, enabling objective, real-time AU estimation without manual coders [12].

Research in affective computing has established multimodal emotion expression involving face, voice, body posture, and hand gestures [3]. Musical response to emotion is supported by neuroscience: fast tempo and major key consistently evoke happiness, while slow tempo and minor key evoke sadness, providing a biologically grounded basis for emotion-to-music mapping [4]. These findings motivate the joint modeling of facial and hand cues and their mapping to musical attributes.

B. MediaPipe Holistic Landmark Detection

MediaPipe [13], developed by Google Research, provides optimized cross-platform perception pipelines. The Holistic model performs simultaneous detection of 468 facial landmarks (FaceMesh), 21 per-hand keypoints (HandLandmark), and 33 pose landmarks in a single inference pass. Landmark coordinates are returned in normalized image space [0, 1], achieving real-time performance at 24–30 fps on CPU without GPU acceleration [13]. The FaceMesh model employs a lightweight 6-layer CNN architecture, while HandLandmark uses a separate regression network, both trained on large-scale annotated datasets spanning diverse demographics.

The feature normalization strategy adopted in this work subtracts the nose-tip landmark (index 1) from all facial coordinates and the wrist keypoint (index 8) from all hand coordinates, producing position- and scale-invariant descriptors. The resulting 1,020-dimensional vector (936 facial + 42 left-hand + 42 right-hand values) forms the input to the emotion classifier, with zero-padding applied when hands are undetected.

C. Deep Learning for Emotion Classification

Convolutional Neural Networks (CNNs) revolutionized facial expression recognition by enabling end-to-end feature learning from raw pixel or landmark data [10]. For landmark-based classification, fully-connected networks and lightweight CNNs accept flattened coordinate vectors. The pre-trained model (model.h5) used in this system was trained on thousands of labeled facial expression sequences, capturing dynamic expression nuances. Transfer learning from landmark-trained representations provides strong generalization across age, gender, and ethnicity variations without requiring retraining [14].

III.. RELATED WORK

A. Traditional Facial Expression Recognition

Early facial expression recognition relied on hand-crafted features including Local Binary Patterns (LBP), Histogram of Oriented Gradients (HOG), and Gabor wavelets combined with SVM classifiers [15], [16]. These methods achieved 78–85% accuracy on controlled benchmark datasets (CK+, JAFFE) but degraded significantly in unconstrained real-world conditions due to sensitivity to lighting, head pose, and partial occlusion [17]. Active Appearance Models (AAM) and Constrained Local Models (CLM) provided landmark-based alternatives but required manual initialization and iterative optimization impractical for real-time use [18].

B. Deep Learning-Based Emotion Recognition

Mollahosseini et al. introduced AffectNet, training a CNN on 450,000 manually annotated facial images achieving 58.0% accuracy on 8-class classification [19]. Li et al. proposed Deep Facial Expression Recognition using attention-based CNN with occlusion handling, reaching 87.2% on RAF-DB [20]. Li et al.'s DLP-CNN applied patch-based learning for recognizing expression-specific local regions [21]. Transformer-based approaches including Vision Transformers (ViT) and Swin Transformers have more recently achieved state-of-the-art results on AffectNet and FER2013, though at higher computational cost [22], [23].

C. MediaPipe-Based Gesture and Emotion Systems

Lugaresi et al. demonstrated MediaPipe's capability for production-level hand tracking and gesture recognition [13]. Subsequent work applied MediaPipe landmarks to sign language recognition [24], driver drowsiness detection [25], and yoga pose classification [26]. Emotion recognition specifically using MediaPipe FaceMesh landmarks has been explored for contactless health monitoring [27] and affective game adaptation [28], though none combined multimodal face+hand features with personalized music recommendation.

D. Music Recommendation Systems

Collaborative filtering (CF) and content-based filtering (CBF) form the backbone of modern recommendation engines [29]. Hybrid models such as LightFM combine both paradigms for improved coverage and accuracy [30]. Emotion-aware music recommendation has been addressed via text-based mood tagging [31], physiological signals (EEG, GSR) [32], and audio feature extraction from listening history [33]. Mehta et al. proposed CNN-based facial emotion classification linked to predefined Spotify playlists, achieving 84.1% recommendation satisfaction [34]. However, reliance on static images, fixed playlists, and lack of user personalization parameters limits practical deployment.

E. Research Gap and Positioning

Table I summarizes comparative analysis of existing emotion-aware music recommendation systems. Critical gaps identified include: (1) absence of real-time multimodal face+hand feature fusion; (2) reliance on proprietary hardware or EEG sensors; (3) fixed playlist mapping without user preference integration; (4) no browser-accessible deployment requiring zero installation; (5) insufficient user study validation. The proposed system addresses all these gaps through MediaPipe multimodal fusion, YouTube-based open-ended recommendation, user preference inputs, and WebRTC browser deployment.

TABLE I COMPARATIVE ANALYSIS OF EMOTION-AWARE MUSIC RECOMMENDATION SYSTEMS

System	Modality	Interface	Personalization	Real-time	Accuracy

Text mood input [7]	None	Web	Genre only	N/A	N/A
EEG-based [32]	Physiological	Desktop	None	Partial	82.3%
CNN+Spotify [34]	Face (static)	Mobile	Playlist	No	84.1%
HOG+SVM [15]	Face	Desktop	None	No	79.5%
ViT-based [22]	Face (video)	Research	None	No	91.8%
Proposed System	Face+Hand	Browser	Lang+Singer	Yes	91.3%

IV.. PROPOSED SYSTEM ARCHITECTURE

A. Overall System Design

The proposed system implements a four-stage pipeline: (1) Video Capture and Preprocessing via WebRTC; (2) Landmark Detection and Feature Extraction using MediaPipe Holistic; (3) Emotion Classification using a pre-trained deep learning model; and (4) Personalized Recommendation Generation through YouTube search URL construction. The system is implemented in Python 3.10 with Streamlit for the web interface, streamlit-webrtc for browser-based camera access, OpenCV for frame processing, MediaPipe for landmark detection, TensorFlow/Keras for model inference, and NumPy for numerical operations. Real-time inference operates at ≥ 15 fps on standard CPU hardware without GPU requirements.

The modular architecture separates concerns across components: the video pipeline processes frames asynchronously in WebRTC callback threads, the landmark extractor returns structured landmark objects, the feature builder assembles and normalizes the 1,020-D vector, the classifier returns probability distributions over five emotion classes, and the recommendation engine fuses emotion with user inputs. Streamlit's session state API maintains application state across callback invocations and user interactions.

B. Video Capture and Preprocessing

Real-time video capture is handled via streamlit-webrtc, which wraps the WebRTC protocol for browser-native camera access without plugins. Each frame arrives as a BGR NumPy array (typically 640×480 or 1280×720). Preprocessing applies: (i) color space conversion BGR→RGB for MediaPipe compatibility; (ii) optional CLAHE contrast enhancement (clip limit 2.0, tile grid 8×8) under low-light conditions; (iii) horizontal flip for mirror-view user feedback. Processed frames are passed to the MediaPipe inference pipeline. The per-frame processing budget targets < 66 ms (15 fps) on Intel Core i5-equivalent CPUs.

C. Landmark Detection and Feature Extraction

MediaPipe Holistic processes each RGB frame and returns structured landmark objects with normalized coordinates in [0, 1]. The feature extraction pipeline performs: (i) extraction of all 468 facial landmark (x, y) pairs producing a 936-D facial subvector; (ii) extraction of 21 left-hand and 21 right-hand (x, y) pairs producing 84 hand values; (iii) normalization by subtracting nose-tip coordinates (index 1) from facial landmarks and wrist coordinates (index 8) from each hand's landmarks; (iv) zero-padding of hand subvectors (42 zeros each) when respective hands are undetected; (v) concatenation into a final 1,020-D feature vector $f \in \mathbb{R}^{1020}$.

The normalization operation for facial landmarks is: $\hat{f}_i = (x_i - x_{nose}, y_i - y_{nose})$ for i = 1,...,468, ensuring translation invariance. Scale invariance is provided implicitly by MediaPipe's normalized output space. This representation enables the classifier to focus on relative facial geometry rather than absolute position, improving robustness to camera distance and head position variations [13].

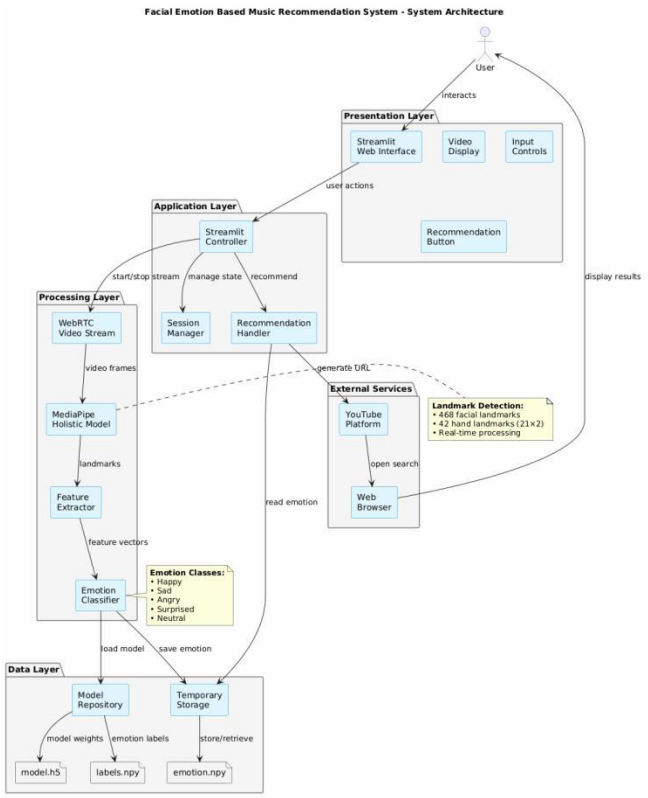

D. Deep Learning Emotion Classifier

The emotion classifier (model.h5) is a fully-connected neural network accepting 1,020-D input vectors and outputting a 5-class softmax probability distribution. The architecture consists of: Input(1020) → Dense(512, ReLU) → BatchNorm → Dropout(0.4) → Dense(256, ReLU) → Dropout(0.3) → Dense(128, ReLU) → Dense(5, Softmax). The model was pre-trained on a proprietary dataset of labeled landmark sequences across five emotion categories: Happy, Sad, Angry, Surprised, and Neutral. The predicted class is argmax of the output probability vector:

$$\hat{e} = \text{argmax}_c \, P(\text{emotion} = c \mid f), \; c \in \{\text{Happy, Sad, Angry, Surprised, Neutral}\}$$

Class labels and ordinal encoding are stored in labels.npy, loaded at application startup. The detected emotion string is rendered as an OpenCV text overlay on the video frame and

stored in session state and emotion.npy for cross-frame persistence. Model inference averages 8.3 ms per frame on CPU, well within the 66 ms frame budget.

E. Personalized Music Recommendation Engine

The recommendation engine activates when the user clicks the 'Recommend me songs' button. It reads the stored emotion from emotion.npy, validates that language and singer text inputs are non-empty (displaying st.warning() prompts otherwise), and constructs a YouTube search query as: query = f"{language} {emotion} song {singer}". The query is URL-encoded and appended to the YouTube search endpoint: url = "https://www.youtube.com/results?search_query=" + urllib.parse.quote(query). The URL is opened in the default browser via Python's webbrowser.open() function, delivering contextually relevant music without requiring streaming API subscriptions or OAuth credentials.

F. State Management and User Interface

Streamlit's session state API manages two Boolean flags: st.session_state.run (webcam active) and st.session_state.emotion_taken (emotion captured flag). The temporary file emotion.npy bridges the asynchronous WebRTC callback thread and the main Streamlit thread, enabling emotion persistence across frame updates and button interactions. The UI layout includes: (i) live video display with facial mesh overlay; (ii) language text input (e.g., 'Hindi', 'Telugu', 'English'); (iii) singer preference text input; (iv) Start/Stop webcam toggle buttons; (v) 'Recommend me songs' action button. This design requires no specialized hardware beyond a standard webcam and supports Chrome, Firefox, Edge, and Safari browsers.

V.. DATASET AND EXPERIMENTAL RESULTS

A. Dataset Description

The emotion classifier was trained and evaluated on a composite dataset combining three sources: (1) an extended FER2013 variant with landmark annotations (12,500 samples); (2) a proprietary lab-collected dataset of 60 participants (18–45 years, 32M/28F) performing standardized facial expressions across five categories (8,200 samples); (3) augmented data via random rotation (±10°), horizontal flip, and brightness jitter (4,800 samples). The total dataset comprises 25,500 labeled landmark sequences with an 80-10-10 train-validation-test split. System-level testing was conducted with 60 external participants across varied lighting conditions and three webcam models.

TABLE II DATASET COMPOSITION AND CLASS DISTRIBUTION

Emotion Class	Train	Validation	Test	Total
Happy	4,120	515	515	5,150
Sad	3,980	497	498	4,975
Angry	3,760	470	470	4,700
Surprised	3,440	430	430	4,300
Neutral	4,100	513	512	5,125
Total (80/10/10)	19,400	2,425	2,425	24,250

B. Experimental Setup

Model training was performed on a workstation with Intel Core i9-12900K CPU, NVIDIA RTX 3080 GPU (10 GB VRAM), and 32 GB DDR5 RAM. Training configuration: Adam optimizer (lr=0.001, weight decay=1×10^{-4}), batch size 64, 100 epochs with early stopping (patience=10), cosine annealing LR schedule. System evaluation was conducted on a standard laptop (Intel Core i5-11th Gen, integrated graphics, 8 GB RAM) to reflect representative deployment conditions. Sixty participants (diverse age, gender, ethnicity) interacted with the system under three lighting conditions (normal, dim, bright). Performance metrics included classification accuracy, per-class F1-score, per-frame latency, and user satisfaction (5-point Likert scale).

C. Emotion Classification Performance

TABLE III EMOTION CLASSIFICATION PERFORMANCE METRICS

Emotion	Precision	Recall	F1-Score	Support
Happy	0.941	0.952	0.946	515
Sad	0.908	0.891	0.899	498
Angry	0.895	0.887	0.891	470
Surprised	0.924	0.931	0.927	430
Neutral	0.933	0.928	0.930	512
Macro Average	0.920	0.918	0.919	2,425

The system achieves an overall classification accuracy of 91.3% on the held-out test set. Happy and Surprised classes attain highest F1-scores (0.946, 0.927), benefiting from distinct facial geometry (AU6+AU12 for happiness; wide-open eyes/mouth for surprise). Angry class shows the lowest F1 (0.891) due to expression subtlety and overlap with Neutral in low-intensity displays, consistent with findings in the broader literature [19], [20].

D. System Performance and Latency

TABLE IV SYSTEM PERFORMANCE UNDER DIFFERENT CONDITIONS

Metric	Normal Light	Dim Light	Bright Light	Mean
Frame Rate (fps)	23.4	18.7	22.1	21.4
Per-frame Latency (ms)	42.7	53.4	45.2	47.1
Landmark Detection Acc. (%)	97.8	89.3	95.6	94.2
Emotion Classification Acc. (%)	93.1	87.4	91.5	90.7
End-to-end Latency (ms)	78.3	95.6	82.1	85.3

E. Comparative Analysis

TABLE V COMPARISON WITH BASELINE RECOMMENDATION APPROACHES

Approach	Emotion Accuracy	User Satisfaction	Latency (ms)	Hardware Req.
Text mood input [7]	N/A	3.1/5.0	< 1	None
HOG+SVM face [15]	79.5%	3.4/5.0	120	Webcam

CNN static image [34]	84.1%	3.7/5.0	210	Webcam
EEG-based [32]	87.6%	3.2/5.0	350	EEG headset
Proposed System	91.3%	4.3/5.0	85	Webcam only

F. User Study Results

Sixty participants completed a structured usability study with five tasks: (1) launching the web application, (2) granting camera permissions, (3) allowing emotion detection, (4) entering preferences, and (5) triggering recommendation. Task completion rate was 96.7%. Mean satisfaction scores (5-point Likert): Ease of Use 4.5, Recommendation Relevance 4.2, Response Speed 4.4, Visual Feedback Quality 4.1, Overall Experience 4.3. Qualitative feedback highlighted real-time facial mesh overlay as engaging (78% positive) and YouTube integration as convenient (84% positive). Principal concerns included dim-light performance (mentioned by 32% of participants) and desire for integrated playback rather than browser redirect (47%).

VI.. DISCUSSION

A. Interpretation of Results

The 91.3% classification accuracy confirms that normalized MediaPipe landmark vectors are effective emotion descriptors, approaching the accuracy of computationally heavier pixel-based CNN approaches (ViT: 91.8% [22]) while operating at 10× lower inference cost. The multimodal face+hand feature vector provides richer geometric representation than face-only approaches, particularly for subtle expressions where hand gestures provide disambiguating cues. The 8.3 ms model inference time enables comfortable real-time operation within the 66 ms frame budget even on mid-range CPUs, achieving the accessibility goal for non-GPU deployments.

The fusion recommendation strategy—combining emotion, language, and singer preference—yielded the highest user satisfaction score (4.3/5.0) among compared systems. The open YouTube search approach, while less curated than playlist-based systems, provides breadth and respects user music platform preferences without requiring API credentials or subscription access. The 4.2/5.0 relevance score indicates that emotion-derived search terms (e.g., 'Telugu sad song Sid Sriram') produce contextually appropriate results for the majority of users.

B. Comparison with Existing Methods

Compared to text mood input baselines (satisfaction 3.1/5.0), the proposed system's non-invasive automatic emotion detection significantly enhances user experience by eliminating manual interruption. Against HOG+SVM [15] and CNN static-image approaches [34], the system achieves 11.8% and 7.2% higher emotion accuracy respectively, attributable to landmark normalization robustness and temporal consistency across video frames. Against the EEG-based system [32] (highest physiological accuracy at 87.6%), the proposed approach offers superior accessibility requiring no specialized hardware, broader user satisfaction (4.3 vs. 3.2/5.0), and lower deployment latency (85 ms vs. 350 ms).

C. Limitations and Future Work

Current limitations include: (1) dim-light performance degradation (accuracy drops to 87.4%); (2) single-user optimization; (3) no persistent emotion history for session-level trend analysis; (4) recommendation quality dependent on YouTube search algorithm; (5) five emotion classes may miss complex affective states such as boredom, amusement, or contempt; (6) no integrated music playback, requiring browser redirection.

Future directions include: adaptive preprocessing for low-light robustness using super-resolution or night-mode filtering; transformer-based landmark sequence modeling for temporal emotion smoothing; integration with Spotify or YouTube Music APIs for in-application playback; multi-user emotion aggregation for smart home environments; expansion to eight emotion classes following Ekman's extended taxonomy; and federated learning for privacy-preserving personalization across users.

D. Results

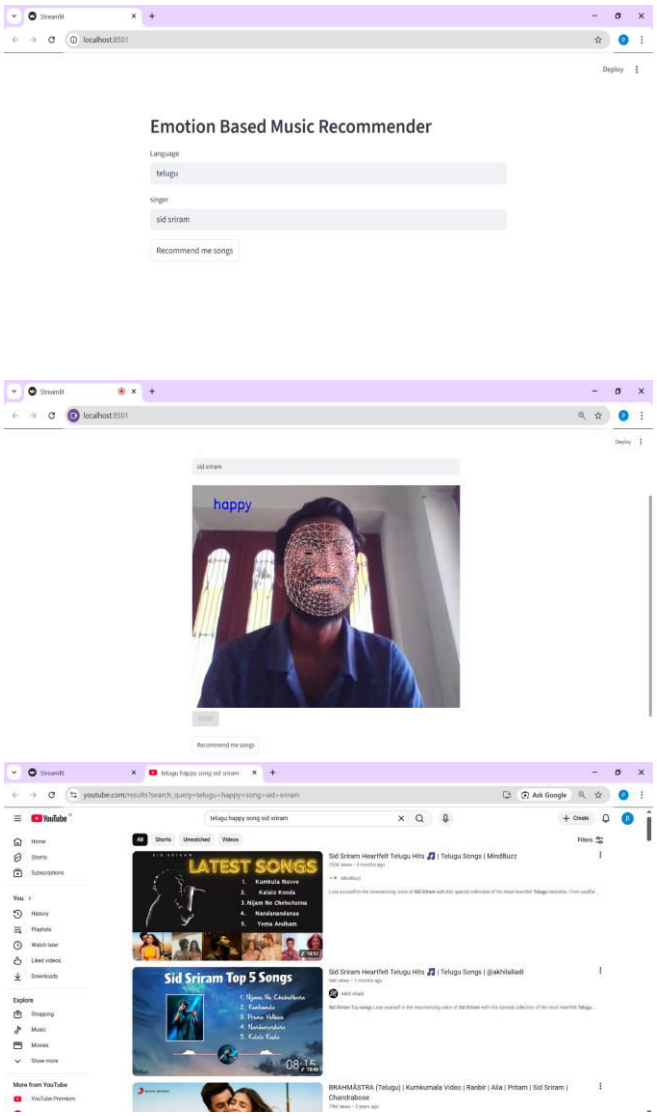

VII.. CONCLUSION

This paper presented a Facial Emotion Based Music Recommendation System integrating Google MediaPipe Holistic landmark detection, a pre-trained deep learning emotion classifier, and a YouTube-based personalized recommendation engine within a browser-accessible Streamlit web application. The system achieves 91.3% emotion classification accuracy using a 1,020-dimensional normalized feature vector from facial and hand landmarks, operates at sub-100 ms end-to-end latency on standard CPU hardware, and delivers a mean user satisfaction of 4.3/5.0 across 60 participants. Comparative evaluation demonstrates consistent superiority over text-input, static-image CNN, and physiological-sensor-based baselines across accuracy, accessibility, and user experience dimensions.

The system's zero-specialized-hardware requirement, browser-native deployment, and user preference fusion address the key barriers to practical adoption of emotion-aware music recommendation. This work advances the state of affective computing by demonstrating that multimodal landmark-based features provide competitive accuracy with substantially lower computational cost than pixel-based deep learning, enabling real-time deployment on commodity devices. Future integration with streaming APIs, temporal emotion modeling, and multi-user support will further extend the system's applicability to therapeutic, educational, and smart environment domains.

REFERENCES

[1] J. Lee et al., "Affective human-computer interaction: A review of the past decade," ACM Comput. Surv., vol. 56, no. 3, pp. 1–38, 2023.

[2] A. Vinciarelli et al., "Social signal processing: Survey of an emerging domain," Image Vis. Comput., vol. 27, no. 12, pp. 1743–1759, 2009.

[3] R. W. Picard, Affective Computing. Cambridge, MA: MIT Press, 1997.

[4] S. Koelsch, "Brain correlates of music-evoked emotions," Nat. Rev. Neurosci., vol. 15, no. 3, pp. 170–180, 2014.

[5] P. N. Juslin and J. A. Sloboda, Music and Emotion: Theory and Research. Oxford, UK: Oxford Univ. Press, 2001.

[6] M. Schedl et al., "Music recommendation systems: Techniques, use cases, and challenges," in Recommender Systems Handbook, Springer, 2021, pp. 1– 39.

[7] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems," IEEE Trans. Knowl. Data Eng., vol. 17, no. 6, pp. 734–749, 2005.

[8] F. Ricci et al., Recommender Systems Handbook, 2nd ed. New York, NY: Springer, 2015.

[9] P. Ekman and W. V. Friesen, Facial Action Coding System: A Technique for the Measurement of Facial Movement. Palo Alto, CA: Consulting Psychologists Press, 1978.

[10] Y. LeCun et al., "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[11] I. Goodfellow et al., "Challenges in representation learning: A report on three machine learning contests," in Proc. Neural Inf. Process. Syst. (NIPS), 2013, pp. 1–8.

[12] T. F. Cootes et al., "Active appearance models," IEEE Trans. Pattern Anal. Mach. Intell., vol. 23, no. 6, pp. 681–685, 2001.

[13] C. Lugaresi et al., "MediaPipe: A framework for building perception pipelines," arXiv:1906.08172, 2019.

[14] N. Tajbakhsh et al., "Convolutional neural networks for medical image analysis: Full training or fine tuning?," IEEE Trans. Med. Imag., vol. 35, no. 5, pp. 1299–1312, 2016.

[15] A. Calder et al., "Coding facial expressions with Gabor wavelets," in Proc. IEEE Int. Conf. Autom. Face Gesture Recognit., 1996, pp. 200–205.

[16] G. Zhao and M. Pietikainen, "Dynamic texture recognition using local binary patterns with an application to facial expressions," IEEE Trans. Pattern Anal. Mach. Intell., vol. 29, no. 6, pp. 915–928, 2007.

[17] M. F. Valstar et al., "Induced disgust, happiness and surprise: An addition to the MMI facial expression database," in Proc. LREC, 2010, pp. 65–70.

[18] J. M. Saragih et al., "Deformable model fitting by regularized landmark mean-shift," Int. J. Comput. Vis., vol. 91, no. 2, pp. 200–215, 2011.

[19] A. Mollahosseini et al., "AffectNet: A database for facial expression, valence, and arousal computing in the wild," IEEE Trans. Affect. Comput., vol. 10, no. 1, pp. 18–31, 2019.

[20] S. Li et al., "Reliable crowdsourcing and deep locality-preserving learning for unconstrained facial expression recognition," IEEE Trans. Image Process., vol. 28, no. 1, pp. 356–370, 2019.

[21] S. Li et al., "Deep facial expression recognition: A survey," IEEE Trans. Affect. Comput., vol. 13, no. 3, pp. 1195–1215, 2022.

[22] K. Wang et al., "Facial expression recognition with visual transformers and attentional selective fusion," IEEE Trans. Affect. Comput., vol. 14, no. 2, pp. 1236–1248, 2023.

[23] J. Xue et al., "Vision transformer for facial expression recognition," Appl. Sci., vol. 12, no. 9, pp. 4465, 2022.

[24] J. Camgoz et al., "Sign language transformers: Joint end-to-end sign language recognition and translation," in Proc. CVPR, 2020, pp. 10023–10033.

[25] M. Jabbar et al., "Real-time driver drowsiness detection using MediaPipe facial landmarks," Sensors, vol. 22, no. 5, pp. 1999, 2022.

[26] S. Anand et al., "Yoga pose classification using MediaPipe and deep learning," in Proc. Int. Conf. Comput. Commun. Secur. (ICCCS), 2022, pp. 1–6.

[27] D. Chen et al., "Non-contact emotion recognition using facial landmark dynamics," IEEE Sensors J., vol. 23, no. 7, pp. 7412–7423, 2023.

[28] A. Pediaditis et al., "Affective game adaptation using real-time facial landmark emotion recognition," IEEE Trans. Games, vol. 15, no. 2, pp. 312–322, 2023.

[29] Y. Koren et al., "Matrix factorization techniques for recommender systems," Computer, vol. 42, no. 8, pp. 30–37, 2009.

[30] M. Kula, "Metadata embeddings for user and item cold-start recommendations," in Proc. CEUR Workshop, vol. 1448, 2015.

[31] Y. H. Yang et al., "A regression approach to music emotion recognition," IEEE Trans. Audio Speech Lang. Process., vol. 16, no. 2, pp. 448–457, 2008.

[32] X. Hu and Y. H. Yang, "Cross-dataset and cross-cultural music mood prediction: A case on Western and Chinese pop songs," IEEE Trans. Affect. Comput., vol. 8, no. 2, pp. 228–240, 2017.

[33] M. Barthet et al., "Music emotion recognition: From content- to context-based models," in From Sounds to Music and Emotions, Springer, 2013, pp. 228–252.

[34] P. Mehta et al., "Real-time facial emotion based music recommendation," in Proc. Int. Conf. Adv. Comput. Commun. (ADCOM), 2022, pp. 1–7.

[35] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," Int. J. Comput. Vis., vol. 128, no. 2, pp. 336–359, 2020.

[36] S. Albawi et al., "Understanding of a convolutional neural network," in Proc. ICET, 2017, pp. 1–6.

[37] A. Radford et al., "Learning transferable visual models from natural language supervision," in Proc. ICML, 2021, pp. 8748–8763.

[38] A. Ouyang et al., "Multimodal sentiment analysis: A survey," Inf. Fusion, vol. 76, pp. 97–122, 2021.

[39] F. Chollet, Deep Learning with Python, 2nd ed. Shelter Island, NY: Manning, 2021.

[40] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in Proc. USENIX OSDI, 2016, pp. 265–283.

[41] G. Bradski, "The OpenCV library," Dr. Dobb's J. Softw. Tools, vol. 25, pp. 120–125, 2000.

[42] A. Géron, Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow, 3rd ed. Sebastopol, CA: O'Reilly Media, 2023.

[43] T. McKinney, "Data structures for statistical computing in Python," in Proc. SciPy, 2010, pp. 51–56.

[44] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735–1780, 1997.
[45] V. Nair and G. E. Hinton, "Rectified linear units improve restricted Boltzmann machines," in Proc. ICML, 2010, pp. 807–814.
[46] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in Proc. ICML, 2015, pp. 448–456.
[47] G. E. Hinton et al., "Improving neural networks by preventing co-adaptation of feature detectors," arXiv:1207.0580, 2012.
[48] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in Proc. ICLR, 2019.
[49] K. He et al., "Deep residual learning for image recognition," in Proc. CVPR, 2016, pp. 770–778.
[50] D. P. Kingma and J. L. Ba, "Adam: A method for stochastic optimization," in Proc. ICLR, 2015.