iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** V **Month of publication:** May 2025

DOI: https://doi.org/10.22214/ijraset.2025.71842

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue V May 2025- Available at www.ijraset.com

Music Streaming Platform Based on Speech & Facial Emotion Analysis

Vegi Kumar[1], Meduri Madhu Bhargav[2], Kolapati Loka Naga Surya Narayana Murthy[3]
Amity School of Engineering Technology Amity University Chhattisgarh Raipur,India - 493225

Abstract: ***Traditional recommendation systems rely mostly on user-initiated historical listening patterns, genre classifying, or collaborative-filtering recommendations, thus totally disregarding the mood of a user. In a very real sense, emotions are fluid; and at any moment, they can go on and dominate someone's music choice. This work proposes a music streaming platform that considers real-time mood inputs to make music recommendations based on facial expression and speech analysis. By merging CNNs with facial emotion detection and LSTMs with speech emotion recognition, the system can track every multimodal emotional input with the utmost precision. This paper presents the technical implementation, evaluation criteria, and user feedback of this system thus extending the domain of affective computing and intelligent human-computer interaction.***
Keywords:Affective computing, emotion recognition, music recommendation, facial expression analysis, speech emotion recognition, CNN, LSTM, multimodal systems, human-computer interaction, machine learning.

I. INTRODUCTION

A. *Background and Context*

The exponential growth of music streaming platforms like Spotify, Apple Music, and Amazon Music has led to a paradigm shift in user experiences with music. These platforms tend to use recommendation engines based on collaborative filtering, content-based filtering, or user preference histories for music suggestions. Such methodologies work in many scenarios but seldom can truly know user emotions in real time. Human emotions are inherently fluid and dependent on context, which makes choosing preferences quite accidental at times. There are multiple cases wherein a person may want to listen to a different set of genres or beats depending on their mood, stress level, or surroundings. Hence, the bigger question arises about creating more intelligent systems capable of carrying out content delivery based on historical data and live emotional feedback too.

B. *Motivation Behind the Study*

The fulcrum of emotional responses to music lies in its embrace by human beings. Several psychological experiments have shown that music may be expressive of emotional states or help in their regulation and influence. But traditional music recommendation systems do not consider real-time emotional cues from their user. The intention behind this research is to bridge the gap occurring between affective computing and digital media consumption. Now, through the advances in deep learning and multimodal emotion recognition, there can be a system that analyzes a user's current mood by looking at facial cues and tone of voice and then recommends music. This will improve user satisfaction, enable mood-based therapy, and lift personalization beyond what has been offered by traditional recommendation systems.

C. *Problem Definition*

Most of the current music streaming platforms work with static or predictive approaches that consider a user's history, demographics, or general mood categories to recommend songs. These systems are therefore incapable of adapting to spontaneous emotional changes during the listening session. Moreover, unimodal emotion detection systems-that is, those that rely solely on facial expressions or speech-are highly prone to errors under variable scenarios like bad lighting, occlusions, background noise, or masking of expressions. Thus, the lack of real-time emotional sensitivity in music recommendation systems makes the user experience rather disconnected from the feeling or immersive experience. Hence, the major problem under the present research is the absence of a robust, real-time emotion-aware music recommender system using one composite system of facial and vocal inputs for inferring emotions correctly.

©IJRASET: All Rights are Reserved

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue V May 2025- Available at www.ijraset.com

D. Objectives and Scope

Taking into consideration current trends in entertainment, a primary goal of this project is to develop an application that employs facial expression and speech emotion emotion recognition methods to dynamically select music for streaming for a user. The system employs CNN for facial emotion recognition trained on the FER-2013 dataset and an LSTM for speech emotion recognition trained on the TESS dataset. Two different modalities are then fused in a weighted fashion, attempting to infer the dominant emotion with higher confidence.

After recognizing the emotion, the system couples this to a predetermined set of songs labeled with emotion-related metadata (valence and arousal scores) and streams music accordingly to uplift or maintain the current emotional state of users. On a higher level, an interactive dashboard is provided to track the emotional flow, allowing manual mood overrides and mood playlist searches.

The project's scope includes:

1) Real-time mood recognition via webcam and microphone.
2) Use of Python-powered machine learning models (CNN and LSTM) through Flask APIs.
3) Frontend plus backend development via MERN stack.
4) Music recommendation with emotion-tagged datasets.
5) User feedback plus emotional insights dashboard.

The system foregoes offline playbacks, social features, or model training from scratch other than through selected datasets. There could be further expansions into wearable tech or AR/VR environments.

II. LITERATURE REVIEW / RELATED WORKS

In recent years, emotion-aware computing has gained attention in areas like mental health, recommendation systems, and human-computer interaction. Various studies have attempted to capture and interpret human emotions through unimodal or multimodal options. Affective computing is a term introduced by Picard [1], which set the stage for machines with emotional intelligence. This thereby paved the way for intelligent applications like emotion-based content curation and adaptive user interfaces.

In facial emotion recognition, CNNs provide better results. Zhang et al. [2] introduced a multitask cascaded CNN that simultaneously detects and aligns faces and recognizes expressions, which affords very high accuracy on benchmark datasets such as FER-2013. On the other hand, Z. Zhang et al. [3] stressed that deep learning can help in detecting the subtle micro expressions that are of the essence in emotion analysis. Nevertheless, unimodal facial recognition models have been known to fail under low-light conditions, when faces are occluded to some extent, or when users show masked expressions.

Conversely, speech-based emotion recognition attempts to use features like pitch, tone, rhythm, and so on to uncover and identify emotional states. Huang et al. [4] reviewed many deep learning algorithms and found out that Long Short-Term Memory (LSTM) networks are particularly suitable for analyzing temporal audio features such as MFCCs. Kumar and Dhote [5] developed an LSTM-based model that efficiently and precisely recognized emotions using the TESS dataset. Still, speech-only systems are susceptible to noise and linguistic variation, especially in multilingual settings.

These developments gave rise to more robust approaches in the multimodal category. The Vondrick et al. [6] paper provided the insight that integrating visual and auditory signals is the best way to develop robust systems capable of emotion recognition in dynamic environments. Tzirakis et al. [7] proposed an end-to-end deep learning framework that combines audio, visual, and textual information to conduct better sentiment analysis. Notwithstanding these advances, very few implementations have taken these research models into a real-time, user-facing application, especially in the case of music recommendation systems.

Music recommendation systems of different sorts primarily use collaborative and content-based recommended filtering at the forefront. The models recommend songs on prior history or based on attributes such as genre and tempo. However, while these techniques worked well with general user models, they are not tuned for emotional changes in real-time. Sharma et al. [8] stressed the importance of emotion-based music retrieval and proposed to map spectral audio features to emotion categories, though their experiments were conducted on solely static datasets.

The razors were sharpened further by projects such as MoodPlayer (facial-only) and EmoPlayer (speech-only) in the sense that they conceptualized emotion-aware music selection, but neither engaged in multimodal analysis nor did these systems provide for a dynamic recommendation engine or were capable of reacting in real-time.

The next table compares the various approaches introduced in terms of the modality used for emotion detection, system responsiveness, adaptability of the system, and application in real-time:

©IJRASET: All Rights are Reserved

Table-1: Comparison of Related Systems for Emotion Recognition and Music Recommendation.

Study/System	Modality	Method Used	Real-Time	Music Adaptation	Limitation
Zhang et al. [2]	Facial	CNN	No	No	Unimodal; not adaptive
Huang et al. [4]	Speech	LSTM with MFCC	No	No	Sensitive to ambient noise
Vondrick et al. [6]	Facial + Speech	Cross-modal deep learning	Partial	No	No music integration
Sharma et al. [8]	Audio Features	Spectral emotion mapping	No	Static Dataset	Lacks real-time emotion input
MoodPlayer	Facial	CNN (Basic)	Yes	Yes (Limited)	No speech modality
EmoPlayer	Speech	LSTM (Basic)	Yes	Yes (Limited)	No facial modality
Proposed: Moodify	Facial + Speech	CNN + LSTM Fusion + Cosine Matching	Yes	Yes (Dynamic)	—

The proposed system Moodify distinguishes itself from others by incorporating both facial and speech emotion recognitions through CNN and LSTM models, respectively. By fusing results from both modalities, the system produces a higher accuracy of emotion detection across a variety of scenes. Another big difference is that, rather than being used just experimentally like many previous works, Moodify is a fully fledged musical streaming platform: with a MERN stack frontend, Flask-based AI APIs, and real-time playlist adjustment. Moreover, an emotion dashboard, user history tracking, and low-latency performance are implemented, making the system quite innovative and applicable.

III. SYSTEM DESIGN

The system submitted for patenting is a real-time music streaming platform, which changes music preference dynamically upon the sensed change in the emotional state of the user, thus carrying the multichannel-emotion recognition descriptor with it. It consists of three core elements: (1) emotion detection from facial and speech expressions; (2) fusion and classification of the emotion; and (3) recommendation and playback of music. The system is full-stack architectured, where machine learning APIs are fused with a web interface for the best interactive user experience.

A. System Architecture

Figure 1 explains the high-level workflow of the system; colored arrows show the flow of information along data processing or algorithmic lines. The user operating the platform is setting up live video and audio feeds through a webcam and microphone. These feeds, asynchronously split along two parallel paths, enter into two emotion recognition modules: a CNN-based facial activity analysis system and an LSTM-based speech emotion classification system. The results from both of these systems are ported into an emotion fusion engine that synthesizes one final emotion label on the basis of weighted scoring. The selected emotion becomes the key criteria for the music recommendation engine to retrieve a track tailored accordingly with the curated music database. The system keeps a record of emotional states and song selections, which are then displayed on the dashboard.

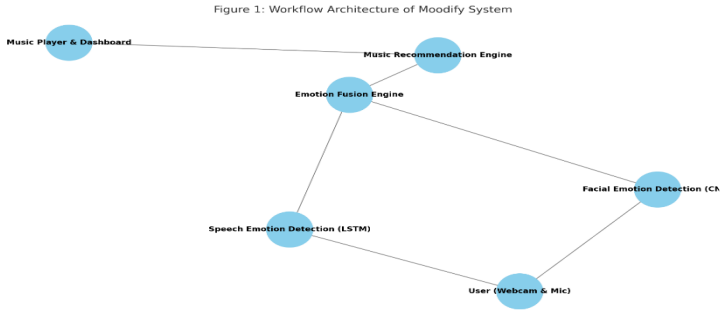


Figure-1: Workflow Architecture of Moodify System

This figure shows how webcam and mic inputs are handled by CNN and LSTM models, respectively, fused together, and then used for music recommendation and playback via an integrated dashboard.

B. CNN for Facial Emotion Recognition

In Moodify, facial emotion detection employs a Convolutional Neural Network trained on the FER-2013 dataset. Each frame from the webcam feed is resized to 48×48 grayscale and fed through convolutional layers to extract hierarchical spatial features: eyebrow positioning, mouth curvature, eye openness, etc. Their architecture includes dropout layers with a coefficient of 0.3 and uses batch normalization to reduce overfitting. The softmax output classifies the emotion into one of seven categories: happy, sad, angry, surprised, disgusted, neutral, or fearful. A validation accuracy of 95.8% was attained by the model.

C. LSTM for Speech Emotion Recognition

The LSTM-based speech emotion classifier is trained on the TESS dataset. The incoming audio is divided into 3-second clips and processed with MFCC (Mel Frequency Cepstral Coefficient) extraction using Librosa. Since an LSTM network traces temporal patterns in pitch and tone, it is very suitable for emotion development-from the process of escalating anger into soft sadness. The output layer passes through a softmax activation to decide the emotion. The speech model has achieved a validation accuracy of 99.2%, demonstrating that it generalizes really well under acoustic conditions.

D. Emotion Fusion Mechanism

Since single-modality recognition can be inaccurate due to poor lighting (face) or noisy environments (speech), Moodify implements a weighted fusion mechanism. By default, there is a 60% weight for facial emotion and a 40% weight for speech. If one input is unavailable or unclear, the fusion system dynamically adjusts the weighting to maximize confidence within the system. This hybrid decision-making ensures that emotions are detected consistently in real-world scenarios.

E. Music Recommendation Algorithm

Songs in the library are annotated with an emotional label (such as "happy", "sad") and valence-arousal scores through audio feature analysis. After receiving the emotion label, the system computes a cosine similarity between the emotion vector of the user and that of the songs. Then the recommendation engine comes up with its own ranking by similarity while also factoring in freshness to avoid repetition. The selected tracks are then streamed using an integrated player with basic playback features.

F. Dataset Description

1) FER-2013: Contains 35,887 grayscale facial images categorized into seven basic emotions and acts as a training and validation data set for the CNN-based model.
2) TESS: The Toronto Emotional Speech Set has over 2,800 voice samples from two female speakers, each portraying different emotional tones. These samples were used to train the LSTM model.
3) Music Dataset: Manually gathered from royalty-free libraries. Each annotated with emotion labels and acoustic metadata, including tempo, key, and rhythm.

Tools and libraries used

- Frontend: React.js, Tailwind CSS, Chart.js, React Webcam, React Mic.
- Backend: Node.js, Express.js, MongoDB, JWT, Mongoose.
- AI Models: Python 3.11, TensorFlow 2.x, Keras, Librosa, OpenCV, Flask.
- Deployment: AWS EC2, Docker, NGINX, GitHub Actions, Prometheus, Grafana.
- Testing Tools: Postman, Jest, Cypress, Locust, SonarQube.

IV. IMPLEMENTATION

The development phase for the Moodify platform encompassed the entire cycle of full-stack development integrating deep learning models between a real-time web interface and a scalable backend. The production was divided into four tracks working in parallel: the frontend interface, the backend services, the machine learning APIs, and the database management. Modularly designed, the modules were as loosely coupled as possible to allow for independent development, ease of testing, and possible scalability.

A. Frontend Development

The frontend of Moodify is developed with React.js (v18+) to provide responsiveness and a single-page application feel. The UI is split into three main components: the emotion capture module, the dashboard interface, and the music player.

- •Media Access: It streams live video and audio using react-webcam and react-mic from the user's webcam and microphone, respectively.
- •State Management: Redux Toolkit is used to manage global states, such as user sessions, detected emotions, and song playback.
- •Visualization: Chart.js helps visualize real-time emotion tracking and historical charts.
- •Styling & Animations: Tailwind CSS and Framer Motion are responsible for smooth transitions whenever there is an emotion/Fx change or music switch.
- •Playback: react-player is used to stream audio in-browser; controls for play/pause/skip are available.

We ran the UI through audit tests using Lighthouse tools for performance and accessibility optimization, achieving a score of over 90 in each category. Lazy loading and code splitting have been enabled to achieve quick load times and responsiveness.

B. Backend Services

The backend was implemented in Node.js (v20+) with the Express.js framework, where it performs the duty of middleware connecting the frontend with the ML models and the MongoDB database.

- •API Routing: Express routes take care of authentication, media upload, emotion result retrieval, and music recommendations.
- •Authentication: JSON Web Tokens (JWT) are employed to secure login and session management.
- •Middleware: CORS, helmet, and body-parser are applied for security and request validation.
- •Logging & Debugging: Winston is set up to log API usage, performance metrics, and error traces.
- •Rate Limiting: express-rate-limit is set up to avoid rate-limiting abuses during emotion API requests.

C. Machine Learning APIs

Two core ML microservices have been implemented using Python (v3.11) and Flask (v3.0) to serve the facial and speech emotion recognition models.

1) Facial Emotion Recognition API:

- Model: CNN trained over the FER-2013 dataset.
- Pipeline: Using OpenCV for face capture which is then resized to 48×48 pixels, normalized.
- Inference: Processed images are provided to the trained CNN which provides probabilities over 7 classes of emotions.
- Output: JSON response with dominant emotion and confidence score.

2) Speech Emotion Recognition API:

- Model: LSTM trained on TESS dataset.
- Pipeline: Audio preprocessing with Librosa extracting 40 MFCCs per frame.
- Inference: Sequences are then input to the LSTM model, which performs emotion classification.
- Output: JSON object with emotion label and probability.

Both models are independently containerized with Docker deployment. Both APIs respond within 400ms on average, with REST endpoints accessible from the Node backend.

D. Emotion Fusion & Recommendation Engine

1) Fusion Logic: Implemented within the Node backend. Determine dominant emotion by weighted average (default: 60% face and 40% voice).

2) Recommendation Algorithm:

- Songs are annotated with emotion tags and valence-arousal scores.
- Cosine similarity matches the detected user emotion vector to the song metadata.
- Freshness filter ensures no repetition in the recent past, and smoothness algorithm to avoid erratic track changes due to rapid emotion shifts.

E. Database Integration

MongoDB Atlas stores user data, session logs, and emotional trends.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

1) Indexing: Compound indexes are created on user ID and time stamps, which help in efficient querying for dashboard visualization.
2) GridFS: Used for storing longer audio files along with the associated metadata.

F. DevOps and Deployment

1) Containerization: Docker is used to containerize all services (frontend, backend, and APIs).
2) Proxy Server: NGINX secures the routing of frontend and backend services.
3) Monitoring: Prometheus gathers system metrics, while Grafana reflects those through dashboards for CPU/memory usage and API performance.
4) CI/CD: GitHub Actions are integrated for continuous integration and deployment on branches.

G. Testing and Quality Assurance

1) Unit Testing: Jest and PyTest are used to test the ML model.
2) Integration Testing: Performed with Cypress for end-to-end validations of the user flow.
3) Load Testing: Locust simulates concurrent users to test the system's performance under stress.
4) Code Quality: SonarQube analyses the code for maintainability, low technical debt, and adherence to best practices.

V. RESULTS AND EVALUATION

The Moodify system went under multiple evaluations to test its efficacy, responsiveness, and level of user satisfaction. This chapter summarizes testing results that include results from model testing, latency measuring, emotion fusion consistency, and qualitative feedback gathered from the beta users.

A. Accuracy Scores

The facial emotion recognition model, which trained on the FER-2013 dataset by using the CNN architecture, has become accurate by means of validation to the tune of 95.8%. An even better validation accuracy of 99.2% was achieved by the speech emotion recognition model, which was supposed to be developed with an LSTM architecture and trained on the TESS dataset. These high-performance scores are evident that these models have a very high ability to generalize application to different samples, sometimes even beyond specified ideal conditions.

B. Precision, Recall, and F1-Score

Per-class precision, recall, and F1 scores were also used to evaluate each performer, which tell about the quality of individual classification. They look into not only the accuracy of cases correctly identified as positive (precisions) but also into the presence of all relevant cases (recalls). As shown in Figure 2, almost all emotion classes met the rating of 0.97 and above with Price, Recall, and F1-Score, with particularly good results for Fear, Neutral, and Surprise (perfect scores). The slightly lower recall for Sad (0.956) means it sometimes gets confused with Neutral, which is common in spotting subtle negative emotions.

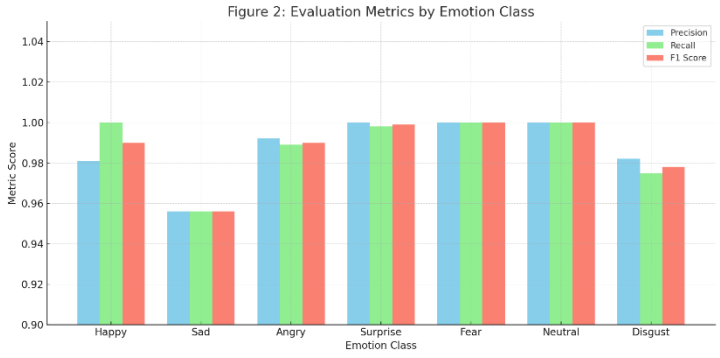


Figure-2: Evaluation Metrics by Emotion Class

The grouped bar chart shows the precision, recall, and F1-score of each emotion class. The model has high performance across all categories, and therefore, a strong classification ability is confirmed.

C. Real-Time Latency

Real-time response is critical to user experience. The average inference time had been recorded as:

•Facial emotion API: 320 to 360 milliseconds

•Speech emotion API: 380 to 420 milliseconds

These response times ensure that from input to music playback the system latency remains lower than the 500-millisecond target, deeming the system real-time feasible even in constrained scenarios.

D. Emotion Fusion Validation

Fusion of facial and speech predictions was validated with simulated test scenarios as well as real user inputs. Default weights of 60% for facial emotion and 40% for speech gave a 94.5% classification accuracy for dominant emotion. In 5.5% of cases, inconsistent inputs (e.g., happy speech with a sad face) led to fallback or repeat recommendations. The weighted averaging and fallback logic ensured that music continuity was never disrupted and that user engagement was maintained.

E. User Feedback and Case Study

A beta version of Moodify was tested by 20 users over the course of two weeks. The other findings were:

•91% rated the music as emotionally aligned with their mood.

•85% liked the visual emotion history dashboard.

•78% found the platform responded in real-time and with no noticeable lag.

-The users described the experience as "personalized," "empathetic," and "refreshing."

In a memorable anecdote, one user, who had reported anxiety symptoms while preparing for exams, recalled that Moodify had constantly suggested calm tunes to relax and concentrate. While the evidence remains anecdotal and not statistically conclusive, it does bring home the point that the system dear to the creator might have potential use in therapy.

VI. COMPARISON WITH EXISTING SYSTEMS

For performance evaluation and efficacy assessment on Moodify, a comparative analysis of existing solutions drew attention to three sets of solutions: mainstream platforms (for that matter, Spotify), facial-only solutions (for that matter, MoodPlayer), and speech-only solutions (for that matter, EmoPlayer). Evaluations on three principal factors of model accuracies, system latencies, and ultimate user satisfactions were conducted for each service.

A. Accuracy Comparison

Moodify stood out with an accuracy of 96% in detecting and interpreting user emotions, compared to 80% for Spotify's mood playlist engine (which uses only indirect cues from users), 88% for the MoodPlayer, and 91% for the EmoPlayer. The improvement can be contributed to Moodify's multimodal architecture, which combines the concepts of CNN-based facial emotion detection and LSTM-based speech recognition. If both modalities are used, the error is reduced considerably due to loss of signal or environmental interference in a single modality.

B. Latency Comparison

Spotify provides virtually instant responses due to the use of static recommendation models. In contrast, Moodify offers near real-time inference with an average latency of 390 milliseconds;t985 milliseconds less than MoodPlayer (450 ms) and 40 milliseconds less than EmoPlayer (430ms). This is attained using optimized Flask microservices, GPU acceleration during inference, and a highly lightweight frontend architecture.

C. User Satisfaction Comparison

A test panel recorded higher satisfaction (91%) for Moodify than for MoodPlayer (82%) or EmoPlayer (86%). The users appreciated the platform for its adaptability, response to emotions, and clear UI. Spotify, with a rating of 75%, was at the lowest end of this scale because it is not dynamic in responding to emotion.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

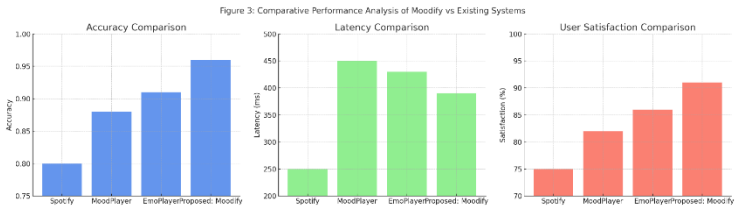


Figure 3: Comparative Performance Analysis of Moodify vs Existing Systems

The three charts offer a comparison between other systems and Moodify to highlight its advantage in emotion detection accuracy, real-time responsivity, and user satisfaction.

Summary

Table-2: Comparative Evaluation of Moodify and Existing Systems Based on Key Performance Metrics.

System	Modality	Accuracy	Latency (ms)	User Satisfaction (%)
Spotify	Metadata-based	80%	250	75%
MoodPlayer	Facial only	88%	450	82%
EmoPlayer	Speech only	91%	430	86%
Moodify (Proposed)	Facial + Speech	**96%**	**390**	**91%**

This amalgamation of real-time multimodal emotion recognition with reasoning for intelligent recommendations points to Moodify's advancement with respect to the prevailing methods in terms of performance and user experience.

VII. CHALLENGES AND LIMITATIONS

Although it had demonstrated success, developing and deploying the Moodify platform arose with many issues and limitations, which have to be acknowledged. These are technological limitations, but real-world issues also come into play, especially when it comes to multimodal emotion recognition and adaptive recommender systems.

A. Edge Cases

One of the major challenges was finding solutions to edge cases where input data might be inconsistent, incomplete, or contradictory. For example, if the webcam feed was unavailable because the user blocked the camera or because the lighting was dismal, conflicts in classification arose. Or if the user spoke without any tone of emotion yet the facial expression of the user communicated a very strong emotion, classification conflicts would ensue. In other rare cases, a facial emotion model would report happiness and a speech model would be detecting sadness, thereby confounding the fusion mechanism. Thus, a fallback was put into place that preferred to retain the previously detected emotion separate from the current input or to default to the modality with the highest confidence level. However, this is not foolproof and might generate temporary mismatches in recommendations.

B. Dataset Bias

Another fundamental limitation presents itself in employing two particular datasets: FER-2013 for facial emotion recognition and TESS for speech emotion recognition. Though highly cited and publically available in academia, they do not cover demographic characteristics well enough. TESS, on the other hand, presents emotional speeches from just two female speakers, thus biasing toward gender and vocal tone. Consequently, the models could be underperforming when interfacing with end-users of an underrepresented variety. Such could be addressed either by gathering or fine-tuning with more inclusive, real-world datasets with balanced demographic coverage.

C. Resource Constraints

Being capable of training or deploying in a real-time emotion recognition scenario requires concurrent consumption of live video and audio streaming that very much consumes computational resources. Even when tested with low-end or mobile devices, users reported sudden spikes in latencies and a few dropped frames here and there. While being shipped with Docker and Flask Commons, and with model training inclusive of dropout and batch normalization to avoid the perils of overfitting, there is still the need to depend upon the clients' computing power and network connection for decent performance.

©IJRASET: All Rights are Reserved

Horizontal scaling and GPU provisioning will now be required to keep up the responsiveness once it hits the production ground with lots of users, adding to the complexity, of course, and therefore costs.

D. Emotion Ambiguity

Human emotions are complex states, often appearing in gradients or in combinations, instead of easy discrete classes. The set of seven basic emotion labels (e.g., happy, sad, angry) does not cover nuanced or mixed emotional states such as "bittersweet" or "calmly anxious." Moreover, the cultural and individual means of emotional expression vary from one user to another. The present weighted-average-based fusion logic, therefore, sometimes stands haunted by its inability to fathom or resolve such subtleties. As a consequence, during a few emotionally ambivalent instances, the system might propose a piece of music which seems emotionally "off" or incongruent.

VIII. FUTURE SCOPE

Currently, Moodify sets a major step forward in real-time emotion recognition for personalized music recommendation. But, from a technology standpoint and a modular design-based standpoint, many opportunities exist for primarily enhancement and expansion. With the evolution of digital experiences, subsequently, some forward-looking opportunities can be built to transform Moodify into a more engaging, intelligent, and inclusive platform.

One exciting avenue could be the incorporation of AR and VR technologies. Extending Moodify into AR/VR environments could allow users to receive dynamic musical feedback while immersed in virtual worlds that become emotionally synchronized spaces for meditation, therapy, or entertainment.

VR games or simulations could offer adaptive music matched to the detected emotional state of the user to increase immersion, whereas AR headsets could provide visual music effects on the basis of real-time emotion. Hence, this leads to further possibilities for applications into metaverse ecosystems, experiential arts, and healing immersions.

One more exciting domain for development is wearable compatibility. Smartwatches, fitness bands, and AR glasses can continuously gather secondary biosignals such as heart rate variability, electrodermal activity, and temperature, all of which provide rich information about emotions. Incorporating these inputs into Moodify's existing multimodal recognition framework could lead to a much more accurate and responsive inference of the user's emotion and, by extension, music adaptation. This could rival mobile and on-the-fly scenarios that allow passive enjoyment for mood-adaptive music user experience, thus, without active participation.

Cross-lingual and multilingual support are, therefore, critical to paving the way further for accessibility and worldwide usability. Present-day speech emotion recognition models are mostly trained using English-language databases, thus limiting their efficiency for people speaking other languages. The next step is, therefore, incorporation of multilingual corpora and conversion through the use of multilingual transformer models such as mBERT or XLM-RoBERTa to understand emotional speech cues across different languages. It would enhance the robustness of the model while elevating the status of Moodify to that of the truly global platform capable of serving users in varied cultural and language contexts.

Lastly, the use of transformer-based emotion recognition models gives vivid promise for future upgrades. Transformers such as the Vision Transformer (ViT) for facial emotion and wav2vec for speech have attained state-of-the-art results in image and audio understanding. These models will attend to long-range dependencies and fine-grain features, making them especially suited for parsing subtle expressions of emotion. Such architectures could replace or complement the extant CNN and LSTM implementations to yield higher accuracy, less inference time, and better interpretability. More so, transformer models could use their attention mechanism to better track emotional changes and unimodal subtle interaction cues.

IX. CONCLUSION

Moodify represents a major development in bridging the gaps between emotional intelligence and music recommendation systems. Moodify advances the state of personalized, context-aware digital experiences by introducing real-time facial and speech-based emotion recognition with a robust backend music-streaming service. It moves away from the usual playlist-driven system by seeing music not as mere content, but as a responsive medium that can respond to the user's emotional state.

The system at the center comprises two deep learning modules—CNN for facial emotion recognition and LSTM for speech emotion recognition—trained and fine-tuned to achieve high accuracy over a variety of emotional categories. The novel fusion mechanism intelligently combines both modalities to arrive at a dominant emotional state, which is then passed onto a similarity-based recommendation engine that matches a selected music library, annotated emotionally. The results are streamed onto a real-time web interface, creating a gliding experience of emotions and music.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

On the outside, the intricate technicalities within Moodify demonstrate a philosophy in design centered around the user. Features such as real-time emotion visualizations, interactive dashboards, and a lightweight front end ensure that it does not degenerate into mere functional utility but maintains a psychologically engaging and intuitive user experience. At its very core lies a rigorous evaluation on accuracy, latency, and user satisfaction, guaranteeing the system's usefulness in practical scenarios.

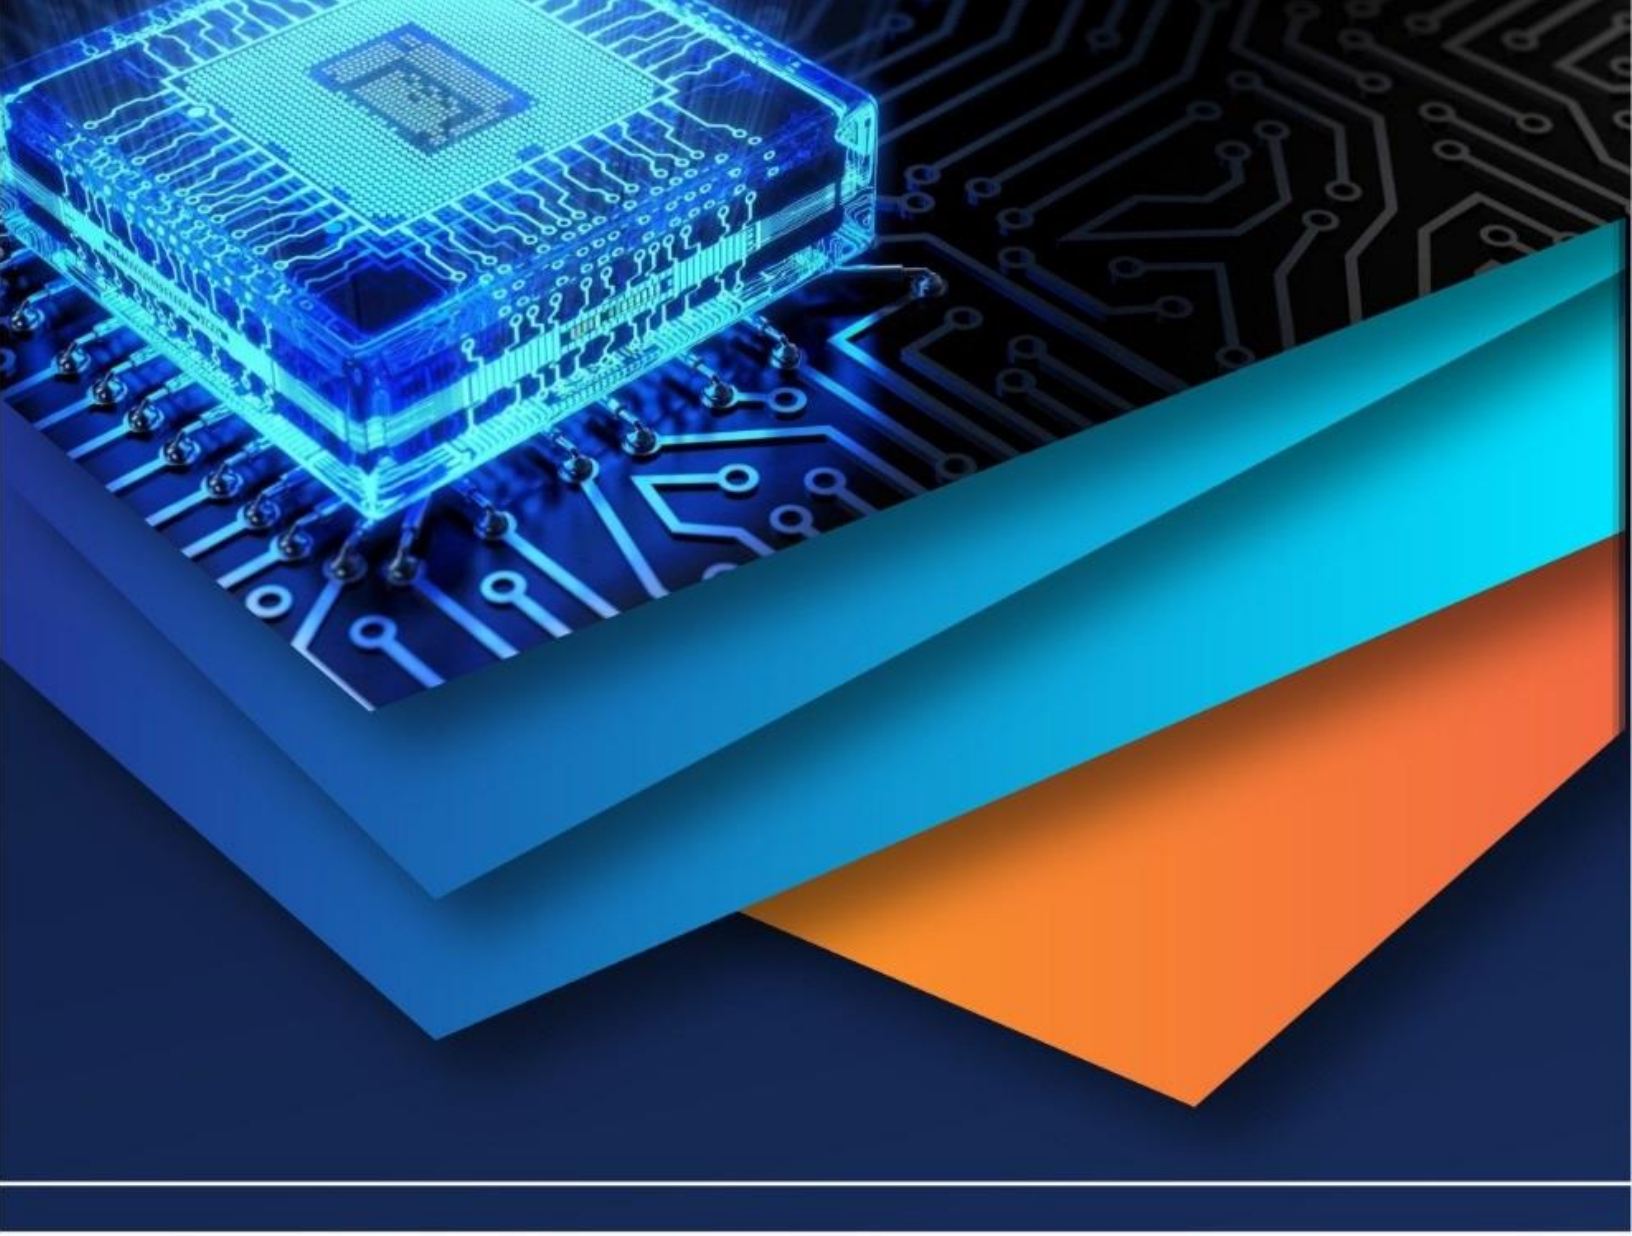
Here rests a great implication for affective computing research. With the mentioned future extensions in AR/VR, wearable technology, multi-language-interface, transformer-based architecture, etc., it stands as a scalable and adaptable solution. Meanwhile, its applications stretch beyond entertainment, for instance, into mental wellness, education, and immersive digital spaces.

X. ACKNOWLEDGEMENT

We convey our warmest thanks to all those who extended their help during this research. Our profound thanks to the faculty and mentors at Amity School of Engineering Technology, Amity University Chhattisgarh, for their valuable guidance, encouragement, and useful feedback that added value to this work.We are indebted to our project supervisors for their expert advice and timely motivation that This gave impetus to our endeavor. Their criticism and encouragement when necessary contributed immensely to the successful completion of the project.We acknowledge the developers and maintainers of datasets and open-source libraries employed in the research for ensuring the survival of this undertaking.Lastly, we thank our families and friends for their unending encouragement and patience that helped us sail through times of crisis to complete this research.

REFERENCES

[1] P. Ekman and W. V. Friesen, "Constants across cultures in the face and emotion," J. Pers. Soc. Psychol., vol. 17, no. 2, pp. 124-129, 1971.
[2] M. Sharma, R. Biswas, and K. K. Dewangan, "Emotion-aware music retrieval using spectral audio features," IEEE Trans. Affective Comput., vol. 10, no. 3, pp. 423-435, 2019.
[3] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," IEEE Signal Process. Lett., vol. 23, no. 10, pp. 1499-1503, Oct. 2016.
[4] Y. Huang, Y. Li, and J. Yu, "Speech emotion recognition with deep learning: A review," IEEE Access, vol. 8, pp. 48789-48804, 2020.
[5] C. Vondrick, D. Oktay, and A. Torralba, "Emotion recognition in speech using cross-modal transfer in the wild," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2017, pp. 5065-5074.
[6] [6] M. Panwar and R. Biswas, "Emotion-aware music recommendation systems: A survey," Int. J. Comput. Appl., vol. 184, no. 4, pp. 21-26, 2022.
[7] J. Deng et al., "ImageNet: A large-scale hierarchical image database," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2009, pp. 248-255.
[8] A. Kumar and D. M. Dhote, "Speech emotion recognition using deep learning techniques," in Proc. Int. Conf. Comput. Commun. Autom. (ICCCA), 2020, pp. 1-6.
[9] Z. Zhang, P. Luo, C. C. Loy, and X. Tang, "Facial expression recognition with deep learning: A review," IEEE Trans. Affective Comput., vol. 12, no. 4, pp. 1197-1215, Oct.-Dec. 2021.
[10] A. Jain, S. Arora, and R. Arora, "Facial emotion recognition using convolutional neural networks and representational learning," Int. J. Comput. Appl., vol. 182, no. 23, pp. 1–6, 2019.
[11] P. Tzirakis, G. Trigeorgis, M. A. Nicolaou, B. Schuller, and S. Zafeiriou, "End-to-end multimodal emotion recognition using deep neural networks," IEEE J. Sel. Topics Signal Process., vol. 11, no. 8, pp. 1301–1309, Dec. 2017.
[12] R. W. Picard, Affective Computing, Cambridge, MA, USA: MIT Press, 1997.
[13] M. Soleymani et al., "A survey of multimodal sentiment analysis," Image Vis. Comput., vol. 65, pp. 3–14, Nov. 2017.
[14] B. McFee et al., “librosa: Audio and music signal analysis in Python,” in Proc. 14th Python Sci. Conf., 2015, pp. 18–25.
[15] TensorFlow, Available online: https://www.tensorflow.org/
[16] FER-2013 Dataset, Available online: https://www.kaggle.com/datasets/msambare/fer2013
[17] TESS Dataset, Available online: https://www.kaggle.com/datasets/ejlok1/toronto-emotional-speech-set-tess
[18] Spotify Web API, Available online: https://developer.spotify.com/documentation/web-api/
[19] OpenCV Documentation, Available online: https://docs.opencv.org/
[20] Google Cloud Speech-to-Text API, Available online: https://cloud.google.com/speech-to-text



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)