Mood Swing Analysis

Ms. Shrutika Suresh More., Shridevi Amol Nandi

Student, Assistant Professor Department of Computer Science & Engineering, Vidya Vikas Pratisthan's Institute of Engineering & Technology, Solapur, India

DOI: https://doi.org/10.51244/IJRSI.2026.1306000474

Received: 04 July 2026; Accepted: 09 July 2026; Published: 17 July 2026

ABSTRACT

Affective computing has emerged as a cornerstone of human-computer interaction (HCI), healthcare analytics, and digital mental health monitoring. Traditional emotion recognition frameworks heavily rely on unimodal architectures—analyzing either text logs, facial expressions, or acoustic patterns in isolation. However, unimodal systems are inherently prone to environmental noise, semantic ambiguities, and cross-channel context blindness, which restrict their real-world reliability. This paper presents a systematic review of contemporary advancements in automated mood swing analysis, focusing on the evolution from handcrafted unimodal classifiers to deep-learning-driven multimodal architectures. We dissect the structural components of feature extraction across linguistic, visual, and acoustic domains, evaluate Early, Late, and Hybrid fusion mechanics, and analyze the deployment bottlenecks in transitioning from complex, resource-intensive models to lightweight web-based frameworks. Finally, we highlight critical gaps in current literature, particularly regarding the handling of cross-modal emotional inconsistencies and real-world framework deployments.

INTRODUCTION

Human emotional states are intrinsically complex, dynamic, and non-linear. Rapid oscillations in emotion, or mood swings, profoundly impact human cognitive decision-making, interpersonal communication, and overall psychiatric well-being. With the recent surges in academic pressure, occupational burnout, and digital lifestyle shifts, computational monitoring of emotional variations has transitioned from a theoretical pursuit to an essential clinical and practical application.

Historically, affect evaluation relied on subjective psychological questionnaires or continuous manual observation, both of which suffer from cognitive bias, temporal lag, and an inability to scale. The maturation of Artificial Intelligence (AI), Deep Learning (DL), and Natural Language Processing (NLP) has opened new frontiers for automated, non-intrusive emotion categorization.

Despite these advancements, a major paradigm constraint remains: Unimodal Dependency. Human beings express affect concurrently through spoken text, facial micro-expressions, and acoustic prosody (tone, pitch, inflection). A system tracking only a single channel remains vulnerable to information deficits. For instance, a person experiencing stress might type grammatically neutral sentences while exhibiting facial or vocal indicators of distress.

To overcome this, Multimodal Affective Computing unifies disparate channels to synthesize a holistic, mathematically robust representation of an individual's true emotional state. This review maps the historical trajectory of emotion recognition architectures, examines cross-domain feature extraction methodologies, analyzes current multi-channel fusion dynamics, and evaluates practical web deployment implementations.

Taxonomy of Domain-Specific Feature Extraction

Automated affect analysis is structurally dependent on transforming high-dimensional, raw data feeds into low-dimensional, numerically dense, and semantically rich feature representations.

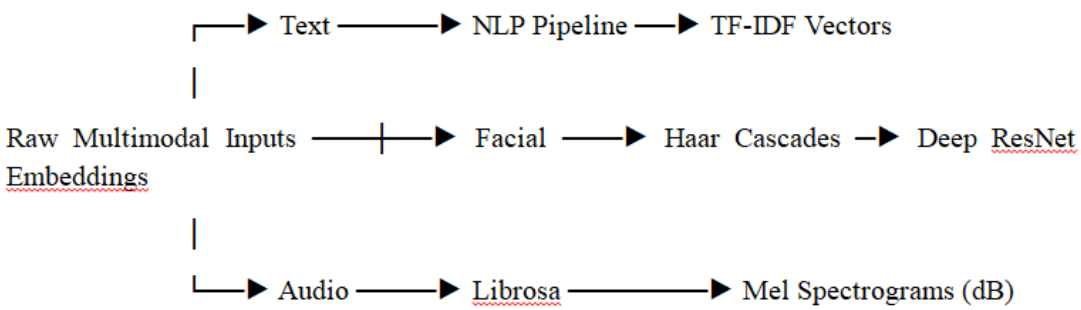

Linguistic Feature Extraction (Text Domain)

Text-based sentiment computation aims to decode explicit and implicit semantic markers within textual data. Early frameworks utilized Lexicon-based matching (e.g., SentiWordNet), which lacked contextual awareness. Modern machine learning pipelines favor vectorized semantic models:

TF-IDF (Term Frequency-Inverse Document Frequency): Evaluates the statistical importance of tokens relative to a corpus, highly effective for sparse, high-dimensional textual logs when paired with statistical classifiers.

Advanced Embeddings: Word2Vec, GloVe, and bidirectional transformer tokens capture dense semantic context, though they require significant memory footprints during deployment.

Visual Feature Extraction (Image/Video Domain)

Facial expression recognition is rooted in spatial pattern matching. The paradigm has shifted from structural geometric constraints to deep spatial representations:

Haar Cascade Classifiers: Deployed in initial preprocessing pipelines for rapid, low-overhead localization of facial regions within bounding boxes.

Deep Residual Architecture (ResNet): By introducing residual mapping arrays and identity shortcut connections, deep architectures solve the vanishing gradient problem, enabling the extraction of multi-layered, abstract visual expressions (e.g., muscle contortions mapped to basic emotional states).

Acoustic Feature Extraction (Audio Domain)

Acoustic emotion classification decodes paralinguistic elements independent of spoken text.

Handcrafted Audio Descriptors: Mel-Frequency Cepstral Coefficients (MFCCs), pitch variance, and zero-crossing rates represent early standard feature sets.

Mel Spectrogram Representation: Transforming raw temporal audio waveforms into a time-frequency log-power representation (mapped to the human-auditory Mel scale) allows audio processing to be treated as a spatial vision task. This makes convolutional topologies highly effective for speech affect analysis.

Multi-Channel Fusion Architectures

The defining challenge of multimodal computing is Fusion—the process through which independent modal feature vectors are unified to output a single prediction.

Feature-Level (Early) Fusion

Early fusion concatenates raw or preprocessed feature vectors from text, visual, and acoustic streams into a singular, high-dimensional joint vector prior to model training. While this preserves cross-channel correlations, it scales poorly due to dimensional mismatching and varying sampling frequencies.

Decision-Level (Late) Fusion

Late fusion processes each data channel independently through specialized modal networks (e.g., Logistic Regression for text, ResNet for images). Each network computes independent class probabilities, which are then unified via aggregation strategies such as:

- Max-Voting
- Confidence-Weighted Averaging
- Meta-Classification

This approach isolates modal errors but may miss subtle cross-channel correlations.

Hybrid Fusion:

Hybrid fusion combines early and late approaches by passing subsets of features through local fusion layers while processing raw signals in isolated, deep channels. This provides a balance of accuracy and flexibility, though it increases model structural complexity.

Critical Research Gaps and Vulnerabilities

A comprehensive review of the literature reveals several systemic gaps in modern affective architectures:

Absence of Tri-Modal Integration Frameworks: Most contemporary architectures explore dual modalities (Text + Video or Audio + Video). Unified platforms that integrate linguistic text vectors, visual residual embeddings, and log-power acoustic spectrograms concurrently remain under-researched.

Handling Cross-Modal Inconsistencies: Current fusion engines are largely optimized under the assumption that all input modalities point to the same emotion. When cross-channel contradictions arise (e.g., text contains high-joy tokens but the vocal input features flat, depressive acoustic frequencies), existing models frequently output corrupted or over-smoothed predictions.

Deployment Bottlenecks: Modern models are heavily weighted toward parameter-heavy transformers requiring constant GPU support. There is a clear need for lightweight, low-latency, and modular web framework configurations (such as Django backends coupled with optimized architectures like ResNet18) capable of processing asynchronous multi-channel uploads without server degradation.

CONCLUSION

Multimodal mood swing analysis marks a major evolutionary step over traditional, error-prone unimodal models. By structuring individual streams into targeted processing blocks—utilizing TF-IDF for linguistic processing, residual CNN structures for spatial visual frames, and Mel scale representations for paralinguistic audio signatures—modern frameworks achieve significantly higher reliability.

Future research must focus on optimizing fusion layer mechanics to better handle conflicting cross-channel context, and on designing lightweight, web-deployable systems that bring academic affective computing into real-world applications.

REFERENCES

1. Ekman, P. (1993). Facial Expression and Emotion. American Psychologist, 48(4), 384-392.
2. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770-778.
3. Tripathi, J., & Beigi, H. (2018). Multi-Modal Emotion Recognition Using Deep Learning Architectures. arXiv preprint arXiv:1808.04926.