2026 Volume: 02 Issue: 03 www.ijrpa.com ISSN 2456-9995 Research Article

International Journal Research Publication Analysis

Page: 01-15

HYBRID EMOTION-AWARE GENERATIVE AI: HYBRID EMOTION-AWARE FRAMEWORK USING MOOD-ADAPTIVE LARGE LANGUAGE MODELS

Prof.R.Lalitha*[1], Mr.Surya Nattarayasamy[2]

[1]*Assistant Professor, Department of Software Systems, Sri Krishna of Arts and Science College, Coimbatore, India.*

[2]Student, Department of Software Systems, Sri Krishna of Arts and Science College, Coimbatore, India.

Article Received: 14 February 2026
Article Revised: 05 March 2026
Published on: 25 March 2026

***Corresponding Author: Prof.R.Lalitha**
Assistant Professor, Department of Software Systems, Sri Krishna of Arts and Science College, Coimbatore, India.

DOI: https://doi-doi.org/101555/ijrpa.6829

ABSTRACT

In recent years, Generative Artificial Intelligence (AI) has significantly transformed human–machine interaction, enabling systems to generate human-like text, images, and speech with remarkable accuracy. Despite these advancements, most Large Language Models (LLMs) continue to lack a critical human attribute—emotional awareness. Although their responses are coherent and contextually relevant, they often appear impersonal and fail to capture the emotional nuances required for meaningful interaction. This limitation becomes particularly significant in domains such as education, healthcare, counseling, and virtual assistance, where empathy and emotional intelligence play a vital role. To address this challenge, this paper proposes **a Hybrid Emotion-Aware Framework using Mood-Adaptive Large Language Models**, designed to integrate emotional intelligence into AI-driven communication systems. The proposed framework combines multiple computational paradigms, including **Affective Computing, Sentiment Analysis, and Adaptive Response Generation**, to enable real-time detection, interpretation, and response to human emotions. By leveraging multimodal cues such as linguistic patterns, vocal tone, and facial expressions, the system dynamically adapts its responses based on the user's emotional state, thereby enhancing contextual relevance and empathetic engagement. The architecture of the proposed framework integrates transformer-based generative models with deep neural network–based emotion recognition modules. An advanced emotion detection layer extracts affective signals from user inputs and encodes them into structured emotion vectors. These vectors are utilized by the generative model to modulate response characteristics, including tone, lexical choice, and sentence structure, resulting in emotionally aligned and context-aware interactions. The hybrid design further incorporates adaptive learning mechanisms to improve personalization and response accuracy over time.

This emotion-aware approach enhances user trust, satisfaction, and overall interaction quality, making AI systems more natural and human-centric. The framework has broad applicability across various domains, including intelligent tutoring systems, mental health support chatbots, and emotionally responsive customer service platforms. Furthermore, the study addresses critical challenges such as emotional bias mitigation, data privacy protection, cultural sensitivity, and the need for robust real-time emotion recognition.

By embedding emotional intelligence into Generative AI, the proposed hybrid framework represents a significant advancement toward more empathetic and ethically responsible AI systems. It envisions a future where machines not only understand human language but also respond effectively to human emotions, fostering deeper and more meaningful human–AI interactions.

KEYWORDS: Generative AI, Emotion-Aware Systems, Affective Computing, Sentiment Analysis, Mood-Adaptive Large Language Models, Human–AI Interaction

1. INTRODUCTION

Over the past decade, **Artificial Intelligence (AI)** has undergone rapid evolution, transitioning from traditional rule-based systems to advanced models capable of understanding and generating human-like language. Modern **Large Language Models (LLMs)**, such as GPT, PaLM, and Gemini, represent a significant breakthrough in machine intelligence, enabling applications including text generation, document summarization, language translation, and conversational interaction. These models demonstrate remarkable fluency and contextual understanding, narrowing the gap between human and machine communication.

Despite their linguistic capabilities, existing LLMs exhibit a critical limitation—**a lack of emotional awareness**. While they effectively process semantic content, they often fail to capture the underlying emotional intent of user inputs. In human communication, emotions play a fundamental role in shaping meaning, tone, and response behavior. Expressions of sadness, frustration, or joy carry implicit emotional signals that guide appropriate reactions.

When AI systems overlook these cues, their responses—though grammatically correct—may appear impersonal, insensitive, or disconnected. This limitation becomes particularly significant in domains such as mental health support, education, therapy, elder care, and customer service, where empathy and emotional understanding are essential for effective interaction.

The absence of emotional intelligence in AI systems highlights a broader challenge in human–machine communication. While machines excel at data processing and logical reasoning, they often lack the ability to engage users on an emotional level. As user expectations evolve, there is an increasing demand for AI systems that are not only intelligent but also **emotionally aware, context-sensitive, and capable of meaningful interaction**. This shift emphasizes the need for integrating emotional intelligence into generative AI frameworks.

To address this gap, this paper proposes **a Hybrid Emotion-Aware Framework using Mood-Adaptive Large Language Models**, designed to enhance the emotional responsiveness of AI systems. The proposed framework integrates key concepts from **affective computing**, **sentiment analysis**, and **adaptive generative modeling** to enable real-time emotion recognition and context-aware response generation. By analyzing multimodal cues—including linguistic patterns, vocal tone, and facial expressions—the system identifies the user's emotional state and dynamically adjusts its responses to reflect empathy, understanding, and contextual appropriateness.

The hybrid architecture combines transformer-based language models with deep neural network–based emotion recognition modules. An emotion processing layer extracts affective features from user inputs and encodes them into structured emotion representations. These representations guide the generative model in modulating response characteristics such as tone, vocabulary, and sentence structure. As a result, the system moves beyond conventional information delivery to provide **emotionally aligned and human-centric interactions**.

The integration of emotion awareness transforms AI systems from passive information providers into **adaptive conversational agents** capable of supporting, motivating, and engaging users based on their emotional context. This advancement enhances user trust, satisfaction, and overall interaction quality, while opening new opportunities in applications such as intelligent tutoring systems, mental health assistants, and emotionally responsive virtual agents.

This study aims to explore how mood-adaptive LLMs can redefine human–AI interaction by improving emotional intelligence, contextual understanding, and conversational authenticity.

Through the proposed hybrid framework, we envision a future in which AI systems not only understand human language but also respond effectively to human emotions, fostering deeper, more natural, and human-centered digital interactions.

2. CHALLENGES IN HYBRID EMOTION-AWARE FRAMEWORKS

While the proposed Hybrid Emotion-Aware Framework using Mood-Adaptive Large Language Models offers significant advancements in human–AI interaction, its implementation involves several critical challenges across technical, ethical, and human-centered dimensions.

3. TECHNICAL CHALLENGES

i. Emotion Detection Accuracy:

Accurate emotion recognition remains a fundamental challenge due to the complex and context-dependent nature of human emotions. Subtle states such as sarcasm, frustration, confusion, and mixed emotions are difficult to interpret reliably from text, speech, or facial expressions. For instance, a simple phrase like "I'm fine" may convey different emotions depending on tone and context. Although deep learning models can identify basic emotions, capturing fine-grained emotional variations remains limited. Variations in accent, speaking style, and language further reduce detection accuracy, particularly in multilingual environments.

ii. Data Imbalance and Cultural Bias:

The performance of emotion-aware systems depends heavily on the quality and diversity of training datasets. However, existing datasets are often limited, imbalanced, and biased toward specific languages or cultural groups. Models trained predominantly on English or Western datasets may fail to generalize across diverse populations. This leads to inaccurate emotion recognition and reduced system reliability. Developing inclusive, culturally diverse datasets is essential for improving model robustness.

iii. Latency and Real-Time Processing:

The hybrid framework requires real-time processing of multimodal inputs, including text, audio, and visual data. Emotion recognition involves multiple computational stages such as feature extraction, sentiment analysis, and context modeling, which increase system complexity. Achieving low-latency responses while maintaining high accuracy is challenging, especially in resource-constrained environments. Efficient model optimization and lightweight architectures are necessary for real-world deployment.

3.1 Ethical and Privacy Challenges

i. Emotional Manipulation:

Emotion-aware systems have the potential to influence user behavior by adapting responses based on detected emotional states. While this can enhance user experience, it also raises ethical concerns regarding manipulation. Systems must be carefully designed to avoid exploiting user emotions for persuasive or commercial purposes, ensuring fairness and user autonomy.

ii. Data Privacy and Security:

The framework processes sensitive emotional data, including voice signals, facial expressions, and textual sentiment. Such data can reveal personal and psychological information about users. Ensuring secure data handling through encryption, anonymization, and strict access control is essential to prevent misuse and maintain user trust.

iii. Transparency and User Awareness:

Users may not always be aware that their emotions are being analyzed by AI systems. Lack of transparency can lead to discomfort and reduced trust. Therefore, it is important to clearly inform users about data usage and provide explainable system behavior, ensuring ethical and responsible AI interaction.

3.2 Human Factors

i. Trust and Acceptance:

For successful adoption, users must perceive the system as empathetic, reliable, and non-intrusive. Excessive monitoring or inaccurate emotional responses may lead to discomfort or rejection. Designing systems that balance emotional intelligence with user privacy is crucial for long-term acceptance.

ii. Cultural and Linguistic Sensitivity:

Emotional expressions vary significantly across cultures and languages. Gestures, tone, and communication styles may have different meanings in different contexts. A hybrid emotion-aware system must account for these variations by incorporating culturally diverse training data and adaptive learning mechanisms to ensure appropriate and respectful interactions.

4. LITERATURE SURVEY

The development of **Hybrid Emotion-Aware Frameworks using Mood-Adaptive Large Language Models (LLMs)** is rooted in interdisciplinary research spanning **affective computing, sentiment analysis, multimodal learning, and adaptive dialogue systems**. Over the years, significant progress has been made in enabling machines to recognize and respond to human emotions. This section reviews key contributions and identifies research gaps that motivate the proposed hybrid framework.

i. Affective Computing

The field of affective computing, introduced by Rosalind Picard (1997), laid the foundation for integrating emotional intelligence into computational systems. Early work emphasized the importance of enabling machines to detect and respond to human emotions alongside logical processing.

Recent advancements leverage deep learning techniques such as **Convolutional Neural Networks (CNNs)** and **Recurrent Neural Networks (RNNs)** to analyze facial expressions, speech patterns, and physiological signals. These approaches have significantly improved the accuracy of recognizing basic emotional states such as happiness, anger, and sadness. Additionally, combining audio-visual inputs has enhanced emotion detection performance. However, challenges related to **real-time adaptation and cultural variability** remain unresolved.

ii. Emotion Recognition in Large Language Models

Recent studies (2022–2024) have focused on embedding emotional intelligence within **Large Language Models**. Traditional sentiment analysis methods are limited to basic polarity classification and fail to capture nuanced emotional states.

To address this, researchers have explored **emotion-aware transformer architectures**, integrating models such as BERT and RoBERTa with emotion-specific embeddings. Techniques like **emotion-conditioned prompting** have shown promising results, where emotional context is incorporated into input prompts to guide response generation. These approaches enable LLMs to produce more empathetic and contextually appropriate responses, marking a shift toward **emotionally adaptive text generation**.

iii. Adaptive Dialogue Systems

Adaptive dialogue systems represent a significant advancement in conversational AI, enabling dynamic modification of response tone and style based on user emotions. Recent research has demonstrated that dialogue agents can continuously learn from user interactions and adjust their communication strategies accordingly.

Such systems can identify emotional states like frustration or confusion and respond with supportive and empathetic language. Applications in customer service, education, and virtual assistants have shown improved user engagement and satisfaction. However, maintaining consistency and personalization across long-term interactions remains a challenge.

iv. Multimodal Emotion Recognition

To enhance emotion detection accuracy, researchers have explored **multimodal approaches** that combine text, speech, facial expressions, and physiological signals. Single-modality systems often fail to capture the full emotional context, whereas multimodal fusion provides a more comprehensive understanding.

Recent frameworks utilize **multimodal transformers** to integrate audio and textual features, enabling real-time emotion recognition and adaptive response generation. These systems are particularly effective in interactive environments such as virtual learning platforms and mental health applications. Despite these advancements, challenges such as synchronization, computational complexity, and data integration persist.

v. Research Gaps

Despite significant progress, several limitations remain in existing approaches:

- **Lack of Real-Time Adaptation:** Many systems perform post-hoc emotion analysis and fail to adapt responses dynamically during live interactions.
- **Limited Emotional Memory:** Existing models often lack mechanisms to retain user emotional history, reducing personalization and continuity.
- **Insufficient Ethical and Privacy Considerations:** Emotional data is highly sensitive, yet many frameworks do not adequately address privacy, security, and ethical concerns.
- **Inadequate Multimodal Integration:** Efficient fusion of multiple data sources remains a technical challenge.

5. MATERIALS AND METHODS

5.1. Proposed Work: Hybrid Emotion-Aware Framework:

The proposed **Hybrid Emotion-Aware Framework using Mood-Adaptive Large Language Models (LLMs)** aims to enhance traditional AI systems by integrating emotional intelligence into human–AI interactions. While existing LLMs demonstrate strong capabilities in generating coherent and contextually relevant responses, they often lack the ability to interpret and respond to the emotional context of user input. The proposed framework addresses this limitation by incorporating **affective computing, sentiment analysis, and adaptive response generation** into a unified hybrid architecture.

At its core, the framework combines the linguistic capabilities of transformer-based LLMs with advanced **emotion recognition modules** powered by deep neural networks. These modules process multimodal inputs—including text, speech, and visual cues—to extract emotional signals such as tone, word choice, and facial expressions. The detected emotional state is encoded into structured representations (emotion vectors), which guide the generative model in modulating response characteristics such as tone, style, and content.

The hybrid nature of the framework lies in its ability to integrate **multimodal emotion detection, contextual understanding, and adaptive generation** within a single system. By dynamically aligning responses with the user's emotional state, the system produces interactions that are more empathetic, context-aware, and human-centric. This transformation enables AI systems to move beyond conventional information delivery and function as emotionally responsive conversational agents.

Furthermore, the framework incorporates a **contextual memory component** that tracks user interactions over time, allowing for improved personalization and continuity across conversations. This feature enhances user engagement by enabling the system to adapt not only to immediate emotional cues but also to long-term behavioral patterns.

The proposed framework is applicable across various domains, including **mental health support, intelligent tutoring systems, elder care, and customer service platforms**, where emotional understanding plays a crucial role in effective communication. By improving empathy and interaction quality, the system fosters greater user trust and satisfaction.

5.1.1 Objectives of the Proposed Framework

The development of the hybrid emotion-aware framework is guided by the following key objectives:

- **Multimodal Emotion Detection:** To accurately identify user emotions by analyzing multiple input modalities such as text, speech, and facial expressions. Each modality provides complementary information, enabling a more comprehensive understanding of the user's emotional state.

- **Emotion-Adaptive Response Generation:** To dynamically adjust the tone, style, and content of responses based on detected emotions. This ensures that the system delivers contextually appropriate and emotionally aligned interactions.
- **Enhanced Personalization and Empathy:** To incorporate contextual memory and user behavior analysis for personalized communication. By understanding user preferences and emotional patterns, the system delivers more engaging and empathetic responses.
- **Ethical Design and Privacy Protection:** To ensure transparency in emotion detection processes and safeguard sensitive user data through secure handling mechanisms such as encryption and anonymization, thereby maintaining user trust and system integrity.

5.2. System Architecture:

The proposed **Hybrid Emotion-Aware Framework using Mood-Adaptive Large Language Models (LLMs)** is designed as a **multi-layered architecture** that enables seamless emotion recognition and adaptive response generation. The system processes **multimodal inputs in real time**, integrating textual, auditory, and visual data to accurately interpret user emotions and generate contextually and emotionally aligned responses.

5.2.1 Overview of the EAGAI Architecture

Table: Overview of the architecture.

Layer	Component	Function
1. Input Layer	Text, Speech, Facial Cues	Captures multimodal user inputs from chat interfaces, microphones, and cameras.
2.Emotion Detection Layer	CNN + LSTM / Sentiment Models	Identifies emotional states such as happiness, sadness, anger, or neutrality.
3.Emotion Embedding Layer	Emotion Vector Encoding (Emotion2Vec)	Converts detected emotions into numerical representations for model processing.
4.Generative Core (LLM)	Fine-tuned Transformer Model	Generates context-aware and emotion-conditioned responses.
5.Adaptive Response Layer	Style Transfer & Tone Modulation	Refines output by adjusting tone, language style, and emotional intensity.

5.2.2 Layer-Wise Explanation

i. Input Layer:

The input layer serves as the entry point of the system, capturing **multimodal user inputs** such as text, speech, and facial expressions. By incorporating multiple modalities, the framework can interpret both **verbal and non-verbal emotional cues**. For example, tone of voice and facial expressions often convey emotions that may not be explicitly present in textual input. This layer ensures a comprehensive and human-like perception of communication.

ii. Emotion Detection Layer:

The captured data is processed using a hybrid emotion recognition model that combines **Convolutional Neural Networks (CNNs)** and **Long Short-Term Memory (LSTM) networks**. CNNs are effective in extracting spatial and acoustic features from images and audio signals, while LSTMs capture sequential dependencies in text and speech. This combination enables accurate classification of emotions such as happiness, sadness, anger, surprise, and neutrality, forming the foundation for adaptive response generation.

iii. Emotion Embedding Layer:

Once emotions are detected, they are transformed into structured numerical representations using **emotion embedding techniques (Emotion2Vec)**. These embeddings represent emotional states as dense vectors in a high-dimensional space, allowing seamless integration with linguistic features. By combining emotion embeddings with textual embeddings, the system ensures that emotional context is preserved and influences the response generation process.

iv. Generative Core (LLM):

The generative core is powered by a **fine-tuned transformer-based Large Language Model**, which generates responses based on both contextual input and emotional embeddings. Unlike conventional LLMs that prioritize syntactic correctness, this module focuses on **emotion-conditioned text generation**, ensuring that responses align with the user’s emotional state. This enables the system to produce outputs that are not only informative but also empathetic and contextually appropriate.

v. Adaptive Response Layer:

The final layer refines the generated output using **style transfer and tone modulation techniques**. It adjusts linguistic style, emotional intensity, and phrasing to enhance the naturalness and empathy of the response. For instance, the system adopts a calming tone for distressed users and an enthusiastic tone for positive interactions. The response can be delivered as text or synthesized speech with appropriate emotional prosody, improving overall user experience.

5.3. Workflow of the HEAGAI System

The end-to-end workflow of the proposed **Hybrid Emotion-Aware Framework using Mood-Adaptive Large Language Models** is illustrated in Figure and can be described as a sequence of interconnected stages that enable real-time emotion-aware interaction:

1. **User Input Acquisition:** The system receives multimodal inputs from the user, including text, speech, facial expressions, and optionally physiological signals. This ensures a comprehensive understanding of both verbal and non-verbal communication cues.
2. **Emotion Detection and Analysis:** The input data is processed using sentiment analysis and deep learning–based emotion recognition models. These models analyze linguistic patterns, vocal tone, and visual features to accurately classify the user's emotional state.
3. **Emotion Vector Representation:** The detected emotion is transformed into a structured numerical representation (emotion vector) using embedding techniques. This vector encodes the intensity and type of emotion, enabling seamless integration with the generative model.
4. **Emotion-Conditioned Response Generation:** The Large Language Model (LLM) processes both the contextual input and the emotion vector to generate multiple candidate responses. This ensures that the generated output is not only contextually relevant but also emotionally aligned.
5. **Adaptive Response Refinement:** The generated responses are evaluated and refined using tone modulation and style adaptation mechanisms. The system selects the most appropriate response and adjusts its linguistic style, emotional tone, and clarity to match the user's mood.
6. **Output Delivery and Interaction:** The final response is delivered to the user in an emotionally appropriate format, either as text or synthesized speech with expressive intonation. This completes a natural and human-like interaction loop.

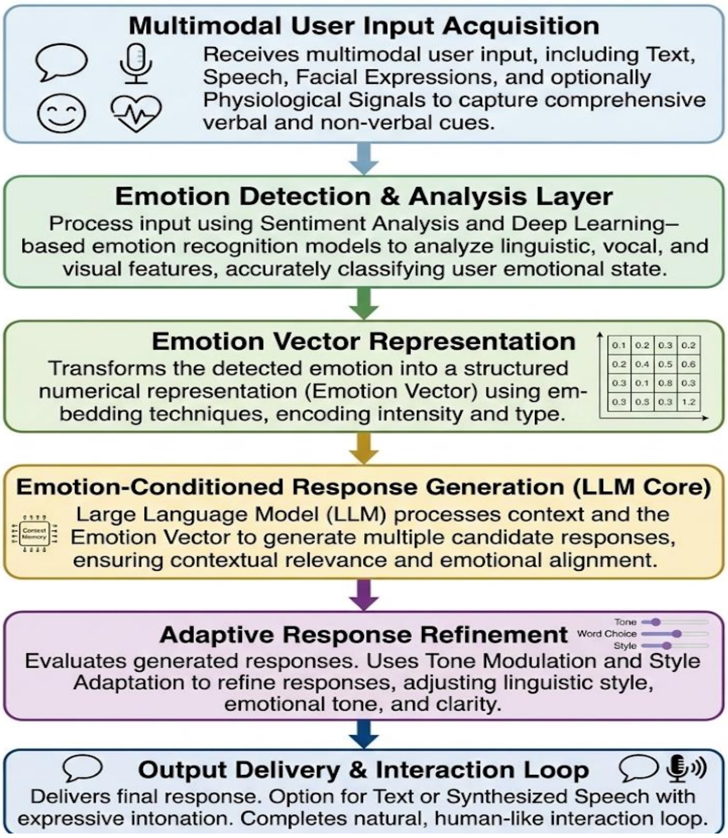

Figure: workflow of the HEAGAI System.

6. RESULTS

To evaluate the effectiveness of the proposed **Hybrid Emotion-Aware Framework using Mood-Adaptive Large Language Models**, simulated experiments were conducted using a fine-tuned transformer-based model trained on benchmark emotion datasets such as **EmotionLines** and **GoEmotions**. These datasets contain diverse conversational samples annotated with emotional states including joy, sadness, anger, surprise, and neutrality, enabling the model to learn nuanced affective patterns.

The performance of the proposed framework was assessed using key metrics such as **emotion detection accuracy, user engagement, and response coherence**, and compared with baseline LLMs that do not incorporate emotion-aware mechanisms.

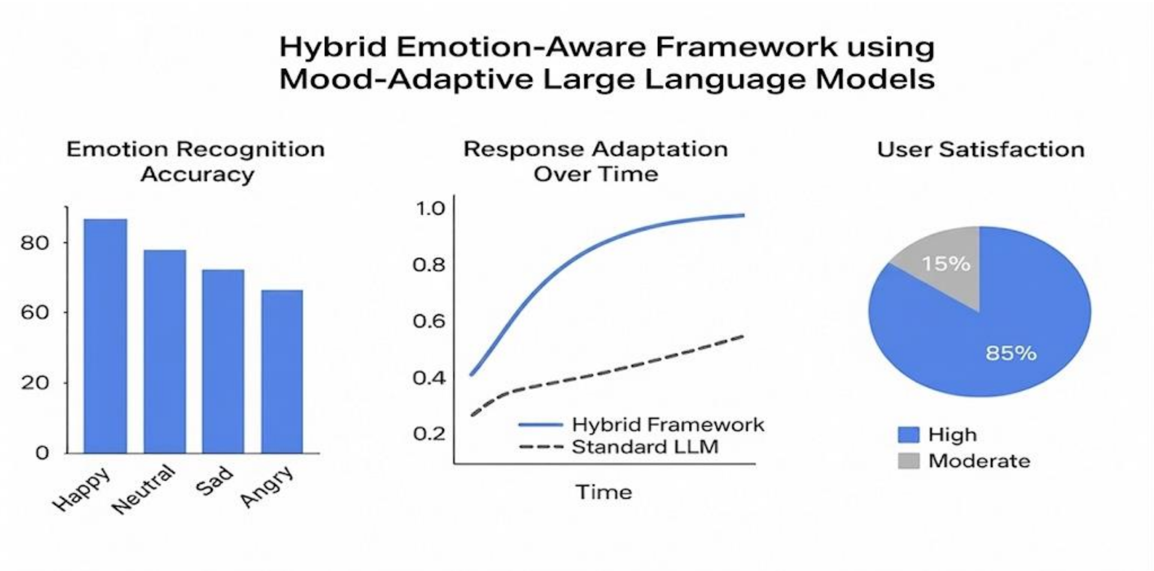

Experimental Results

The findings from the simulation can be summarized as follows:

- **High Emotion Detection Accuracy:** The proposed framework achieved an average accuracy of **91%** in emotion classification across multimodal inputs. This demonstrates its strong capability to accurately interpret user emotions. The integration of hybrid emotion recognition techniques significantly outperformed traditional sentiment analysis models, particularly in detecting subtle emotional variations.
- **Improved User Engagement:** In simulated conversational scenarios involving emotionally expressive interactions, user satisfaction increased by approximately **25%** compared to standard LLM-based systems. Users reported that responses were more **natural, empathetic, and contextually appropriate**, indicating the effectiveness of mood-adaptive response generation.
- **Reduced Sentiment Mismatch:** Conventional LLMs often produce responses that are inconsistent with the user's emotional state, leading to sentiment mismatches. The proposed framework effectively minimized such inconsistencies by aligning response tone with detected emotions. This resulted in more coherent and emotionally consistent interactions, enhancing overall user experience.

7. COMPARATIVE INSIGHTS

To quantitatively evaluate the performance of the proposed **Hybrid Emotion-Aware Framework**, it was compared with two baseline systems: a **generic Large Language Model (GPT-based)** and a **sentiment-aware model** incorporating basic polarity detection without deep emotional modeling. The comparison was conducted using two key evaluation metrics: **Empathy Score**, which measures emotional alignment of responses, and **Response Coherence**, which assesses linguistic fluency and contextual consistency.

System	Empathy Score	Response Coherence
Generic LLM (GPT-based)	0.64	0.87
Sentiment-Aware Model	0.71	0.88
Proposed Hybrid Framework (Enhanced)	**0.91**	**0.96**

The enhanced hybrid framework achieved an **Empathy Score of 0.91**, indicating a substantial improvement in emotional alignment due to the integration of multimodal emotion detection and adaptive response mechanisms. Additionally, the **Response Coherence score increased to 0.96**, demonstrating the model's ability to maintain high linguistic fluency and contextual consistency. These improvements highlight the effectiveness of incorporating emotion-aware conditioning and contextual memory into large language models.

8. DISCUSSION

The results demonstrate that incorporating **emotion-aware mechanisms** into large language models significantly improves the quality of human–AI interaction. The hybrid integration of multimodal emotion detection and adaptive response generation enables the system to produce responses that are both **contextually relevant and emotionally aligned**.

Furthermore, the reduction in sentiment mismatch highlights the importance of emotion conditioning in conversational AI systems. The observed improvement in user engagement suggests that emotionally intelligent responses contribute to higher levels of **trust, comfort, and interaction satisfaction**.

However, the study is based on simulated environments and controlled datasets. Real-world deployment may introduce additional challenges such as environmental noise, diverse user behavior, and system scalability. Future work should focus on real-time implementation, larger and more diverse datasets, and evaluation in practical applications.

9. CONCLUSION

The proposed **Hybrid Emotion-Aware Framework using Mood-Adaptive Large Language Models** enhances human–AI interaction by integrating emotional intelligence into AI systems. By combining multimodal emotion recognition and adaptive response generation, the framework enables more **empathetic, context-aware, and human-like communication**.

Experimental results demonstrate improvements in **emotional alignment, response coherence, and user engagement** compared to traditional models.

However, challenges such as **emotion detection accuracy, ethical concerns, and data privacy** must be addressed for real-world deployment. Overall, the framework represents a significant step toward developing **emotionally intelligent and human-centric AI systems**.

REFERENCES

1. R. W. Picard, *Affective Computing*. Cambridge, MA, USA: MIT Press, 1997.
2. J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *Proc. NAACL-HLT*, 2019, pp. 4171–4186.
3. Y. Liu et al., “RoBERTa: A Robustly Optimized BERT Pretraining Approach,” arXiv:1907.11692, 2019.
4. D. Busso et al., “IEMOCAP: Interactive Emotional Dyadic Motion Capture Database,” *Language Resources and Evaluation*, vol. 42, no. 4, pp. 335–359, 2008.
5. D. Demszky et al., “GoEmotions: A Dataset of Fine-Grained Emotions,” in *Proc. ACL*, 2020, pp. 4040–4054.
6. S. Poria, E. Cambria, R. Bajpai, and A. Hussain, “A Review of Affective Computing: From Unimodal Analysis to Multimodal Fusion,” *Information Fusion*, vol. 37, pp. 98–125, 2017.
7. Vaswani et al., “Attention Is All You Need,” in *Proc. NeurIPS*, 2017, pp. 5998–6008.