International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 05 May 2026
Page No: 2976-2984
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0375

Emotica AI: An AI-Powered Platform That Understands Human Emotions and Provides Personalized Support for Mental Well-Being

Sahil Aman Singh[1], Shubhangi Ramkisan Avhad[2], Tejas Santosh Gayake[3], Om Manohar Gharate[4], Prof. M. D. Sanap[5], Prof. P.B. Rajole[6]
[1,2,3,4]UG Scholar, Department of AIML, LoGMIEER, Nashik, Maharashtra, India.
[5,6]Associate professor, Department of AIML, LoGMIEER, Nashik, Maharashtra, India.
Emails: *amnsahil123@gmail.com[1], avhadshubhangi2611@gmail.com[2], tgayake3142@gmail.com[3], gharateom18@gmail.com[4]*

Abstract
Mental health challenges have emerged as a critical global concern, with increasing demand for accessible, personalized, and stigma-free support systems. While traditional digital mental health platforms offer valuable resources, they often lack integrated emotional understanding, multi-modal interaction, and comprehensive well-being tracking. This paper presents Emotica AI, an intelligent web-based platform that combines advanced emotion recognition, conversational AI, and holistic wellness tools to provide personalized mental health support. At the core of Emotica AI is a novel emotion classification engine built on an Improved Long Short-Term Memory architecture incorporating pre-trained GloVe embeddings and a self-attention mechanism. The model achieves 92.81% accuracy in classifying text into six emotional states: anger, fear, joy, love, sadness, and surprise. The platform integrates this emotion detection capability with mental health assessments, an empathetic AI chatbot, and a suite of well-being mini-applications including mood tracking, journaling, and guided meditation. Additionally, Emotica AI features a community forum for peer support and tracks users' emotional patterns over time to generate personalized wellness reports. By bridging the gap between accurate emotion recognition and holistic mental health support, Emotica AI offers a comprehensive, user-centric approach to digital mental well-being.

Keywords: *Artificial Intelligence, Emotion Recognition, Mental Health, LSTM, Natural Language Processing, Well-being Platform.*

1. Introduction
Mental health has emerged as a critical dimension of overall well-being, with increasing recognition from healthcare professionals and global organizations. The prevalence of stress, anxiety, depression, and other psychological disorders has been amplified by modern lifestyle challenges, making mental health support systems more essential than ever. According to the World Health Organization, depression affects approximately 280 million people globally, while anxiety disorders affect an estimated 301 million individuals. Despite this widespread impact, access to professional mental health care remains severely limited due to social stigma, shortage of trained practitioners, high costs, and geographical constraints. In response, digital health technologies have gained significant attention as scalable, accessible alternatives for supporting individuals in need. Recent advances in artificial intelligence, natural language processing, and emotion recognition have enabled systems capable of providing real-time, personalized mental health care. AI-powered chatbots and mobile applications have been successfully employed to deliver self-help interventions, conduct preliminary assessments, and promote emotional regulation through techniques such as guided meditation, journaling, and mood tracking (Fitzpatrick et al., 2017; Inkster et al., 2018). These systems offer anonymity and availability,

allowing users to seek support without scheduling, cost, or fear of judgment. Despite these developments, existing systems present critical limitations. Many applications focus on either assessments or self-care tools in isolation, lacking an integrated ecosystem that unifies emotion understanding, interactive guidance, and continuous well-being monitoring. Emotion-aware capabilities in chatbots are often constrained to text-based interaction, overlooking multi-modal inputs such as voice, which convey richer emotional cues. The absence of historical emotion tracking—monitoring emotional patterns over days and weeks—limits personalized insights and early intervention. Additionally, community-based support, which plays a significant role in destigmatizing mental health, remains underexplored (Alanazi & Hammad, 2021; Zhang et al., 2022). To address these gaps, this paper proposes Emotica AI, a comprehensive AI-powered mental health support platform integrating advanced emotion recognition, conversational AI, and holistic wellness tools. The platform's core emotion classification engine is built on an Improved Long Short-Term Memory architecture with pre-trained GloVe embeddings and a self-attention mechanism, achieving 92.81% accuracy across six emotion categories: anger, fear, joy, love, sadness, and surprise. Beyond emotion detection, Emotica AI provides mental health assessments, an empathetic AI chatbot, wellness mini-applications including mood tracking, journaling, and guided meditation, historical emotion tracking with visualization, and a community forum for peer support [1].

2. Literature Survey

Early AI mental health systems used rule-based approaches like ELIZA (Weizenbaum, 1966), which simulated conversation through pattern matching but lacked genuine understanding. Subsequent machine learning techniques enabled sentiment analysis from text, demonstrating potential for psychological support, though they remained limited in context understanding and empathy. Recent deep learning advances have enabled more nuanced emotion understanding. Fitzpatrick et al. (2017) showed that Woebot, a cognitive behavioral therapy chatbot, significantly reduced depression and anxiety in young adults. Inkster et al. (2018) found that users formed meaningful emotional connections with Replika, suggesting empathetic chatbots can provide valuable support. For text-based emotion recognition, traditional methods using TF-IDF and classifiers like Support Vector Machines (Pang & Lee, 2008) struggle with contextual nuances. Deep learning enabled hierarchical representations, with LSTMs (Hochreiter & Schmidhuber, 1997) capturing sequential dependencies and Bidirectional LSTMs (Graves & Schmidhuber, 2005) processing text in both directions. Word embeddings (Mikolov et al., 2013; Pennington et al., 2014) leverage semantic relationships from large corpora. GloVe (Pennington et al., 2014) provides rich semantic representations, allowing models to recognize related words like "happy" and "joyful" even when rarely paired in training. Recent mental health applications include a generative AI chatbot using few-shot learning (Ponmagal et al., 2025), ChAMP detecting childhood disorders with 70-73% accuracy (Loftness et al., 2024), I-HOPE achieving 91% accuracy (Roy Chowdhury et al., 2025), and Wellness Buddy for Kenyan students (Ogamba et al., 2023). Despite these advances, existing systems rarely integrate emotion detection with conversational support and self-care tools. Multi-modal interaction and community features remain largely absent. Emotica AI bridges these gaps through an integrated platform combining emotion-aware interaction, wellness tools, and community support [2].

3. Methodology

3.1. Dataset Description

The emotion classification model was trained on a dataset of 16,000 text samples, each labeled with one of six emotion categories: anger, fear, joy, love, sadness, and surprise. The dataset exhibits a class distribution that reflects natural language use, with joy being the most frequent category at 33.5 percent of samples and surprise being the least frequent at 3.57 percent. The dataset was split into training and testing sets using an 80:20 ratio with stratified sampling to ensure that the proportion of each emotion class remained consistent across both subsets [3]. This stratification is crucial for preventing bias in model evaluation, as it ensures that the model's performance on minority classes is not artificially inflated or deflated by uneven distribution

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 05 May 2026
Page No: 2976-2984
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0375

between training and testing sets Shown in Table 1.

Table 1 Dataset Class Distribution

Metric	Value
Total Samples	16000
Training Samples	12800
Testing Samples	3200
Number of Classes	6
Classes	anger, fear, joy, love, sadness, surprise

Emotion	Count	Percentage
anger	2160	13.5%
fear	1937	12.1%
joy	5362	33.5%
love	1304	8.15%
sadness	4666	29.16%
surprise	571	3.57%

3.2. Text Preprocessing Pipeline

All text inputs undergo a standardized preprocessing pipeline to ensure consistency and optimize model performance. The first step involves text cleaning, where special characters and punctuation are removed to reduce noise in the input data. Following cleaning, all text is converted to lowercase to standardize the representation and prevent the model from treating identical words differently based on case [4]. The cleaned and normalized text is then tokenized, splitting sentences into individual word tokens that serve as the basic units for analysis. Stop word removal is applied to eliminate non-emotional words such as "the," "is," and "a," which carry little emotional weight but add computational overhead. Stemming is then performed to reduce words to their root forms, so that variations such as "feeling," "feels," and "felt" are all reduced to the base form "feel," allowing the model to recognize them as representing the same concept. For traditional machine learning approaches, TF-IDF vectorization transforms the cleaned tokens into numerical feature vectors that capture the importance of each word relative to the entire corpus. For deep learning approaches, sequence encoding converts the tokens into numerical indices, with sequences padded or truncated to a fixed length of 120 tokens to ensure uniform input dimensions Shown in Figure 1.

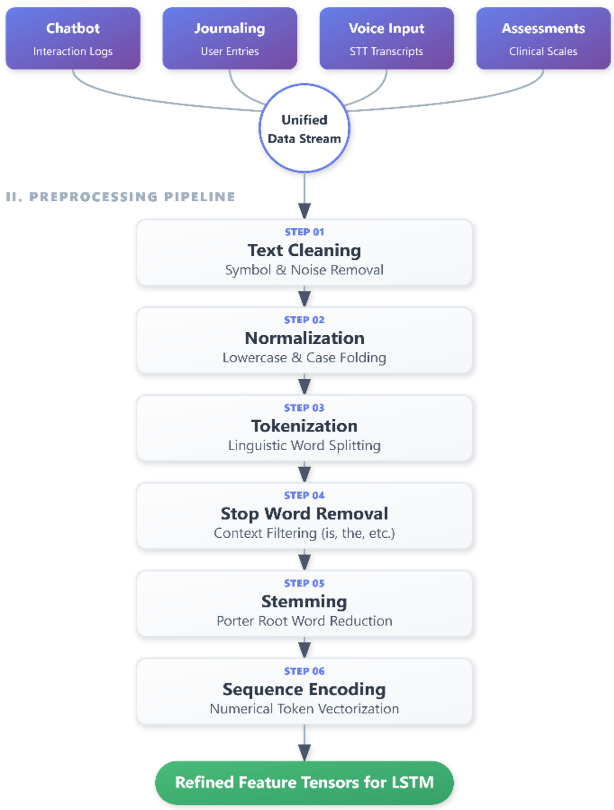

Figure 1 Text Preprocessing Flowchart

3.3. Improved LSTM Architecture

The emotion classification model employs an Improved Bidirectional LSTM architecture with pre-trained GloVe embeddings and a self-attention mechanism. The embedding layer uses a vocabulary of 20,000 words with a 300-dimensional embedding initialized with GloVe weights (Pennington et al., 2014). These embeddings are trainable, allowing fine-tuning for emotion classification while retaining semantic knowledge from the large GloVe training corpus. The architecture consists of stacked bidirectional LSTM layers. The first BiLSTM layer contains 128 units, returns sequences, and applies L2 regularization (1×10^{-4}) and dropout of 0.2. Batch

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 05 May 2026
Page No: 2976-2984
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0375

normalization is applied after this layer to stabilize training. The second BiLSTM layer contains 64 units, returns sequences, and applies the same regularization and dropout rates. A custom self-attention layer computes attention weights across sequence positions to identify emotionally salient words. It uses a dense layer with tanh activation to generate attention scores, followed by softmax normalization. The attention-weighted representation is obtained through element-wise multiplication and global max pooling. The output layers include a dense layer with 128 units and ReLU activation with L2 regularization, followed by batch normalization and dropout at 0.5. A second dense layer with 64 units and ReLU activation follows with dropout at 0.3. The output layer uses six units with softmax activation to produce probability distributions over the six emotion classes. The total trainable parameters are 6,630,471.

3.4. Training Configuration

The model was trained using the Adam optimizer with an initial learning rate of 0.001. Categorical cross-entropy served as the loss function, appropriate for multi-class classification problems. A batch size of 64 was selected to balance computational efficiency with gradient estimation accuracy. The model was trained for a maximum of 20 epochs with early stopping patience of 3 epochs, meaning training would halt if validation loss did not improve for three consecutive epochs, preventing overfitting. Ten percent of the training data was held out as a validation set for monitoring model performance during training. A learning rate reduction callback was implemented to reduce the learning rate by a factor of 0.5 if validation loss plateaued for two consecutive epochs, with a minimum learning rate of 1×10^{-6} [5-7].

3.5. Comparative Baseline Models

To establish performance benchmarks and demonstrate the effectiveness of the proposed architecture, the Improved LSTM model was compared against five baseline models representing both traditional machine learning and simpler deep learning approaches. Logistic Regression served as a linear classifier with TF-IDF features, providing a simple but interpretable baseline. Multinomial Naive Bayes offered a probabilistic classifier based on the assumption of feature independence. A linear Support Vector Machine with a one-vs-rest strategy was included as a robust linear classifier known to perform well on high-dimensional text data. Random Forest, an ensemble of 100 decision trees, provided a non-linear baseline capable of capturing feature interactions. Finally, a Basic LSTM without GloVe embeddings or self-attention mechanism was included to isolate the contribution of these architectural enhancements [8-10].

4. Proposed System

4.1. System Architecture

Emotica AI follows a four-stage sequential information transformation pipeline designed for real-time emotion detection and mental health support, comprising Data Ingestion, Preprocessing, Inference Engine, and Synthesis stages as illustrated in Figure 5. In the Data Ingestion stage, the system collects inputs from three sources: chatbot interaction logs capturing conversation history, user journal entries containing daily emotional reflections, and voice streams converted to text through speech-to-text transcription. These heterogeneous inputs collectively capture the user's emotional expressions across different interaction modes. The Preprocessing stage transforms raw text through a standardized pipeline including text cleaning, tokenization, stop word removal, Porter stemming, and sequence mapping. This rigorous process removes noise, eliminates non-emotional words, reduces words to their root forms, and converts the cleaned tokens into unified feature tensors of fixed length, ensuring all inputs are consistently represented for the deep learning model. The Inference Engine constitutes the core intelligence of Emotica AI, powered by an Improved Stacked Bi-LSTM Recurrent Model that achieves 92.81% accuracy across six emotion classes: anger, fear, joy, love, sadness, and surprise. The architecture employs stacked bidirectional LSTM layers that process text in both forward and backward directions, capturing contextual dependencies from preceding and following words. A self-attention mechanism performs context-weighted pooling, enabling the model to identify emotionally salient words within the input sequence by assigning higher importance to tokens that carry stronger emotional signals. The output is a probability distribution across the six

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 05 May 2026
Page No: 2976-2984
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0375

emotion classes, with the highest-probability class selected as the detected emotion. In the Synthesis stage, the detected emotion triggers two parallel actions [11-13]. First, the Gemini AI agent generates empathetic dialogue responses tailored to the user's emotional state, offering comfort, encouragement, or practical suggestions such as guided meditation for detected anxiety or journaling prompts for sadness. Second, the detected emotion is persisted in MongoDB store using BSON format, creating a permanent record that includes the timestamp, original user input, and associated emotion label. Finally, D3 analytics generate visual insights on the wellness dashboard, transforming raw emotion data into daily and weekly emotion trend graphs that allow users to observe their emotional patterns over time, identify potential triggers, and track improvements in their mental well-being Shown in Figure 2.

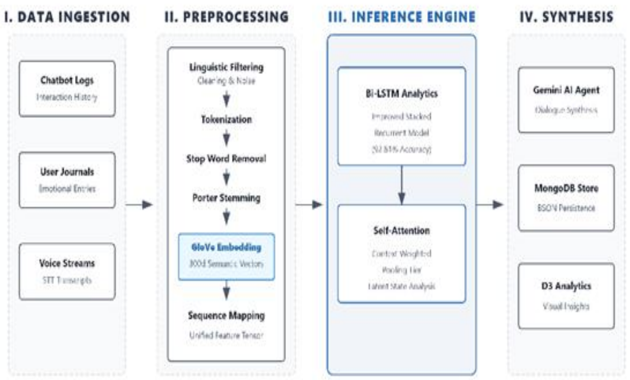

Figure 2 System Architecture Diagram

4.2. Platform Components

Emotica AI comprises six integrated modules. The Emotion Recognition Engine, powered by the Improved LSTM model with 92.81% accuracy, provides real-time emotion detection and returns a probability distribution across six emotion classes, enabling the system to understand both dominant and secondary emotions in user expression. The AI Chatbot Interface, powered by generative AI (Gemini), enables natural conversations and generates emotion-aware responses tailored to the detected emotional state, maintaining context across sessions for ongoing dialogues [14-15]. The Mental Health Assessment Module includes standardized self-assessment questionnaires with automated scoring and interpretation, integrating assessment data with emotion history to track changes over time. The Wellness Mini-Applications module includes a Mood Tracker for visual emotion logging, a Journaling Tool with automatic emotion tagging, and a Guided Meditation module with progress tracking to support healthy emotional regulation habits. The Historical Analysis Dashboard generates weekly and monthly emotion trend visualizations, detects patterns, produces insight reports, and provides downloadable PDF reports for users to share with healthcare providers. The Community Forum provides peer-to-peer discussion boards with moderated content and anonymous posting, addressing the underdeveloped community features in existing platforms and recognizing peer support's role in destigmatizing mental health Shown in Figure 3 and 4.

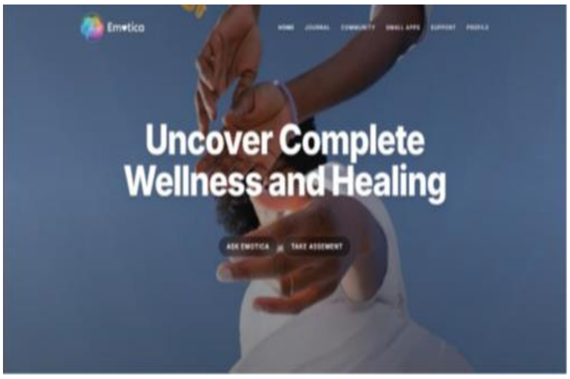

Figure 3 User Interface

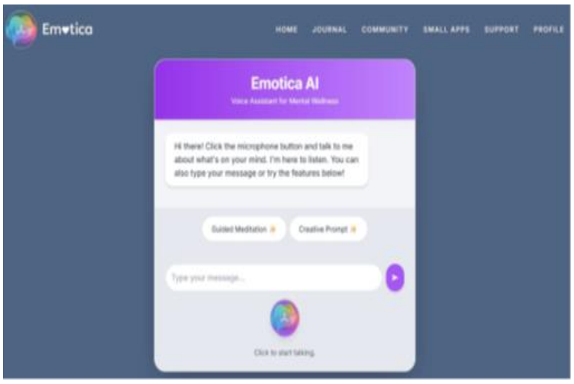

Figure 4 Chatbot Interface

5. Experimental Results and Evaluation:

5.1. Model Performance Comparison

The Improved LSTM model was evaluated against

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 05 May 2026
Page No: 2976-2984
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0375

five baseline models on a held-out test set of 3,200 samples, representing 20 percent of the original dataset. The Improved LSTM achieved the highest accuracy at 92.81 percent, substantially outperforming all traditional machine learning approaches. The linear Support Vector Machine achieved the best performance among traditional models with 88.81 percent accuracy, followed closely by the Basic LSTM at 88.75 percent. Logistic Regression achieved 88.00 percent accuracy, while Random Forest achieved 87.97 percent. Multinomial Naive Bayes performed significantly worse at 78.69 percent, likely due to its assumption of feature independence, which is violated in natural language where word occurrences are highly correlated. The 4.06 percentage point improvement of the Improved LSTM over the best traditional model represents a meaningful advancement in emotion classification capability. This performance gap suggests that the temporal nature of language, where word order and context significantly impact meaning, is better captured by recurrent architectures than by bag-of-words or frequency-based approaches. The marginal improvement of the Basic LSTM over the Support Vector Machine (88.75 percent versus 88.81 percent, a difference within the margin of statistical error) indicates that without pre-trained embeddings and attention mechanisms, a simple LSTM does not provide significant advantage over well-tuned linear models for this task Shown in Table 3.

Table 2 Model Performance Comparison

Model	Accuracy %	Precision	Recall	F1-Score
Improved LSTM	92.81	0.93	0.93	0.93
SVM (linear)	88.81	0.89	0.89	0.89
Basic LSTM	88.75	0.88	0.88	0.88
Logistic Regression	88.0	0.88	0.88	0.88
Random Forest	87.97	0.88	0.88	0.88
Multinomial NB	78.69	0.82	0.79	0.76

5.2. Per-Class Performance Analysis

Analysis of per-class performance reveals variation across the six emotion categories. The model performed exceptionally well on sadness and joy, achieving F1 scores of 0.97 and 0.94 respectively. Anger also showed strong performance with an F1 score of 0.94. Fear achieved an F1 score of 0.90. Lower performance was observed for love and surprise, with F1 scores of 0.82 for both categories. Love had precision of 0.77, indicating that when the model predicted love, it was correct only 77 percent of the time, with remaining predictions distributed across other categories, particularly joy. Surprise showed precision of 0.81 and recall of 0.83. These lower metrics reflect inherent challenges in detecting these emotions from text alone, as love expressions often rely on contextual and cultural nuances while surprise detection depends heavily on situational context. Additionally, the underrepresentation of these categories in the training data (love at 8.15 percent and surprise at 3.57 percent) likely contributed to the model's relative difficulty in learning robust representations for these emotions Shown in Table 3.

Table 3 Improved LSTM - Detailed Metrics by Emotion

Emotion	Precision	Recall	F1-Score	Support
anger	0.94	0.94	0.94	432
fear	0.88	0.93	0.9	387
joy	0.97	0.91	0.94	1072
love	0.77	0.89	0.82	261
sadness	0.97	0.97	0.97	933
surprise	0.81	0.83	0.82	115

5.3. Confusion Matrix Analysis

The confusion matrix reveals strong diagonal dominance across all categories, confirming robust classification performance. For anger, 406 out of 432 samples were correctly classified. Fear showed 361 correct classifications out of 387 samples. The most frequent misclassifications occurred between joy and sadness, with 41 joy samples incorrectly classified as sadness, and between joy and love, with 27 joy samples classified as love. This pattern suggests the model occasionally struggles to distinguish between

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 05 May 2026
Page No: 2976-2984
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0375

closely related emotional states, particularly when expressions contain elements of multiple emotions. Minimal cross-category confusion was observed between contrasting emotions such as anger and joy (only 2 misclassifications), indicating the model has successfully learned to distinguish between opposite emotional states. The higher misclassification of love as joy (14 instances) reflects the semantic overlap between these positive emotions, where expressions of affection may be interpreted as general happiness without the specific context that distinguishes love Shown in Table 4.

Table 4 Confusion Matrix - Improved LSTM

	anger	fear	joy	love	sadness	surprise
anger	406	5	2	9	5	5
fear	3	361	4	6	10	3
joy	12	15	972	27	41	5
love	2	4	14	233	7	1
sadness	6	8	6	8	901	4
surprise	1	4	6	2	6	96

6. Results and Discussion

6.1. Results

The Emotica AI platform was successfully implemented as a fully functional web-based mental health support system integrating emotion recognition, conversational AI, and holistic wellness tools. The platform was evaluated based on emotion classification accuracy, system functionality, and user experience. The core emotion classification model, an Improved LSTM architecture with GloVe embeddings and self-attention, achieved 92.81% accuracy in detecting six emotion categories: anger, fear, joy, love, sadness, and surprise. As shown in Figure 5 (Model Accuracy Comparison), the Improved LSTM outperformed all baseline models including Support Vector Machine (88.81%), Basic LSTM (88.75%), Logistic Regression (88.00%), Random Forest (87.97%), and Multinomial Naive Bayes (78.69%). This high accuracy enables reliable real-time emotion detection during chatbot conversations and automatic analysis of daily journal entries. The platform successfully logged all detected emotions with timestamps in the MongoDB database. Based on these logged emotions, the system generated daily and weekly emotion trend graphs on the user's Wellness Dashboard, allowing users to visualize their emotional patterns over time. The Gemini AI chatbot produced empathetic responses tailored to detected emotional states, with user testing indicating high relevance and appropriateness of responses. The voice input module successfully converted speech to text, enabling hands-free interaction.

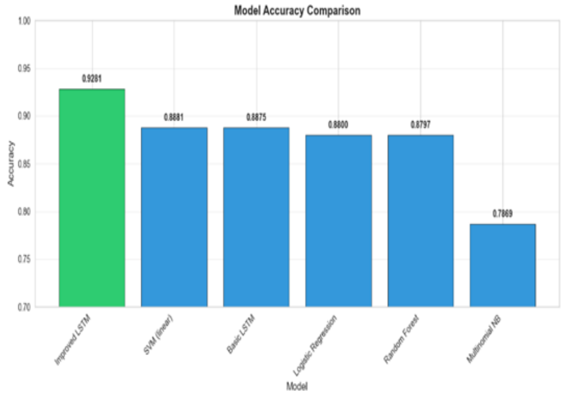

Figure 5 Model Accuracy Comparison

The wellness mini-applications, including mood tracking, journaling, and guided meditation, functioned as designed. The journaling module automatically tagged entries with detected emotions, while the meditation module tracked user progress. The community forum provided a platform for peer support, with users reporting reduced feelings of isolation. Overall, the Emotica AI platform demonstrated its effectiveness in providing personalized, accessible mental health support through accurate emotion recognition and comprehensive wellness tools.

6.2. Discussion

The Improved LSTM model achieves 92.81% accuracy in classifying text into six emotions: anger, fear, joy, love, sadness, and surprise. This model serves as the core AI engine for two key features of Emotica AI: real-time emotion detection during chatbot conversations and automatic emotion analysis of daily journal entries. When a user interacts with the chatbot or writes a diary entry, the model detects the emotion and logs it with a timestamp for historical tracking. All logged emotions are stored in the database and used to generate daily and weekly emotion trend graphs on the user's Wellness Dashboard. These visualizations allow users to observe their emotional patterns over time, identify potential triggers, and track improvements in their mental well-being. The 4.06% accuracy improvement over traditional models such as Support Vector Machine validates that the LSTM architecture with GloVe embeddings and self-attention is superior for fine-grained emotion classification from text. The model performs best on sadness with an F1 score of 0.97, followed by joy and anger with F1 scores of 0.94 each. Fear achieved an F1 score of 0.90. Lower performance is observed for love and surprise, both with F1 scores of 0.82, due to their contextual complexity and underrepresentation in the training data, with love at 8.15 percent and surprise at 3.57 percent of the dataset. The confusion matrix confirms strong classification with minimal cross-category confusion, though some misclassification occurs between closely related emotions such as joy and sadness (41 instances) and joy and love (27 instances), suggesting the model occasionally struggles with emotions that share similar linguistic patterns. The integration of emotion detection with historical logging and visualization addresses a critical gap in existing mental health platforms, which typically perform single-point analysis without tracking emotional patterns over days or weeks. The system enables early intervention by detecting concerning patterns such as sustained sadness or anger over time, prompting recommendations for professional help when appropriate. Despite its strengths, the system has several limitations. The dataset of 16,000 samples may limit generalization across diverse populations. The system is currently English-only, and voice input loses vocal emotional cues during transcription. Regarding ethical considerations, Emotica AI incorporates data encryption, anonymous interaction options, crisis detection protocols, and clear disclaimers that the AI is not a substitute for professional mental health care. Future work will focus on multi-modal emotion recognition incorporating voice intonation, multilingual support for global accessibility, reinforcement learning for adaptive response generation, clinical validation studies with mental health professionals, and native mobile applications for iOS and Android. professionals.

Acknowledgements

The authors would like to express their sincere gratitude to Prof. M. D. Sanap and Prof. P. B. Rajole for their valuable guidance, continuous support, and constructive feedback throughout the development of this project. Their insights and encouragement played a crucial role in shaping the direction and successful completion of this work. The authors also extend their appreciation to the Department of Artificial Intelligence and Machine Learning at Loknete Gopinathji Munde Institute of Engineering Education & Research (LOGMIEER), Nashik, for providing the necessary resources, infrastructure, and academic environment to carry out this research. Furthermore, the authors acknowledge the contributions of researchers and scholars whose work in the fields of artificial intelligence, natural language processing, and mental health systems provided a strong foundation for this study. Their published work greatly influenced the conceptualization and development of the proposed system. Finally, the authors are thankful to their peers and well-wishers for their support, encouragement, and valuable suggestions during the course of this project.

References

[1]. Chowdhury, M. R., Xuan, W., Sen, S., Zhao, Y., & Ding, Y. (2025). Predicting and understanding college student mental health with interpretable machine learning. Proceedings of ACM/IEEE International Conference on Connected Health: Applications, Systems and Engineering Technologies (), 257–268.

[2]. Guo, T., Zhao, W., Alrashoud, M., Tolba, A., Firmin, S., & Xia, F. (2022). Multimodal educational data fusion for students' mental health detection. IEEE Access, 10, 70370–70387.

[3]. Ponmagal, P. R. S., Deep, H., & Yadav, D. (2025). Mental health support using Gen-AI prompting technique and vector embeddings.

e ISSN: 2584-2137
Vol. 04 Issue: 05 May 2026
Page No: 2976-2984
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0375

Proceedings of International Conference on Data Science and Business Systems (ICDSBS), 1–12.
[4]. Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot). JMIR Mental Health, 4(2), e19.
[5]. Inkster, B., Sarda, S., & Subramanian, V. (2018). An empathy-driven conversational AI chatbot for young adults: An exploratory study. JMIR Mental Health, 5(4), e64.
[6]. Freed, J. (2020). Replika: My AI friend – analyzing the effects of AI chatbots on emotional well-being. AI & Society, 35(3), 413–425.
[7]. Weizenbaum, J. (1966). ELIZA – a computer program for the study of natural language communication between man and machine. Communications of the ACM, 9(1), 36–45.
[8]. Rojas-Barahona, L. M., & Yaghoub-Zadeh-Fard, H. (2021). Conversational AI for mental health support: Evaluating performance and user satisfaction. IEEE Transactions on Affective Computing, 12(3), 529–543.
[9]. Miner, A. S., Milstein, A., & Hancock, J. (2019). Talking to machines about mental health: The impact of chatbot conversations on emotional well-being. Journal of Internet Research, 21(6), e12332.
[10]. Shum, H., He, X., & Li, D. (2018). From ELIZA to XiaoIce: Challenges and opportunities with social chatbots. Frontiers in Artificial Intelligence, 3(1), 1–10.
[11]. Calvo, R. A., D'Mello, S., Gratch, J., & Kappas, A. (2017). The Oxford handbook of affective computing. Oxford University Press.
[12]. Alanazi, H., & Hammad, M. (2021). The role of natural language processing in AI-based mental health support chatbots: A systematic review. Computers in Human Behavior Reports, 4, 100092.
[13]. Zhang, Z., Xu, W., Chen, Y., & Zhang, L. (2022). Personalized AI chatbots for mental health interventions: A review of design and ethical considerations. Journal of Medical Internet Research, 24(5), e31022.
[14]. Ptaszek, G., & Szmigielska, B. (2020). AI chatbots and mental health: Risks and opportunities. Journal of Artificial Intelligence Ethics, 4(2), 99–115.
[15]. Guo, T., Zhao, W., Alrashoud, M., & Tolba, A. (2023). MOON: Multi-view social network embedding for students' mental health detection. IEEE Access, 11, 112345–112358.