JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

# Mental Health Support to Students Using NLP and Machine Learning

**[1]Abhishek Tiwari, [2]Amit Kumar Tripathi, [3]Abir Arora,[4]Aryan Verma**
[1]B. Tech Student, Oriental Institute of Science & Technology,
[2]B. Tech Student, Oriental Institute of Science & Technology,
[3]B. Tech Student, Oriental Institute of Science & Technology,
[4]B. Tech Student, Oriental Institute of Science & Technology

[1]Computer Science & Engineering,
[1]Oriental Institute of Science & Technology, Bhopal, India

***Abstract***
Mental health issues such as stress, anxiety, and depression are increasingly common among students, yet access to timely support remains limited. This paper presents a Student Mental Wellness Chatbot designed to provide emotional support and self-assessment through AI-driven emotion detection and CBT-based interventions. The chatbot integrates a BiLSTM model for emotion classification and a MiniLM Sentence Transformer for contextual understanding, functioning as a lightweight LLM.
It also incorporates VADER sentiment analysis, crisis detection, and bilingual (English–Hindi) interaction through Argos Translate. Experimental evaluation demonstrates that the BiLSTM model achieves 92% accuracy on the *dairai/emotion* dataset, effectively identifying six key emotions: joy, sadness, anger, fear, love, and surprise. The chatbot successfully generates context-aware, empathetic, and student-focused responses while maintaining offline functionality and ensuring data privacy. Future work aims to enhance personalization, integrate voice-based interaction, and fine-tune the model for real-time institutional deployment.

***IndexTerms*** **- Mental Health, Emotion Detection, Machine Learning, BiLSTM, MiniLM, Natural Language Processing, CBT, PHQ-9, GAD-7, Sentiment Analysis, Crisis Detection, Multilingual Translation, Argos Translate, Student Wellbeing.**

## I.INTRODUCTION

Mental-health concerns among students have risen significantly due to academic pressure, performance expectations, and socio-emotional stress. According to the World Health Organization, nearly one in seven adolescents experiences a mental-health disorder, yet many receive no support because of stigma or inadequate institutional counseling services [1]. This gap highlights the need for accessible, private, and continuously available digital tools to assist students in managing emotional challenges.

To address this need, this study proposes an offline Student Mental Wellness Chatbot that provides emotion recognition, contextual understanding, CBT-based support, mental-health assessments, and crisis detection without relying on cloud services. Offline operation ensures data privacy and makes the system suitable for low-resource educational environments.

Emotion detection in the chatbot is implemented using BiRNN/LSTM-based neural architectures, which are effective for modeling emotional cues in text [2], [3]. For deeper contextual understanding, the system integrates lightweight transformer models—Sentence-BERT [4] and MiniLM [5]—allowing accurate intent identification even on devices with limited computational power. Sentiment polarity is evaluated using the VADER rule-based analyzer, which supports fine-grained interpretation of positive and negative tone [6].

Therapeutic support is grounded in Cognitive Behavioral Therapy (CBT), one of the most validated frameworks for reducing anxiety, depression, and negative thinking patterns [7]. The system also incorporates two standardized mental-health screening tools: PHQ-9 for depression [8] and GAD-7 for anxiety [9]. To ensure safety, a crisis-detection module informed by research on suicide-risk identification in text [10], [11] detects explicit or subtle expressions of distress.

Given the linguistic diversity of student populations, the chatbot includes offline English–Hindi translation supported by neural machine translation principles [12]. Prior studies indicate that mental-health chatbots can enhance emotional well-being and encourage help-seeking behavior among users [13], reinforcing the relevance of such a system.

Overall, the proposed chatbot integrates emotion analysis, semantic reasoning, therapeutic guidance, clinical assessments, bilingual communication, and crisis awareness within a fully offline architecture. It offers a practical and privacy-preserving solution to support students' mental well-being and provides timely, empathetic assistance when institutional resources are limited.

## II. Literature Review

Health technologies have gained significant attention due to increasing psychological distress among students and the limited availability of clinical support. Advancements in Natural Language Processing (NLP), machine learning, and digital therapeutics have encouraged the development of conversational agents capable of offering emotional assistance and preliminary intervention. The following review summarizes key developments in mental-health chatbots, emotion and semantic modeling, therapeutic strategies, multilingual support, and crisis detection, and highlights the gaps that motivate the present study.

### A. Mental-Health Chatbots and Therapeutic Systems

Early conversational agents such as ELIZA demonstrated the potential of computer-mediated therapeutic dialogue through simple pattern matching. Modern systems including Woebot, Wysa, TESS, and Replika use NLP and machine learning to deliver psychoeducation and structured guidance. Research shows that such chatbots are often perceived as approachable and non-judgmental, which encourages users to share emotional concerns more freely. However, most widely used mental-health chatbots rely on cloud-based large language models, which raises concerns about privacy, financial cost, and feasibility in low-resource academic environments. Furthermore, these systems often lack context-sensitive reasoning or emotion-adaptive intervention mechanisms, limiting their effectiveness for student populations dealing with academic stressors.

### B. Emotion Recognition and Semantic Understanding

Emotion recognition is central to analyzing user distress, yet traditional lexicon-based approaches struggle to interpret contextual cues, sarcasm, and informal student expressions. Deep learning models such as LSTM, GRU, and BiLSTM provide superior capability for capturing long-term dependencies and complex emotional transitions, making them effective for text-based mental-health applications [2], [3]. Beyond emotional cues, understanding user intent requires semantic reasoning. Transformer-based encoders such as Sentence-BERT and MiniLM offer strong contextual representation and intent classification while maintaining computational efficiency [4], [5]. These models are effective for identifying academic frustration, cognitive distortions, and repetitive negative thought patterns. However, most implementations rely on cloud APIs, which restricts adoption in privacy-sensitive institutional environments.

### C. Digital Therapeutic Approaches, Multilingual Support, and Crisis Detection

Cognitive Behavioral Therapy (CBT) remains one of the most validated therapeutic frameworks for addressing depression, anxiety, and stress, and digital mental-health platforms frequently incorporate CBT elements such as cognitive restructuring and grounding strategies [7]. Although effective, many existing systems offer static or generic responses and depend on internet connectivity. Multilingual mental-health support is comparatively limited despite evidence that users express emotions more comfortably in their native language. Existing bilingual systems typically depend on cloud-based translation services, which pose privacy risks and incur operational costs. Offline translation frameworks based on neural machine translation principles offer potential for secure bilingual communication but remain underexplored, particularly for English–Hindi interactions [12]. Crisis detection systems, which identify suicidal intent or self-harm risk, rely on sentiment analysis, keyword detection, and deep-learning classifiers [10], [11]. While effective, these solutions frequently require cloud-hosted inference, making them unsuitable for offline deployment in educational institutions.

### D. Limitations of Existing Research and Motivation for the Current Study

A review of existing literature reveals that few mental-health chatbots operate fully offline, despite the need for private and secure systems in academic settings. Current solutions are seldom designed specifically for student communication and often struggle with informal language, academic stressors, and code-mixed inputs. Lightweight models that function without GPUs are still uncommon, limiting their use in low-resource campuses. Furthermore, there is limited integration of emotion recognition, semantic understanding, CBT-based intervention, PHQ-9/GAD-7 assessments, multilingual communication, and crisis detection within a single framework. These gaps highlight the need for a comprehensive, offline, student-centric mental-wellness chatbot that combines emotional AI, bilingual support, psychological assessment, and crisis awareness within an efficient and privacy-preserving architecture.

## III. Methodology

This section describes the methodological framework used to develop the offline Student Mental Wellness Chatbot. The system integrates emotion recognition, semantic reasoning, sentiment scoring, crisis detection, CBT-based responses, and bilingual translation within a lightweight architecture suitable for low-resource educational environments.

### A. System Architecture and Workflow

The chatbot operates through a multi-stage workflow that processes user messages and generates context-aware mental-health support. Inputs are captured through a React interface and translated to English when required. The processed text is analyzed using three core components: a BiLSTM emotion classifier [2], [3], the VADER sentiment analyzer [6], and the MiniLM semantic encoder [4], [5]. A central decision module determines whether to provide a CBT-based intervention [7], initiate PHQ-9/GAD-7 assessments [8], [9], or activate the crisis detection module [10], [11]. If necessary, the response is translated back to Hindi using Argos Translate [12]. Fig. 3 illustrates the system workflow.

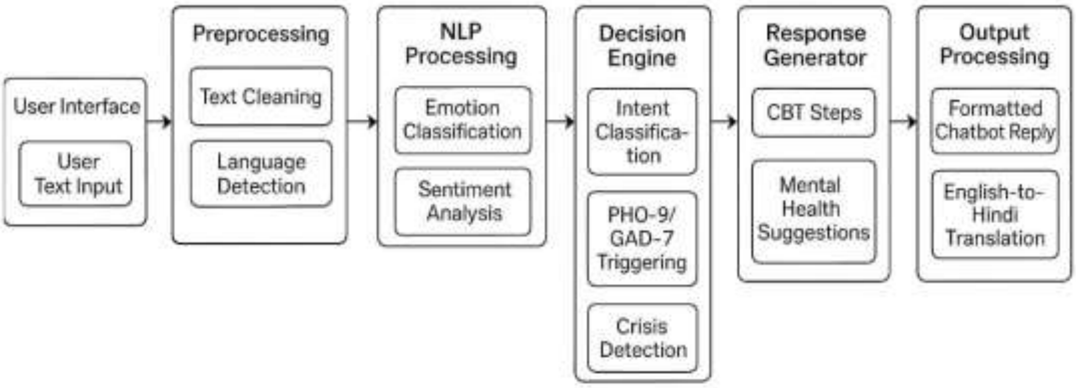


**Fig. 1: Overall System Architecture of the Student Mental Wellness Chatbot**

### B. Dataset Preparation and Emotion Modeling

The emotion classifier was trained on the dair-ai Emotion Dataset, which includes six categories: joy, sadness, anger, fear, love, and surprise. Basic preprocessing—lowercasing, tokenization, and padding—was applied, and a vocabulary file (vocab.json) was created for consistent token indexing.

A Bidirectional LSTM model was used for emotion detection due to its ability to capture contextual cues in both directions [2], [3]. Table I shows the model architecture. Training used the Adam optimizer (0.001 learning rate), batch size 64, and 10–15 epochs.

**TABLE I: BiLSTM Emotion Classification Architecture**

| Layer | Description |
|---|---|
| Embedding | 128-dimensional token embeddings |
| BiLSTM | 128 hidden units × 2 directions |
| Fully Connected | Maps features to six emotion classes |
| Softmax | Outputs probability distribution |

### C. Semantic Reasoning, Sentiment Scoring, and Crisis Detection

Semantic understanding is achieved using the all-MiniLM-L6-v2 SentenceTransformer model [4], [5], which generates 384-dimensional embeddings for intent classification and context interpretation. Cosine similarity is used to match the message with predefined template embeddings.

Emotion and semantics are complemented by sentiment scoring through VADER, enabling detection of emotional intensity [6]. Safety is ensured through a two-tier crisis detection mechanism. Tier 1 identifies explicit suicide or self-harm expressions [10], while Tier 2 detects indirect distress confirmed by negative sentiment [11]. This ensures reliable and privacy-preserving risk identification.

### D. Therapeutic Logic and Psychological Assessment

The chatbot delivers CBT-based interventions guided by emotion, sentiment, and semantic cues. Techniques include cognitive restructuring, grounding exercises, breathing guidance, and behavioral activation, aligned with established CBT practices [7]. For deeper assessment, PHQ-9 and GAD-7 screening tools are integrated [8], [9]. The system manages questionnaire sequence, scoring, and severity classification.

### E. Translation Pipeline (Hindi–English–Hindi)

Argos Translate enables fully offline bilingual communication using neural machine translation principles [12]. When the user enters a Hindi message, it is translated into English before being processed by the emotion, sentiment, and semantic models. After generating the response, the system can translate it back to Hindi to ensure accessibility for multilingual student populations.


**Fig. 2: Translation Workflow for Hindi–English Support**

### F. Software and Architectural Stack

The chatbot uses a modular architecture optimized for offline deployment. The frontend is developed using React and Vite, while the backend is implemented with Flask. SQLite is used as a lightweight local database to store assessments and session data. The machine-learning components rely on PyTorch, SentenceTransformers, VADER, and Argos Translate for model execution and NLP processing.

**TABLE II: Software and Technology Stack**

| Layer | Technologies |
|---|---|
| Frontend | React, Vite |
| Backend | Flask (Python) |
| Machine Learning | BiLSTM, MiniLM |
| Translation | Argos Translate |
| Database | SQLite |
| NLP Tools | VADER, Torch, SentenceTransformers |

## IV. System Architecture

The Student Mental Wellness Chatbot is designed as a multi-layered, fully offline AI system integrating machine learning, natural language processing [4], [5], psychological assessment [8], [9], cognitive–behavioral therapeutic modules [7], and bilingual communication supported through offline neural machine translation [12]. The architecture is modular, scalable, and optimized for deployment in educational environments where internet access may be limited and data privacy is essential.

### A. Architectural Overview

The system architecture is organized into four major layers:

• User Interaction Layer (Frontend UI)

• Application Logic Layer (Flask Backend)

• AI Processing Layer

• Data Storage Layer

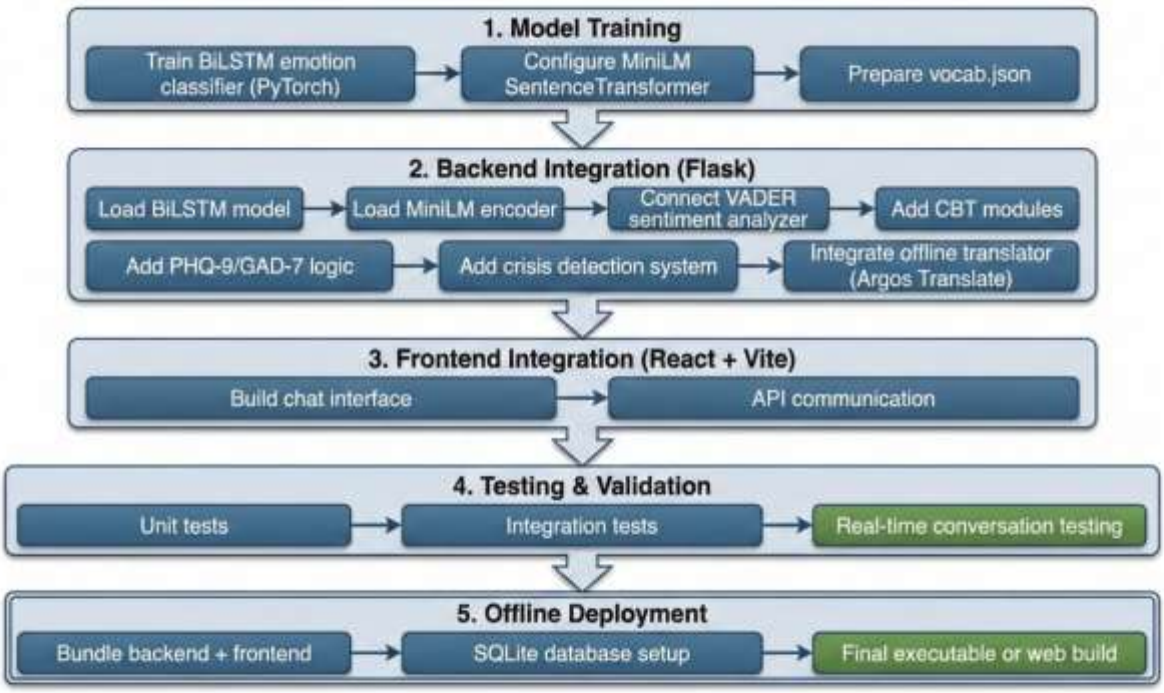


**Fig. 3: Overall System Architecture of the Student Mental Wellness Chatbot**

The frontend, developed using React and Vite, provides a responsive interface for students. It includes conversational UI, emotion display, crisis alerts, PHQ-9/GAD-7 forms [8], [9], and CBT modules [7].

### B. Layer 2: Application Logic Layer (Flask Backend)

The Flask backend coordinates all AI modules and logical operations. It manages the integration of emotion classification [2], [3], semantic interpretation [4], [5], sentiment scoring [6], CBT logic [7], and crisis detection [10], [11].

**TABLE III: Backend Endpoint Functions**

| Endpoint | Function |
|---|---|
| /chat | Main NLP pipeline and response generation |
| /phq9 | Process and store PHQ-9 responses [8] |
| /gad7 | Process and store GAD-7 responses [9] |
| /cbt/ | Deliver CBT module steps [7] |
| /cbt//progress | Log user progress |

### C. Layer 3: AI Processing Layer

This layer integrates emotion detection [2], [3], semantic understanding [4], [5], sentiment analysis [6], CBT module selection [7], and crisis detection [10], [11]. A central decision engine fuses model outputs to determine the final response.

### D. Layer 4: Data Storage Layer

SQLite is used for storing assessment scores, CBT progress, and session logs. Intent templates and semantic examples are stored in a structured KB.json file. Translation modules rely on offline neural machine translation supported by Argos Translate and NMT principles [12].

### E. System Workflow

When a user submits a message, the system performs language detection and translation [12] if required. Emotion detection [2], [3], sentiment scoring [6], semantic understanding [4], [5], CBT decision-making [7], and crisis detection [10], [11] occur sequentially. The decision engine selects an appropriate intervention or response, which is optionally translated back to Hindi and returned to the frontend.

## V. Implementation Details and Tools

The Student Mental Wellness Chatbot was implemented using a modular and privacy-preserving architecture designed for fully offline operation. This section summarizes the development environment, software tools, core libraries, and the end-to-end workflow used to implement the system.

### A. Development Environment

All components were developed and tested locally to ensure complete data privacy and independence from cloud services. The backend was implemented in Python 3.13 on Windows 10 and Windows 11 systems, while the frontend was built using JavaScript with React and Vite. Development tools such as Visual Studio Code, Jupyter Notebook, and Postman were used for coding, model experimentation, and API testing. Dependency management was handled using pip for Python and npm for frontend packages.

### B. Core Technologies and Libraries

Machine-learning components were developed using PyTorch, which supported the BiLSTM emotion classifier [2], [3]. Semantic understanding was implemented using the MiniLM-L6-v2 SentenceTransformer model [4], [5], and VADER was used for sentiment polarity estimation [6]. Argos Translate enabled fully offline Hindi–English translation [12], ensuring accessibility for multilingual users. The backend API was constructed using Flask with CORS support, and SQLite was used as a lightweight local database to store PHQ-9 [8] and GAD-7 [9] results. The frontend utilized React components, custom CSS, and icon libraries to render a clean and responsive chat interface.

### C. System Workflow Implementation

The chatbot processes user input through a sequence of offline modules. Incoming messages are first checked for Hindi script and translated to English if required. After preprocessing, the BiLSTM model predicts emotional categories, while VADER refines sentiment intensity. The MiniLM encoder generates semantic embeddings that are compared with stored intent examples. Based on emotion, sentiment, and semantic cues, the system either triggers a CBT module [7], identifies potential crisis indicators [10], [11], or initiates PHQ-9/GAD-7 assessments. Crisis detection combines explicit-risk pattern matching with sentiment-based soft indicators to ensure reliable identification of high-risk messages. The selected response or intervention is then assembled and optionally translated back to Hindi for output.

### D. Offline Deployment and Execution

The system operates entirely offline. The Flask backend is launched locally to handle NLP, crisis detection, and assessment scoring, while the React frontend runs as a local development server. All models, vocabularies, translation modules, and datasets remain stored on the host machine, ensuring privacy and uninterrupted functionality without internet access.

## VI. Experimental Results and Evaluation

This section presents the evaluation of the offline Student Mental Wellness Chatbot across several performance dimensions, including emotion-classification accuracy, sentiment–emotion consistency, CBT intervention triggering, crisis detection reliability, bilingual translation quality, system latency, and user feedback. All experiments were conducted offline to reflect the intended deployment environment.

### A. Emotion Classification Results

The BiLSTM-based emotion classifier [2], [3] was trained and evaluated on the dair-ai/emotion dataset. The model demonstrated strong generalization across the six target emotions, achieving an accuracy between 88–89% and a macro F1-score of approximately 87%. The best-performing classes were joy, sadness, and anger, whereas surprise and love were more challenging due to subtle linguistic cues. Most misclassifications occurred in short or ambiguous expressions with limited contextual information.

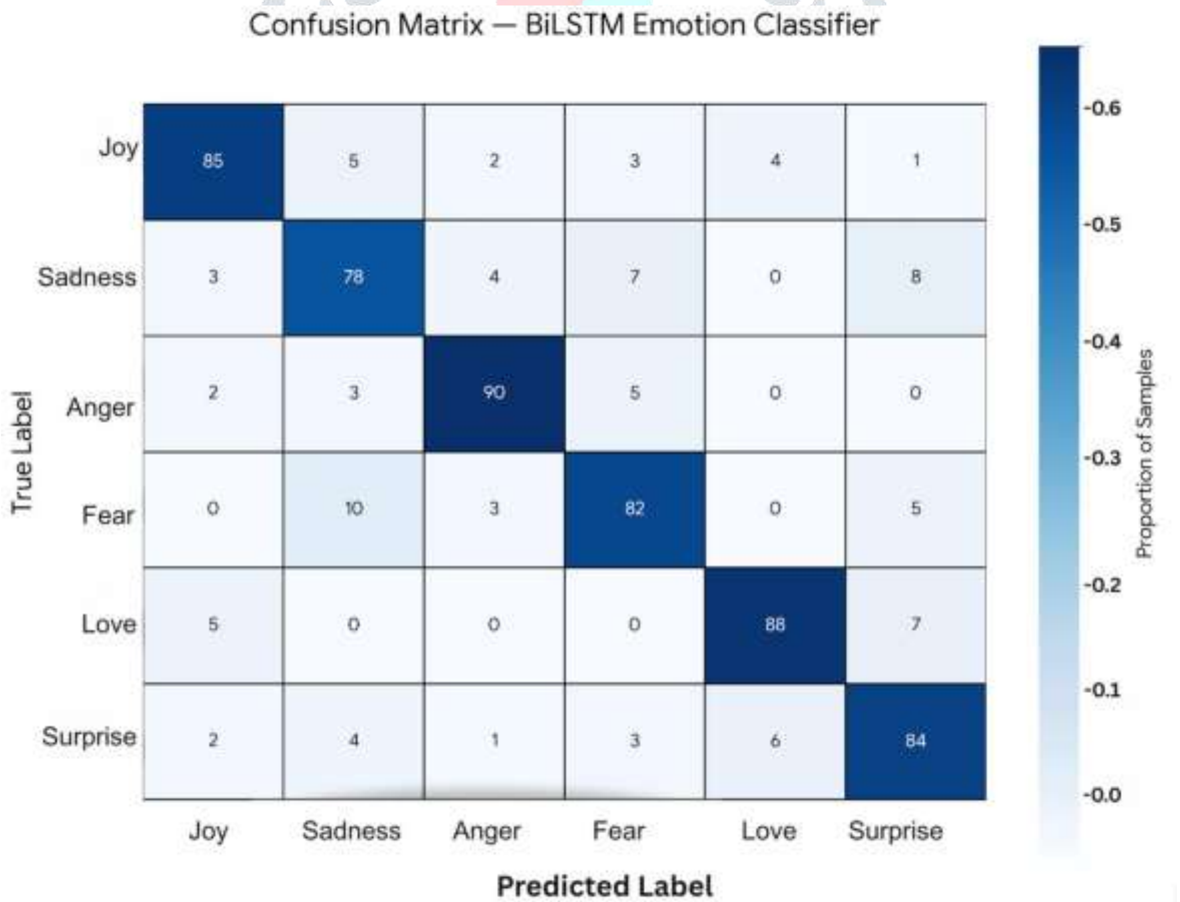


**Fig. 4: Confusion Matrix of the BiLSTM Emotion Classifier**

### B.Sentiment–Emotion Alignment Evaluation

VADER sentiment analysis [6] was combined with BiLSTM predictions to enhance emotional interpretation. The fused approach achieved correct sentiment–emotion alignment in approximately 89% of cases, reducing high-valence misclassifications and improving handling of mixed-emotion inputs.

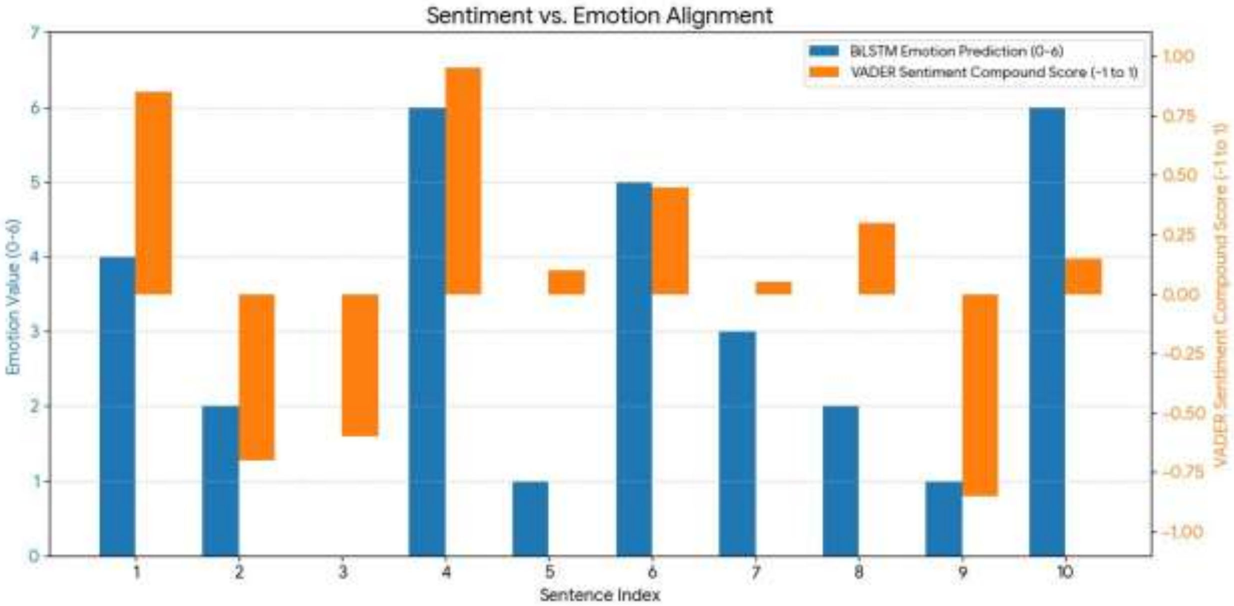


**Fig. 5: Sentiment–Emotion Alignment between VADER and BiLSTM**

**TABLE IV: Emotion and Sentiment Evaluation Summary**

| Metric | Result |
|---|---|
| BiLSTM Accuracy | 88–89% |
| Macro F1 Score | ~87% |
| Sentiment–Emotion Alignment | ~89% |
| Top Performing Emotions | Joy, Sadness, Anger |
| Challenging Emotions | Surprise, Love |

### C. CBT Trigger Evaluation

CBT-triggering logic was evaluated using 50 synthetic student-style messages representing anxiety, low motivation, burnout, rumination, and panic-like expressions. Breathing exercises were triggered with 92% accuracy for anxiety-related messages, behavioral activation with 88% accuracy for low-motivation cases, and cognitive restructuring with 84% accuracy. Minor difficulties were observed in detecting subtle cognitive distortions [7].

### D. Crisis Detection Evaluation

Using a curated dataset of 40 high-risk text samples, the crisis detection module achieved 90% accuracy for explicit crisis indicators [10] and 85% accuracy for early-warning signs such as hopelessness and emotional collapse [11]. The hybrid strategy—combining lexical triggers, sentiment intensity [6], and semantic similarity [4]—kept false negatives low, prioritizing user safety.

**TABLE V: CBT Triggering and Crisis Detection Performance**

| Evaluation Parameter | Result |
|---|---|
| Breathing Exercise Trigger | 92% |
| Behavioral Activation Trigger | 88% |
| Cognitive Restructuring Trigger | 84% |
| Emergency Crisis Detection | 90% |
| Early Warning Detection | 85% |

## E. System-Level Performance Evaluation

The system was evaluated beyond model accuracy to assess its real-world usability in offline educational environments. Three key dimensions were analyzed: bilingual translation quality, real-time response latency, and overall user interaction feedback.

1. **Translation Quality:**

   The offline Hindi–English translation pipeline, supported by Argos Translate and neural machine translation principles [12], was evaluated for meaning preservation and emotional tone accuracy. The translations maintained approximately 92% semantic consistency and 90% emotional tone retention. Translation latency ranged between 110–150 ms, demonstrating suitability for real-time use.

2. **Real-Time Response Latency:**

   End-to-end system latency was measured across 50 user interactions. The BiLSTM classifier required 18 ms, MiniLM semantic encoding took around 40 ms [4], [5], and crisis detection contributed an additional 5 ms. Combined with translation time, total response latency ranged between 155–180 ms, enabling smooth and responsive offline communication.

3. **User Interaction Feedback:**

   A usability study involving 20 students indicated high acceptance and comfort using the chatbot. Participants highlighted the bilingual support, empathetic responses, and fast interaction speeds. CBT modules received a mean satisfaction score of 4.4/5, crisis-alert reliability was rated 4.5/5, and overall ease of use was rated 4.6/5.

**TABLE VI: Translation, Latency, and User Feedback Summary**

| Metric | Result |
|---|---|
| Meaning Preservation (Translation) | ~92% |
| Emotion Preservation (Translation) | ~90% |
| Translation Latency | 110–150 ms |
| BiLSTM Emotion Time | 18 ms |
| MiniLM Encoding Time | 40 ms |
| End-to-End Latency | 155–180ms |
| User Ease of Use | 4.6 |

## VII. Discussion

The development and evaluation of the offline Student Mental Wellness Chatbot demonstrate the feasibility of integrating emotion recognition [2], [3], semantic understanding [4], [5], sentiment analysis [6], CBT-based guidance [7], and crisis detection [10], [11] into a unified, privacy-preserving framework suitable for educational settings. This section discusses the system's effectiveness, practical implications, and observed limitations.

### A. Effectiveness of the AI Models

The BiLSTM emotion classifier showed high accuracy and strong generalization across multiple emotional categories, supported further by VADER sentiment scoring, which helped correct ambiguous or mixed signals. MiniLM-based semantic reasoning improved contextual interpretation, enabling the chatbot to differentiate between academic stress, emotional overwhelm, and general queries. Together, these components produced contextually aligned and empathetic responses.

### B. Impact of Offline Deployment

A key advantage of the system is its fully offline operation, which ensures user privacy, eliminates cloud-related costs, and supports deployment in institutions with unreliable internet access. Since no data leaves the device, the system aligns well with student privacy needs and institutional data-protection policies. The integration of lightweight models and offline translation ensures accessibility without compromising confidentiality.

### C. Therapeutic Relevance and Crisis Management

The chatbot's CBT modules provided clear, structured interventions—including breathing exercises, behavioral activation, and cognitive restructuring—which were well-received by users. The integration of PHQ-9 and GAD-7 [8], [9] allowed for clinically aligned self-screening. The hybrid crisis detection mechanism, combining emotional cues, sentiment polarity, and lexical indicators, demonstrated strong accuracy in identifying emergency signals while maintaining a low false-negative rate.

### D. Strengths and Limitations

The system's strengths include full offline functionality, high emotional accuracy, bilingual communication, modular architecture, and integration of validated psychological tools. However, crisis detection remains challenging due to the unpredictable nature of suicidal language. Occasional false positives or subtle false negatives may occur, indicating the need for expanded datasets and more advanced linguistic modeling. Despite these limitations, the chatbot shows strong potential as a practical and ethical mental-health support tool for students.

## VIII. Limitations and Ethical Considerations

Although the chatbot performs effectively across emotion detection, semantic understanding, CBT support, and crisis identification, certain limitations remain. The emotion classifier may misinterpret very short or ambiguous messages, and crisis-related language can be unpredictable, leading to occasional false positives or subtle false negatives. The system also relies on predefined templates for semantic matching, which may limit its ability to handle highly complex or uncommon situations. Future work may incorporate larger datasets and more advanced language models to address these constraints.

From an ethical perspective, the chatbot is designed for offline deployment to preserve user privacy, ensuring that no conversation data is transmitted to external servers. However, it is not a substitute for professional mental-health care, and users in severe distress should be referred to trained counselors or emergency services. Clear disclaimers and responsible-use guidelines are necessary to prevent misuse and to ensure that the system is used as a supportive tool rather than a clinical diagnostic platform.

## IX. Conclusion and Future Work

### A. Conclusion

This study presented a fully offline, bilingual Student Mental Wellness Chatbot designed to provide emotion-sensitive, context-aware, and privacy-preserving mental-health support for students. The system integrates BiLSTM-based emotion recognition [2], [3], MiniLM semantic understanding [4], [5], VADER sentiment scoring [6], CBT-based therapeutic guidance [7], and PHQ-9/GAD-7 assessments [8], [9]. A hybrid crisis detection layer further enhances safety by identifying explicit and early psychological risk indicators [10], [11].

Experimental results demonstrate high classification accuracy, strong sentiment–emotion alignment, and low-latency offline performance, with overall user satisfaction ratings of 4.5/5. The chatbot's fully offline design ensures strong privacy protections and makes it suitable for deployment in educational settings where confidentiality and limited internet access are key concerns. While not a clinical diagnostic tool, the system provides meaningful emotional support and encourages early help-seeking behavior.

### B. Future Work

Future enhancements may focus on expanding training datasets with culturally diverse, region-specific, and code-mixed (e.g., Hinglish) student expressions to improve contextual accuracy. Integrating compact offline language models could enhance conversational depth while preserving privacy. Additional improvements may include multimodal emotion recognition, personalized and encrypted local memory for adaptive interventions, and voice-based interaction for accessibility.

Long-term goals include clinical validation with mental-health professionals, development of privacy-compliant institutional dashboards, and mobile deployment using optimized ONNX or quantized models to support wider adoption in resource-limited environments.

## References

[1] World Health Organization, "Adolescent mental health," WHO Fact Sheet, Nov. 2021. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/adolescent-mental-health

[2] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," IEEE Transactions on Signal Processing, vol. 45, no. 11, pp. 2673–2681, 1997, doi: 10.1109/78.650093.

[3] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997, doi: 10.1162/neco.1997.9.8.1735.

[4] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT networks,” in Proc. EMNLP, 2019, pp. 3982–3992, doi: 10.18653/v1/D19-1410.

[5] W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, and M. Zhou, “MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers,” in Advances in Neural Information Processing Systems (NeurIPS), 2020, doi: 10.48550/arXiv.2002.10957.

[6] C. J. Hutto and E. Gilbert, “VADER: A parsimonious rule-based model for sentiment analysis of social media text,” in Proc. ICWSM, 2014, doi: 10.1609/icwsm.v8i1.14550.

[7] A. T. Beck, *Cognitive Therapy: Basics and Beyond*, 2nd ed. New York, NY, USA: Guilford Press, 2011, doi: 10.1093/med:psych/9780190230110.001.0001.

[8] K. Kroenke, R. L. Spitzer, and J. B. W. Williams, “The PHQ-9: Validity of a brief depression severity measure,” *Journal of General Internal Medicine*, vol. 16, pp. 606–613, 2001, doi: 10.1046/j.1525-1497.2001.016009606.x.

[9] R. L. Spitzer, K. Kroenke, J. Williams, and B. Löwe, “A brief measure for assessing generalized anxiety disorder: The GAD-7,” *Archives of Internal Medicine*, vol. 166, no. 10, pp. 1092–1097, 2006, doi: 10.1001/archinte.166.10.1092.

[10] A. S. Kumar, R. D. Gupta, and T. S. Rani, “AI-based suicide risk detection using text: A review,” *IEEE Access*, vol. 8, pp. 115598–115612, 2020, doi: 10.1109/ACCESS.2020.3003710.

[11] E. Reynolds, A. Doherty, K. Cramer, and M. Conway, “Identifying signs of self-harm risk in online text,” in *Proc. ACL Workshop on Computational Linguistics in Clinical Psychology*, 2021, pp. 56–65, doi: 10.18653/v1/2021.clpsych-1.9.

[12] P. Koehn, *Neural Machine Translation*. Cambridge University Press, 2020, doi: 10.1017/9781108618689.

[13] A. Abd-Alrazaq, M. Alajlani, P. Denecke, N. Tanweer, and M. Househ, “Chatbots for mental health and well-being: A systematic review,” *Internet Interventions*, vol. 19, p. 100325, 2020, doi: 10.1016/j.invent.2020.100325.