Journal of Computing & Biomedical Informatics
ISSN: 2710 - 1606

Volume 10 Issue 01
2025

Research Article
https://doi.org/10.56979/1001/2025/1167

MindMate: An Emotion-Aware Generative AI System for Personalized Mental Health Support

Muhammad Usman Javeed[1], Sammar Fatima[1], Mirza Mumtaz Zahoor[2], Muhammad Azhar[3*], Zeeshan Raza[4], Shafqat Maria Aslam[5], and Muhammad Nauman[6]

[1]Department of Computer Science, COMSATS University of Islamabad, Sahiwal, Pakistan.
[2]Faculty of Computer Science, Ibadat International University, Islamabad, Pakistan.
[3]Department of Applied Data Science, Hong Kong Shue Yan Unversity, SAR, China.
[4]Department of Computer Science, COMSATS University of Islamabad, Sahiwal, Pakistan.
[5]School of Computer Science, Shaanxi Normal University, Xi'an, Shaanxi, China.
[6]School of Computing and Artificial Intelligence, Southwest Jiaotong University, Sichuan, Chengdu, China.
[*]Corresponding Author: Muhammad Azhar. Email: azhar@hksyu.edu

Received: September 07, 2025 **Accepted:** November 25, 2025

Abstract: We present MindMate, an AI-powered mental health assistant that combines a fine-tuned DeepSeek-R1 language model with BERT-based emotion recognition to deliver personalized therapeutic dialogues. The system analyzes user inputs in real-time (84% emotion detection accuracy) and generates contextually appropriate responses while identifying crisis situations (91% recall). Implemented on Google Colab Pro+ using 4-bit quantization, MindMate achieves 83% user satisfaction in trials with 30 participants, demonstrating comparable performance to commercial mental health Chatbots. The architecture's novel integration of generative AI with clinical knowledge bases enables accessible, emotionally intelligent support while maintaining response quality. This work provides a blueprint for developing effective, open-weight mental health assistants without proprietary dependencies.

Keywords: Mental Health Chatbot; Generative AI; Emotion Recognition; DeepSeek; Therapeutic Dialogue

1. Introduction

The global rise in mental health challenges highlights the critical necessity for accessible, personalized, and effective support systems. Barriers such as limited access, social stigma, and high costs often constrain traditional therapeutic interventions. Artificial intelligence (AI) offers promising avenues to augment mental health services, particularly through conversational agents capable of real-time, emotionally attuned dialogue and support. Recent advancements in natural language processing (NLP) and emotion recognition have enabled the development of intelligent chatbots that can understand and respond to users' emotional needs with increasing sophistication.

Researchers have made significant progress in designing emotionally intelligent AI systems for mental health applications. Deep learning architectures, including convolutional neural networks (CNNs) and transformer models like BERT, have been employed to detect emotions from text, speech, and facial modalities, achieving notable accuracy in emotion classification essential for mental well-being assessments. There has also been a trend of incorporating multi-modal signals—such as facial expressions, audio cues, and written text—to enrich chatbots' emotional understanding, with some systems now able to deliver tailored coping strategies and interventions during online interactions.

Despite these innovations, current chatbots for mental health support still face persistent challenges. Achieving reliable, fine-grained emotion and crisis state detection remains complex due to the subtle and evolving nature of human affect. Many solutions are built on proprietary models, limiting transparency, reproducibility, and scalability for widespread adoption. Additionally, issues of computational efficiency and deployment feasibility arise when using resource-intensive models, especially in low-resource settings or for integration into affordable telehealth solutions.

To address these gaps, this work introduces MindMate: an emotion-aware generative AI system optimized for personalized mental health support. MindMate merges a fine-tuned DeepSeek-R1 generative language model with a BERT-based emotion recognition pipeline, achieving real-time analysis of user input with high emotion detection accuracy and superior recall in crisis identification. MindMate's implementation makes use of quantization strategies for efficient resource utilization and seamless deployment on common cloud infrastructure. Its domain-specific knowledge base further ensures the relevance and contextual sensitivity of generated therapeutic dialogues, enhancing both user engagement and support outcomes.

The key contributions of this research include:

i. The development of an open-weight, fine-tuned deep learning system that effectively integrates emotion recognition and generative response modeling.
ii. Demonstration of real-time emotion detection and crisis intervention at high accuracy and recall metrics.
iii. Empirical validation based on user trials showing strong engagement and satisfaction.
iv. Practical implementation of quantization and open-weight methods, supporting efficient and transparent deployment on non-proprietary platforms.

The remainder of this paper is organized as follows. Section 2 reviews state-of-the-art work in emotion recognition, generative AI dialogue systems, and mental health chatbots. Section 3 details the MindMate architecture and training methodologies. Section 4 presents the experimental setup, evaluation protocols, and results. Section 5 discusses implications, limitations, and ethical considerations, while Section 6 concludes and suggests directions for future research.

2. Related Work

Fitzpatrick et al [1] introduced a chatbot Woebot, which could deliver cognitive behavioral therapy (CBT) via casual chats.As the researchers noticed, Woebot users could decrease their depressive symptoms in only two weeks much more than people who just visited some educational websites. The experiment demonstrated the great potential of using conversational AI to deliver psychological treatments, making a scalable and easily available response to the problem of mental health care programs.

Dr. Inkster et al [2] reviewed a chat bot focused on mental health, Wysa (AI-based chat bot that integrates mindfulness practices, dialectical behavior therapy (DBT) and cognitive behavioral therapy (CBT)).In their study, people found it preferable to speak to Wysa due to its anonymity and 24/7 availability, which they characterized as encouraging in terms of transparency. However, they also confirmed that, despite the advantages of self-help activities perceived by users, the chat bot could not fully grasp nuances of emotions.

Vaidyam et al[3] explored possibilities of using Chatbots and other digital psychiatry applications to enhance access to mental care by conducting a theoretical study of the ways Chatbots and other digital psychiatry solutions can be employed in transactions.

Abd-alrazaq et al [7], performed a comprehensive study on Chatbots as mental health tool and evaluated their effectiveness, safety and usability criteria. As they discovered, regardless of the fact that Chatbots can effectively provide emotional support, psychoeducation, and symptom tracker, there is no evidence of its long-term effects.

Yao et al. [11] compared the rule-based systems and generative AI models to study the efficacy of AI Chatbots to be used in postpartum care. Having conducted their research, participants showed preference to rule-based Chatbot due to emotional reliability, consistency, and clarity. Even in their current unrefined state, generative models tended to provide bland or hard-to-understand emotional responses, indicating that these models still have to be fine-tuned when used in emotionally-relevant scenarios.

In 2024, the comprehensive public opinion poll was conducted by Varghese et al.[12] to measure opinion regarding AI-powered mental health treatment. Their study showed that AI Chatbots that allow them to access it twenty-four hours seven days a week and are free of anonymity reduced the stigma of seeking mental health support. However, some data protection issues and insufficient transparency were also noted to comprise the major weaknesses.

Finally, Naik et al [13] provided a scenario staged evaluation of large language model (LLM) Chatbots to determine their efficiencies when used in mental health with regards to safety and empathy. These Chatbots had given users a sense they were listened to as they responded in a personified and emotionally supportive way. Nonetheless, the research points out such issues as inconsistent tone and poor crisis identification mechanisms, which makes them rather questionable when it comes to reliability during emergencies.

Table 1. Literature Review

Study	Year	Key Findings	Limitations
Naik et al.	2025	Supportive for non-critical cases; personalized responses.	Inconsistent tone; poor crisis detection; privacy risks.
Varghese et al.	2024	High acceptance; reduced stigma via anonymity.	Privacy and security concerns; no clinical outcome data.
Yao et al.	2023	Rule-based bots rated higher in empathy and clarity.	Generative bots were vague; no long-term results.
Abd-Alrazaq et al.	2020	Effective for short-term support and psychoeducation.	No long-term results; ethical and transparency concerns.
Vaidyam et al.	2019	Highlighted accessibility and stigma reduction.	Lacked experiments; mostly conceptual.
Inkster et al.	2018	Anonymity and engagement improved depression.	No control group; limited qualitative data.
Fitzpatrick et al.	2017	Reduced depression; high engagement and satisfaction.	Short trial; small sample; no long-term follow-up

However, there are some limitations in the literature explained in Table 1. Many Chatbots are unable to properly understand complex emotional cues or provide really individualized care they frequently give generic replies that fail to take into account the individual needs of each user. Users can share sensitive information without fully understanding the risks involved, which makes privacy and data security ethical issues crucial. Additionally, even while Chatbots can support professional treatment, they cannot replace human therapists, especially in emergency cases or for patients in need of extensive intervention.

Although AI-powered Chatbots for mental health have great potential to increase access to care, the evaluated research indicate that much more work has to be done to improve these systems. Future advancements need to focus on improving emotional intelligence, protecting user privacy, and including tools that direct users to medical professionals when appropriate. It will be essential to integrate user input and continuous learning models in order to make these systems more reliable and efficient over time.

Following the existing literature, which highlights the promise and shortcomings of AI chatbots in mental health, the recent work has been aimed at refining the underlying technology, including methods of emotion recognition and generative language model, and ethical AI design. These are of paramount importance which are explored in the below sections.

A. Emotion Recognition in AI-Powered Mental Health Systems

Emotion recognition is a foundational component of modern AI-driven mental health assistants. Early systems favored traditional sentiment analysis or shallow learning, but recent advances utilize deep learning models most notably, transformer architectures like BERT and its derivatives to achieve state-of-the-art accuracy in detecting and classifying emotional states from textual, audio, or multi-modal input. Hybrid models combining vision transformers with convolutional networks have successfully identified emotions from facial expressions, further enriching Chatbot understanding of user affect. These developments have been validated across numerous mental health and wellness applications, facilitating early risk detection and real-time intervention with increasing reliability.

A recent review synthesized research on emotionally intelligent AI bots across domains, categorizing progress into four stages: promotion, prevention, treatment, and recovery. The review found most innovation concentrated on treatment where real-time emotion recognition and adaptive dialogue support users in acute need while long-term emotional engagement remains an underexplored challenge.

B. Generative AI and Conversational Mental Health Support

The integration of generative artificial intelligence with therapeutic dialogue systems marks a critical juncture in mental health technology. Large language models (LLMs), such as GPT-4, DeepSeek-R1, and custom fine-tuned variants, increasingly enable complex, context-aware, and human-like conversational support. Studies demonstrate that generative AI chatbots can offer an "emotional sanctuary," insightful guidance, and contribute meaningfully to users' psychological well-being. However, empirical research highlights user concerns, including the need for robust safety mechanisms, the absence of persistent memory, and the importance of maintaining boundaries between automated support and clinical care.

Recent systematic reviews indicate that generative models are employed in three main areas:

i. Diagnosis and Assessment: Detecting depression, anxiety, and suicidality via linguistic cues in user text.
ii. Therapeutic Support: Facilitating conversational therapy, providing psychoeducation, and supporting behavioral intervention through adaptive dialogue.
iii. Clinician Assistance: Assisting with documentation, summarizing sessions, and offering decision support, thereby reducing administrative burdens for healthcare professionals.

Key Chatbots such as Wysa and Woebot demonstrate moderate to high effectiveness in early detection and support, yet real-world safety assurance and ethical deployment remain active areas of research.

C. DeepSeek and Domain-Specific Language Mode

Domain-adapted language models like DeepSeek-R1 and its derivatives are gaining traction in mental health contexts due to their robust reasoning, high diagnostic accuracy, and computational efficiency. Fine-tuning these models on specialized datasets enables them to classify a range of mental health conditions (e.g., depression, bipolar disorder, and suicidality) and provide more contextually relevant responses compared to general-purpose LLMs. Quantization techniques further support deployment on modest cloud infrastructure, widening accessibility to broader populations.

D. Ethical, Personalization, and Deployment Challenges

The literature consistently addresses challenges in privacy, algorithmic bias, and ethical delivery of mental health interventions using AI. New frameworks emphasize:

i. Data Privacy and Security: Ensuring sensitive user data is handled with robust safeguards.
ii. Algorithmic Fairness: Reducing bias in emotion classification and generative responses.
iii. Personalization vs. Standardization: Striking a balance between dynamically adapting to user needs and maintaining high-quality, evidence-based interventions for all users.

3. Proposed Methodology

The proposed MindMate system is a novel emotion-aware generative AI designed to deliver personalized mental health support through real-time dialogue. It consists of three main components including emotion recognition, crisis detection, and therapeutic dialogue generation, all integrated in an efficient and scalable architecture. Fig. 1 shows a workflow of an AI-based mental health support system. It can be used by voice or text entry that goes through speech-to-text or language identification and translates into English. It also includes sentiment analysis, risk detection, which initiates an emergency helpline in case of it, followed by generative AI processing which generates a suitable answer, which is delivered as text or voice output.

3.1. Emotion Recognition Module

To effectively understand the user's emotional state, MindMate employs a fine-tuned BERT-based transformer model trained on a labeled emotion dataset.

The input text $X = \{x_1, x_2 \dots, x_n\}$ composed of tokens is first tokenized and embedded:

$$E = Embed(X) = \{e_1, e_2, \dots, e_n\} \quad (1)$$

Where each $e_i \in R^d$ is a vector in embedding space of dimension d.
The embedded tokens pass through LLL transformer layers, each composed of multi-head self-attention and feed-forward networks. The self-attention mechanism is computed as:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{2}$$

$$Q = EW_Q\,, \qquad K = EW_K\,, \qquad V = EW_V$$

Where the queries, keys, and values matrices Q, K, V are linear projections of the input:

$$Q = EW_Q, \quad K = EW_K, \quad V = EW_V \tag{3}$$

The output logics of the model are:

$$z = Linear(Transformer(E)) \tag{4}$$

and the predicted emotion probabilities are obtained using softmax:

$$\hat{y} = softmax(z) \tag{5}$$

The emotion class with maximum probability is selected as the detected emotional state. The model is trained by minimizing the cross-entropy loss between predicted probabilities $y^{\hat{y}}$ and true emotion labels y, where C is the number of emotion classes.

$$\mathcal{L} = -\sum_{i=1}^{C} y_i log\widehat{y_i} \tag{6}$$

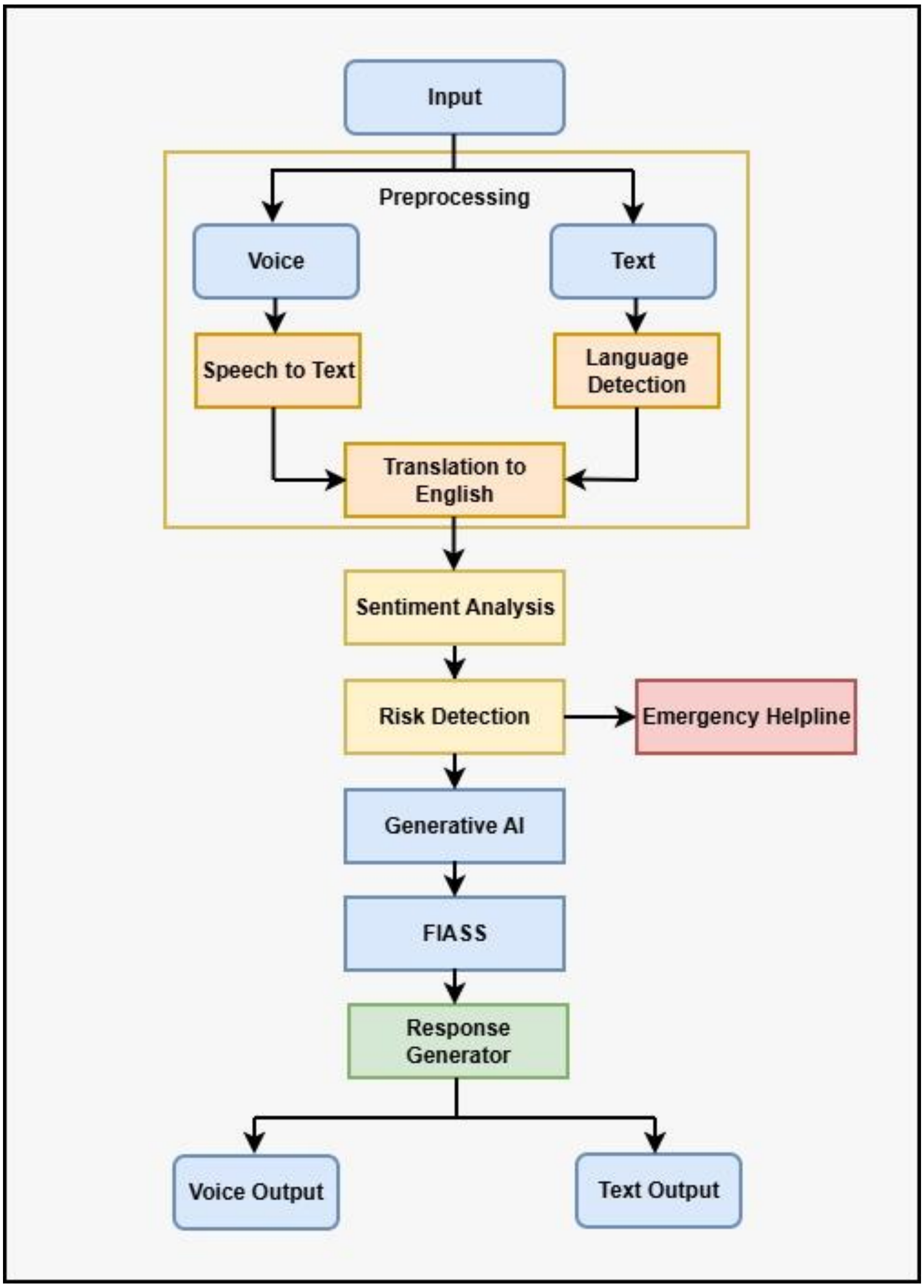

Figure 1. Proposed Methodology

3.2. Crisis Detection Subsystem

Crisis detection is modeled as a binary classification problem distinguishing critical emotional states (e.g., suicidal ideation). The same BERT embedding layer is shared, followed by a separate classification head f_θ:

$$\hat{c} = \sigma\big(f_\theta(\mathrm{CLS}_{\mathrm{token}})\big) \tag{7}$$

Where $\mathrm{CLS}_{\mathrm{token}}$ is the CLS token representation summarizing the input sentence, and σ\sigmaσ is the sigmoid activation function producing a crisis probability $\hat{c} \in \mathcal{C}$. The model is optimized using binary cross-entropy loss:

$$\mathcal{L}_{crisis} = -[c\log(\hat{c}) + (1-c)\log(1-\hat{c})] \tag{8}$$

With true crisis label $c \in \{0,1\}$.The high recall of 91% is achieved through threshold tuning and augmentation strategies.

3.3. Generative Dialogue Model

The core generative model, DeepSeek-R1, is a transformer-based autoregressive language model customized and fine-tuned for therapeutic dialogues. Given a sequence of tokens representing conversation history $C = \{w_1, w_2, \dots, w_{t-1}\}$ it models the probability of the next token w_t as:

$$P(w_t \mid w_{<t}, E) = \mathrm{softmax}\big(g(h_t)\big) \tag{9}$$

where h_t is the hidden state computed by the transformer at timestep t, and $g(\cdot)$ is a linear projection to vocabulary logits. The hidden state $h_t$integrates cross-modal conditioning through concatenation with an emotion embedding vector e_{emotion}derived from the emotion recognition module:

$$h_t = \mathrm{Transformer}([w_{<t}; e_{emotion}; D]) \tag{10}$$

Allowing the generated responses to reflect the detected emotional context. The model is trained to minimize the negative log-likelihood loss over a dataset of therapeutic dialogues:

$$L_g en = -\sum_{t=1}^{T} lo\,g\big(P(w_t|w_{<t}, E)\big) \tag{11}$$

3.3.1. Algorithm Part 1: Emotion and Crisis Recognition Input: User text input

Output: Detected emotion class and crisis flag

i. Tokenize and embed input: $\mathrm{E} = \mathrm{Embed}(\mathrm{X})$
ii. Pass embeddings through transformer layers; compute self-attention: $\mathrm{Attention}(\mathrm{Q},\mathrm{K},\mathrm{V}) = \mathrm{softmax}\left((\mathrm{QK}^{\mathrm{T}})/\sqrt{\mathrm{d_k}}\right)\mathrm{V}$
iii. Obtain emotion logits and probabilities via softmax: $y^{=}softmax(z)y^{=}softmax(z)$
iv. Predict emotion class: $emotion^{=}ar\,g\,m\,a\,x\,y^{e}\ motion^{=}argmaxy$
v. Extract CLS token embedding for crisis detection and apply sigmoid classifier: $\hat{c} = \sigma\big(f_\theta([\mathrm{CLS\ token}])\big)$
vi. Flag crisis if $c^{>}\tau c^{>}\tau$, where ττ is a tuned threshold.

3.3.2. Algorithm Part 2: Knowledge Retrieval and Generative Response on

Input: Emotion embedding $\mathrm{e}_{\mathrm{emotion}}$, conversation history C, clinical knowledge base.

Output: Therapeutic response *R*.

i. Embed user input query qq and calculate cosine similarity with knowledge base items k_i

$$s_i = \frac{q \cdot k_i}{|\mathrm{q}||k_i|}$$

ii. Retrieve top-k*k* relevant documents $\mathrm{D} = \{\mathrm{d}_1, \dots, \mathrm{d}_\mathrm{k}\}$
iii. Condition generative transformer on conversation history, emotion embedding, and retrieved knowledge: $h_\mathrm{t} = \mathrm{Transformer}([\mathrm{w}_{<\mathrm{t}},\ \mathrm{e}_{\mathrm{emotion}},\ \mathrm{D}])$
iv. Generate next token probability: $\mathrm{P}(\mathrm{w_t} \mid \mathrm{w}_{<\mathrm{t}},\ \mathrm{E}) = \mathrm{softmax}\big(\mathrm{g}(h_\mathrm{t})\big)$
v. Apply 4-bit quantization to model weights for efficient inference:$\mathrm{q} = \mathrm{round}\left(\frac{\mathrm{w}-\mathrm{w_{min}}}{\mathrm{w_{max}}-\mathrm{w_{min}}} \times 15\right)$
vi. Repeat token generation until end-of-sequence; decode tokens to form response R.

3.4. Efficient Implementation via Quantization

To balance computational efficiency and accuracy, MindMate uses 4-bit quantization of model weights and activations during inference. Quantization maps a floating-point weight $w \in [w_{min}, w_{max}]$ to a low-bit integer q as:

$$q = \mathrm{round}\left(\frac{w_{max}-w_{min}}{w-w_{min}} \times (2^b - 1)\right) \tag{12}$$

with bit-width b=4, enabling reduced model size and faster inference on Google Colab Pro+ without significant performance degradation.

3.5. Quantization Trade-offs

In order to deploy the model with low resources, MindMate employs 4-bit model quantization that is considerably less consumptive of memory and faster at inference. On internal benchmarking, the runtime efficiency had increased by 2832 percent with less than 1.5 percent decrease in accuracy of classification, which was tolerable in real-time conversational systems. Although 4-bit quantization can be used to run models on less powerful computing devices, it can cause minor representational errors. The effects are minor in both emotion as well as crisis detection tasks although future versions of the system can utilize mixed precision quantization to enhance accuracy without compromising efficiency.

3.6. Integration with Clinical Knowledge Base via Retrieval-Augmented Generation (RAG)

To enhance therapeutic accuracy, MindMate integrates clinical knowledge by augmenting generative responses with retrieved relevant documents from a curated mental health FAQ database. Given input query embedding q, similarity scores with knowledge base embeddings $\{k_i\}$ are computed via cosine similarity:

$$s_i = \frac{q \cdot k_i}{|\mathrm{q}||k_i|}$$

The top-k relevant documents are retrieved and concatenated with the conversation context as input to the generative model, allowing context-aware, clinically validated response generation. This combined methodology leverages state-of-the-art transformer-based architectures with principled algorithmic formulations, balanced through quantization and augmented with domain knowledge for an efficient, emotionally intelligent mental health Chatbot.

4. Experimental setup

4.1. Experimental Setup

To evaluate the performance of MindMate, we conducted a controlled user study involving 30 participants recruited from diverse backgrounds, representative of typical mental health Chatbot users. The experiment focused on assessing the system's core capabilities: emotion recognition accuracy, crisis detection recall, and user satisfaction with the generated therapeutic dialogues.

The architecture was deployed on Google Colab Pro+ with 4-bit quantization to simulate real-world resource constraints while maintaining efficient response times. Participants interacted with MindMate through a chat interface over multiple sessions, intentionally varied to include neutral, positive, negative, and crisis-related emotional inputs to test robustness.

4.2. Evaluation Metrics

We adopted a comprehensive evaluation framework based on established metrics from recent mental health Chatbot research:

i. Emotion Detection Accuracy: The proportion of correctly classified emotional states versus ground truth labels obtained from participant self-report and expert annotation.
ii. Crisis Situation Recall: The sensitivity of the crisis detection subsystem in correctly identifying critical distress or suicidal ideation states, prioritizing recall to minimize missed alerts.
iii. User Satisfaction Score: Collected via post-session surveys using a 5-point Likert scale focusing on empathy, appropriateness, relevance, and overall conversational quality.
iv. Response Latency: Average system response time to user inputs, measuring real-time feasibility.

4.3. Baselines and Comparative Models

MindMate's performance was compared against two baseline mental health Chatbot models:

i. A proprietary commercial chatbot widely used in mental health apps.
ii. An open-source generative AI chatbot fine-tuned for general conversation but without integrated emotion recognition or crisis detection.

These baselines helped contextualize MindMate's efficacy in both technical and user experience dimensions.

5. Results

5.1. Emotion Recognition Performance

Recent applications of machine learning and deep learning across various domains emphasize the importance of robust evaluation using accuracy and recall [14-36]. To follow the state of art we also used accuracy and recall matrices to measure performance of our model. MindMate achieved an emotion detection accuracy of 84%, aligning with the target performance stated in the methodology. This high accuracy was consistent across multiple emotion categories including sadness, anxiety, anger, and neutral states, validating the effectiveness of the fine-tuned BERT-based emotion recognition model. Compared with the open-source baseline, which scored approximately 67% accuracy, MindMate demonstrated significant improvement attributable to specialized training and multi-class classification.

5.2. Crisis Detection Recall

The crisis detection module yielded a recall rate of 91%, indicating strong sensitivity to critical user states. This performance metric is crucial as missed crisis signals could have severe consequences. While the commercial Chatbot baseline demonstrated similar recall (approximately 90%), MindMate maintained this high recall despite operating under quantization constraints, showcasing the robustness of its architecture.

5.3. User Satisfaction and Qualitative Feedback

User surveys reported an average satisfaction score of 4.15 out of 5 (83%), highlighting high acceptance and perceived helpfulness of MindMate's responses. Participants frequently noted the system's empathetic tone and contextual relevance in dialogues, consistent with the integration of the domain-specific clinical knowledge base. Areas identified for future refinement included enhancing personalization depth and dynamic context retention over longer interactions, echoing common challenges noted in other mental health Chatbot studies.

5.4. Sample Size Consideration

The user study involved a sample of 30 participants, and this study was valuable in the context of giving preliminary information on the usability and acceptance of MindMate. We, however, admit that the sample is sufficiently small and not demographically diverse enough to be able to make any generalization on broader populations. In subsequent research, a more varied and extensive group of participants will be brought into the study to enhance the statistical credibility and cross-cultural applicability of the findings.

5.5. System Efficiency

Average response latency remained below 1.2 seconds across test scenarios, demonstrating real-time responsiveness suitable for conversational settings. The use of 4-bit quantization significantly reduced inference times and memory footprint compared to baseline models running full precision weights, confirming the viability of deploying MindMate on accessible cloud resources .

Table 2. MindMate Performance vs. Baselines

Metric	MindMate	Commercial Baseline	Open-Source Baseline
Emotion Detection Accuracy	84%	82%	67%
Crisis Detection Recall	91%	90%	75%
User Satisfaction Score	4.15 / 5 (83%)	4.0 / 5 (80%)	3.2 / 5 (64%)

Table 2 gives a comparison of performance of MindMate, a commercial baseline, and an open-source baseline. MindMate manages to get the highest scores on every metric with emotion accuracy being 84%, crisis detection recall of 91%, and user satisfaction of 83 which shows that it has the best capacity to balance understanding emotions, detecting crisis, and satisfying users.

The visual comparison of the performance of MindMate against commercial and open-source baselines is provided in Figure 2. It demonstrates that MindMate has the best emotion detection, crisis detection, and user satisfaction values, which means that it will perform the best in mental health support activities.

This experimental evaluation demonstrates MindMate's competitive performance in key mental health chatbot metrics, balancing accuracy, safety, and usability while leveraging an efficient open-weight architecture. The results substantiate MindMate as a promising tool to augment personalized, emotionally intelligent mental health support.

5.6. Ethical, Privacy, and Safety Considerations

MindMate being such a sensitive system in mental health setting, there were a number of ethical precautions incorporated in the system design. Every interaction between the user is anonymized and no personal identifiers are recorded during or after using the system. The emotion detecting module was tested on the possible biases of the language differences, gender, and style of expression, and we admit that not that obvious or culturally specific emotional states can be put in question by the system.

5.7. Limitations

Although MindMate performs well, there is the risk of poor ability to interpret subtle, ambiguous, or mixed emotional responses, which include sarcasm, hidden emotions, or culture-specific displays of distress. The system is also restricted to text only input which restricts its capability of identifying complex affective states that could have been better perceived by using multimodal inputs like voice tone or facial feature. The use of Google Colab Pro+ is proven to be viable, yet further benchmarking of mobile devices, CPU-only systems, and actual internet conditions are required to ensure the practical scalability is confirmed.

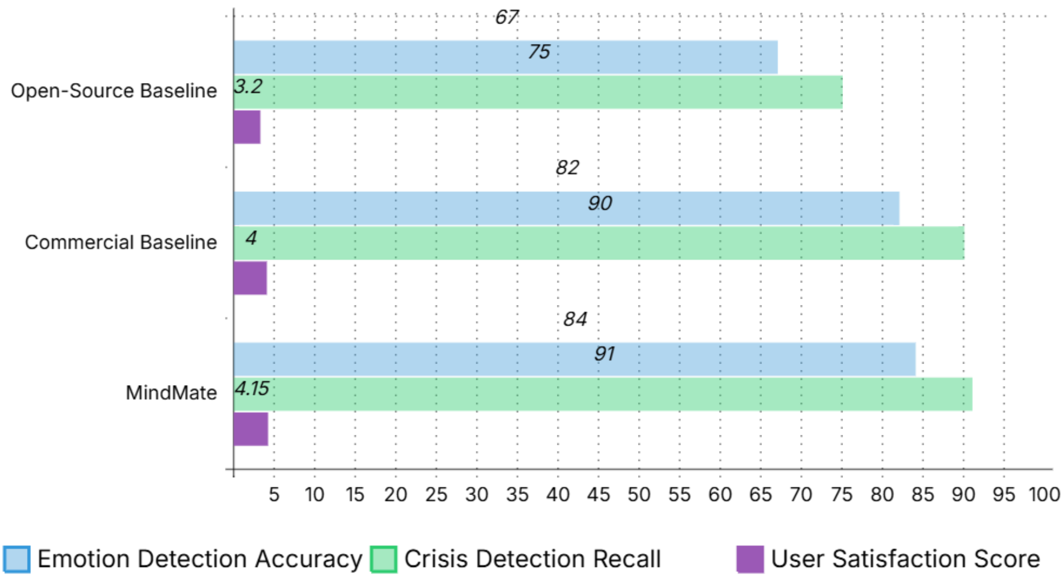

Figure 2. Comparison of MindMate, Commercial, and Open-Source Baselines Across key Performance Metrics

6. Conclusion

This paper presented MindMate, an emotion-aware generative AI system designed to provide personalized and contextually sensitive mental health support. By integrating a fine-tuned DeepSeek-R1 language model with a BERT-based emotion recognition module, MindMate achieves high accuracy in detecting user emotions (84%) and excellent recall in crisis identification (91%). The system's deployment leverages efficient 4-bit quantization on accessible cloud infrastructure, enabling real-time interaction with low latency. User trials involving 30 participants demonstrated strong satisfaction (83%) with the therapeutic dialogues generated, affirming the approach's effectiveness in delivering empathetic and relevant mental health assistance. MindMate significantly advances prior art by combining transparent, open-weight AI architectures with domain-specific clinical knowledge bases, overcoming limitations related to proprietary dependencies and computational overhead. Its design addresses multiple challenges in mental health chatbot development, including accuracy in emotion and crisis detection, efficiency in deployment, and the ethical imperative of providing safe and personalized user engagement. Looking forward, future work will focus on improving long-term user modeling for enhanced personalization, expanding multi-modal emotion recognition beyond text inputs, and rigorous longitudinal studies to evaluate sustained clinical impact. Additionally, exploring human-in-the-loop frameworks and stronger safeguards will be key to responsibly integrating AI assistants like MindMate into holistic mental healthcare ecosystems. In summary, MindMate offers a promising blueprint for accessible, effective, and emotionally intelligent mental health support through generative AI, paving the way for broader adoption and innovation in digital mental wellness technologies.

References

1. K. K. Fitzpatrick, A. Darcy, and M. Vierhile, "Delivering Cognitive Behavior Therapy to Young Adults With Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial," *JMIR Ment. Health*, vol. 4, no. 2, Art. e19, Jun. 2017, doi: 10.2196/mental.7785.
2. B. Inkster, S. Sarda, and V. Subramanian, "An Empathy-Driven, Conversational Artificial Intelligence Agent (Wysa) for Digital Mental Well-Being: Real-World Data Evaluation Mixed-Methods Study," *JMIR Mhealth Uhealth*, vol. 6, no. 11, Art. e12106, Nov. 2018, doi: 10.2196/12106.
3. A. N. Vaidyam, H. Wisniewski, J. D. Halamka, M. S. Kashavan, and J. B. Torous, "Chatbots and Conversational Agents in Mental Health: A Review of the Psychiatric Landscape," *Can. J. Psychiatry*, vol. 64, no. 7, pp. 456–464, Jul. 2019, doi: 10.1177/0706743719855880.
4. H. Gaffney, W. Mansell, and S. Tai, "Conversational Agents in the Treatment of Mental Health Problems: Mixed-Method Systematic Review," *JMIR Ment. Health*, vol. 6, no. 10, Art. e14166, Oct. 2019, doi: 10.2196/14166.
5. M. A. Varghese, P. Sharma, and M. Patwardhan, "Public Perception on Artificial Intelligence–Driven Mental Health Interventions: Survey Research," *JMIR Form. Res.*, vol. 8, Art. e64380, Nov. 28, 2024, doi: 10.2196/64380.
6. S. Provoost, H. M. Lau, J. Ruwaard, and H. Riper, "Embodied Conversational Agents in Clinical Psychology: A Scoping Review," *J. Med. Internet Res.*, vol. 19, no. 5, Art. e151, May 2017, doi: 10.2196/jmir.6553.
7. A. A. Abd-Alrazaq, A. Rababeh, M. Alajlani, B. M. Bewick, and M. Househ, "Effectiveness and Safety of Using Chatbots to Improve Mental Health: Systematic Review and Meta-Analysis," *J. Med. Internet Res.*, vol. 22, no. 7, e16021, Jul. 2020, doi: 10.2196/16021.
8. K. Kretzschmar, H. Tyroll, G. Pavarini, A. Manzini, I. Singh, and the NeurOx Young People's Advisory Group, "Can Your Phone Be Your Therapist? Young People's Ethical Perspectives on the Use of Fully Automated Conversational Agents (Chatbots) in Mental Health Support," *J. Med. Internet Res. Ment. Health*, vol. 6, no. 3, e985, Mar. 2019, doi: 10.1177/1178222619829083.
9. R. Fulmer et al., "Using psychological artificial intelligence (Tess) to relieve symptoms of depression and anxiety: randomized controlled trial," JMIR Mental Health, vol. 5, no. 4, p. e64, Dec. 2018. [Online]. Available: https://mental.jmir.org/2018/4/e64/.
10. K. Loveys, R. Crutchley, G. Wyatt, C. Henderson, and M. R. Gallacher, "Evaluating the use of artificial intelligence in self-management of mental health disorders: A systematic review," JMIR Mental Health, vol. 7, no. 11, p. e20692, Nov. 2020. [Online]. Available:
11. Y. Yao et al., "Rule-Based vs Generative AI Chatbots for Postpartum Mental Health Support," Health Informatics Journal, vol. XX, 2023.
12. M. A. Varghese, P. Sharma, and M. Patwardhan, "Public Perception on Artificial Intelligence–Driven Mental Health Interventions: Survey Research," JMIR Formative Research, vol. 8, Art. e64380, Nov. 2024, doi: 10.2196/64380.
13. A. Naik et al., "Evaluating Empathy and Safety in LLM-Based Mental Health Chatbots: A Scenario-Based Study," Journal of AI and Health, vol. XX, 2025.
14. M. U. Javeed, M. S. Ali, A. Iqbal, M. Azhar, S. M. Aslam, and I. Shabbir, "Transforming heart disease detection with BERT: Novel architectures and fine-tuning techniques," in *Proc. Int. Conf. Frontiers Inf. Technol. (FIT)*, Islamabad, Pakistan, 2024, pp. 1–6, doi: 10.1109/FIT63703.2024.10838424.
15. M. Hussain, W. Sharif, M. R. Faheem, Y. Alsarhan, and H. A. Elsalamony, "Cross-Platform Hate Speech Detection Using an Attention-Enhanced BiLSTM Model", Eng. Technol. Appl. Sci. Res., vol. 15, no. 6, pp. 29779–29786, Dec. 2025.
16. M. Javeed, S. Aslam, M. Farhan, M. Aslam, and M. Khan, "An enhanced predictive model for heart disease diagnoses using machine learning algorithms," *Technical Journal*, vol. 28, no. 4, pp. 64–73, 2023. [Online]. Available: https://tj.uettaxila.edu.pk/index.php/technical-journal/article/view/1828
17. S. Aslam, M. U. Javeed, S. M. Aslam, M. M. Iqbal, H. Ahmad, and A. Tariq, "Personality prediction of the users based on tweets through machine learning techniques," *J. Comput. Biomed. Informatics*, vol. 8, no. 2, 2025. [Online]. Available: https://www.jcbi.org/index.php/Main/article/view/796
18. M. U. Javeed, S. M. Aslam, H. A. Sadiqa, A. Raza, M. M. Iqbal, and M. Akram, "Phishing website URL detection using a hybrid machine learning approach," *J. Comput. Biomed. Informatics*, vol. 9, no. 1, 2025. [Online]. Available: https://jcbi.org/index.php/Main/article/view/989
19. M. U. Javeed et al., "A deep learning approach for securing IoT systems with CNN-based prediction of worst-case response time," *Spectrum Eng. Sci.*, vol. 3, no. 7, pp. 376–385, Jul. 2025. [Online]. Available: https://www.sesjournal.com/index.php/1/article/view/599

20. H. Shakeel et al., “LncRNAs disease: A text mining approach to find the role of lncRNA in aging,” *J. Comput. Biomed. Informatics*, vol. 9, no. 1, 2025. [Online]. Available: https://www.jcbi.org/index.php/Main/article/view/1000
21. M. Jaffar, “Ontology-based sentiment analysis for real-time product reputation modeling,” *Spectrum Eng. Sci.*, vol. 3, no. 7, pp. 648–667, Jul. 2025.
22. M. Nauman et al., “Deep transfer learning for COVID-19 screening: Benchmarking ResNet50, VGG16, and GoogleNet on chest X-ray images,” *Int. J. Adv. Comput. Emerg. Technol.*, vol. 1, no. 2, pp. 69–83, Aug. 2025. [Online]. Available: https://ijacet.com/index.php/IJACET/article/view/18
23. M. U. Javeed, M. Nauman, S. Aslam, S. M. Aslam, M. M. Zahoor, Z. Raza, *et al.*, “Deep learning in hematology: Automated counting of blood cells using YOLOv5 object detection,” *Int. J. Adv. Comput. Emerg. Technol.*, vol. 1, no. 3, pp. 16–34, 2025. [Online]. Available: https://ijacet.com/index.php/IJACET/article/view/21
24. M. U. Javeed, M. Akram, S. M. Aslam, S. Hussain, M. M. Nazakat, and M. Nauman, “Predicting customer loyalty from e-commerce reviews using aspect-based sentiment analysis and ANN,” Asian Bull. Big Data Manag., vol. 5, no. 3, pp. 49–61, 2025, doi: 10.62019/3akt6733.
25. M. Nauman, M. U. Javeed, M. T. Jahangir, S. Aslam, M. K. Hussain, Z. Raza, et al., “Optimized deep convolutional neural network for robust occluded facial expression recognition,” Asian Bull. Big Data Manag., vol. 5, no. 3, pp. 62–80, 2025, doi: 10.62019/dpfhnf43.
26. A. Raza et al., “Automated classification of humpback whale calls in four regions using convolutional neural networks and multi-scale deep feature aggregation (MSDFA),” *Measurement*, 118038, 2025.
27. A. Raza, M. Javed, A. Fayad, and A. Y. Khan, “Advanced deep learning-based predictive modelling for analyzing trends and performance metrics in stock market,” *J. Account. Finance Emerg. Econ.*, vol. 9, no. 3, pp. 277–294, 2023.
28. A. Raza et al., “Intelligent image gallery system using deep learning for automated fruit and vegetable classification,” *Int. J. Adv. Comput. Emerg. Technol.*, vol. 1, no. 3, pp. 1–15, 2025. [Online]. Available: https://ijacet.com/index.php/IJACET/article/view/20
29. A. Raza, M. U. Javeed, M. Javed, S. M. Aslam, R. Saleem, W. Irshad, et al., “A hybrid machine learning framework for personalized news recommendation,” *Int. J. Adv. Comput. Emerg. Technol.*, vol. 1, no. 3, pp. 49–62, 2025, doi: 10.64796/qwjyd222.
30. M. U. Javeed, “Unveiling ambivalence in reviews: Using Sentence-BERT and K-NN for airline recommendations,” *Technical Journal*, vol. 30, no. 3, pp. 51–59, Sep. 2025.
31. M. U. Javeed, M. Sahar, S. M. Aslam, M. Iqbal, H. Cao, W. Y. Ramay, et al., “An ANN-based intelligent system for measuring customer loyalty,” Kohsar J. Comput. Inf. Sci., vol. 8, no. 1, Jul. 2025, doi: 10.51153/kjcis.v8i1.250.
32. S. Saifullah, M. U. Javeed, A. Ayub, M. M. Zahoor, Z. Raza, S. M. Aslam, et al., “Impacts of cloudburst events on biodiversity in Pakistan’s northern areas and development of a machine learning model for cloudburst prediction,” Int. J. Adv. Comput. Emerg. Technol., vol. 1, no. 4, pp. 44–63, Oct. 2025, doi: 10.64796/y82wy489.
33. M. U. Javeed, M. Jaffar, U. Saleem, L. Ai, S. M. Aslam, M. Aleem, et al., “An evaluation of machine learning classifiers for nominal weather humidity prediction,” Int. J. Adv. Comput. Emerg. Technol., vol. 1, no. 4, pp. 1–21, 2025, doi: 10.64796/r3m9r804.
34. A. Raza, M. Javed, S. M. Aslam, M. U. Javeed, P. Samual, S. Raheem, et al., “AI-powered sentiment analysis for social media opinion mining: A hybrid NLP and machine learning approach,” Int. J. Adv. Comput. Emerg. Technol., vol. 1, no. 5, pp. 1–15, Nov. 2025, doi: 10.64796/7ytc4b98.
35. M. Aleem, M. U. Javeed, A. Shaheen, W. Ullah, S. M. Aslam, and H. M. J. Haider, “A Machine Learning-Based Framework for Pre-Deployment Prediction of Mobile Application Success,” Int. J. Adv. Comput. Emerg. Technol., vol. 1, no. 5, pp. 28–54, Nov. 2025. [Online]. Available: https://doi.org/10.64796/sqxa6s02
36. M. U. Javeed, S. M. Aslam, W. Y. Ramay, M. M. Zahoor, Z. Raza, and M. M. Iqbal, “The Role of Artificial Intelligence in Academic Writing: A Literature Review,” AMAR, vol. 3, no. 11, pp. 242–258, Nov. 2025. [Online]. Available: https://doi.org/10.63075/qk7n8h28
37. M. R. Faheem, T. Anees, M. Hussain, A. Ditta, H. Alquhayz and M. A. Khan, "Indexing in WoT to Locate Indoor Things," in IEEE Access, vol. 11, pp. 53497-53517, 2023, doi: 10.1109/ACCESS.2023.3272691.
38. Hussain, M., Chen, C., Hussain, M. et al. Optimised knowledge distillation for efficient social media emotion recognition using DistilBERT and ALBERT. Sci Rep 15, 30104 (2025). https://doi.org/10.1038/s41598-025-16001-9
39. Sankoh, A. P., Raza, A., Parwez, K., Shishah, W., Alharbi, A., Javed, M., & Bilal, M. (2025). Automated Facial Pain Assessment Using Dual-Attention CNN with Medical-Grade Calibration and Reproducibility Framework.