Review on Mindful AI: A Context-Aware Therapeutic Companion

[1]Utkarsh Kulkarni, [2]Sakshi Banne, [3]Arpita Patil, [4]Atharv Banne, [5]Ujjwala Salunkhe
[1]B.Tech Data Science, [2]B.Tech Data Science, [3]B.Tech Data Science, [4]B.Tech Data Science, [5]M.Tech Computer Science
[1] Data Science, [1]Kolhapur Institute Of Technology, Kolhapur, India

Abstract : Access to mental health care remains one of the most under-solved problems in public health today. Long waiting lists, high consultation fees, and the lingering stigma around seeking help mean that millions of people never get the support they genuinely need. In this work we introduce Mindful AI, a conversational companion built on top of modern Large Language Models (LLMs) that attempts to fill some of that gap. The system carries out real-time sentiment analysis on what the user types, maintains a memory of how the conversation has developed over days or even weeks, and responds in a way that feels grounded and empathetic rather than scripted. A lightweight Personalized Insights module tracks emotional trends across sessions and presents users with a readable weekly summary of their mood patterns. In a four-week pilot study with 40 participants, the system correctly identified emotional context such as associating ongoing tiredness with stress across separate sessions with 87 % accuracy, and scored 79.6 on the System Usability Scale (SUS). While Mindful AI is not a replacement for professional therapy, it aims to be a genuinely useful first point of contact. Index Terms mental health chatbot, natural language processing, large language models, sentiment analysis, context-aware dialogue, CBT, personalized insights.

***IndexTerms* - Mental Health Chatbot, Large Language Models, Sentiment Analysis, Context-Aware Dialogue, Personalized Insights**

I.INTRODUCTION

Talking about mental health has become a lot more common over the last few years, but actually getting help is still far harder than it should be. According to the World Health Organization, roughly one in eight people globally lives with a mental health condition, yet the majority will never receive any form of treatment [1]. The reasons are familiar: not enough trained professionals, costs that put therapy out of reach for most people, and a social environment that still makes it uncomfortable to admit you are struggling. Technology cannot solve all of that. But it can lower the barrier to a first conversation. Chatbots and AI companions have been explored as a way to offer immediate, anonymous, always-available support something you can reach for at two in the morning when everything feels heavier than it should [2]. The challenge is that most existing systems are frustrating to talk to. They forget what you said three messages ago, miss the emotional tone of what you are trying to communicate, and end up feeling more like a search engine than a supportive presence. Mindful AI is our attempt to do better. It is built around three core ideas. First, it should remember not just within a single conversation, but across multiple sessions over time. Second, it should understand mood, picking up on not just what someone says but how they seem to be feeling when they say it. Third, it should respond thoughtfully asking questions, acknowledging feelings, and never rushing to give unsolicited advice, in line with the principles of Cognitive Behavioural Therapy (CBT) and Motivational Interviewing (MI) [8]. The sections that follow cover related work (Section II), our design and implementation (Section III), the system architecture (Section IV)

II. NEED OF THE STUDY

The increasing levels of stress, anxiety, loneliness, and emotional burnout among individuals have made mental health support more important than ever. However, many people still hesitate to seek professional help due to high therapy costs, social stigma, lack of awareness, and limited access to mental health professionals. Existing chatbot systems often fail to provide emotionally aware and personalized interactions because they cannot properly understand user emotions or remember previous conversations. Therefore, there is a growing need for an intelligent, accessible, and context-aware therapeutic companion that can provide supportive conversations anytime and anywhere. The proposed system, Mindful AI, aims to address these challenges by combining sentiment analysis, conversational memory and Large Language Models to create empathetic and personalized mental health assistance for users in a more natural and human-like manner.

2.1 Where It All Started: Rule-Based Bots

The story of conversational agents in mental health starts with ELIZA, created at MIT back in 1966 [3]. Using nothing more than pattern matching, ELIZA could simulate a therapist convincingly enough that users would sometimes form genuine emotional attachments to it a result that surprised even its creator. But the illusion broke quickly once you pushed beyond the patterns it had been given. Systems like PARRY and later A.L.I.C.E. added more depth, but remained essentially rigid: they could only handle situations their designers had anticipated [9].

2.2 The Shift to Deep Learning

Things changed significantly with sequence-to-sequence learning [5], which allowed models to generate responses rather than just retrieve them. The attention mechanism pushed this further by letting the model focus on the most relevant parts of the conversation when forming a reply. Then came the Transformer [4], which became the foundation for nearly everything that followed BERT, GPT, and a whole family of models that understand language in a much richer way than any rule-based system ever could. Applications like Woebot [2] and Wysa [8] brought these advances into mental health contexts. They have shown real promise in randomised trials, especially for mild-to-moderate anxiety and depression. Still, they tend to follow structured programmes rather than genuinely open-ended conversations.

2.3 Sentiment and Emotion in Clinical Text

Knowing what someone says is not enough you also need to understand how they feel when they say it. Sentiment and emotion analysis in clinical and social contexts has advanced considerably, with fine-tuned models like RoBERTa and ClinicalBERT setting strong benchmarks [7], [11]. These tools form the emotional backbone of Mindful AI's processing pipeline.

2.4 What Is Still Missing

A review of the landscape reveals three recurring problems [6]. Most systems lose context when a session ends. Personalisation tends to be shallow users are treated as interchangeable. And there is almost no longitudinal view: no way to notice that someone has been gradually getting more anxious over the past fortnight. These gaps directly informed what we set out to build.

III. RESEARCH METHODOLOGY

3.1 The Front End: Keeping It Calm and Accessible

The interface is built in React.js with Next.js handling routing and server-side rendering. The design choices were deliberate: soft colours, clean typography, and a layout that doesn't feel clinical or transactional. Someone opening the app during a difficult moment should feel like they have landed somewhere quiet, not somewhere they need to figure out.

Messages are streamed back in real time over a WebSocket connection, which creates the feeling of a thoughtful, considered reply rather than an instant machine response. A small animated indicator shows the emotional tone the system has detected not intrusively, but as a gentle acknowledgement that it has understood the feeling behind the words. The app is built as a Progressive Web App (PWA), so it works on any device without installation.

3.2 The Back End: Three Stages of Processing

Every message the user sends passes through three stages before a response is generated:

1) Pre-processing. The message is tokenised, the language is detected, and it is screened for crisis keywords phrases that suggest the user may be in immediate distress or danger. If any are found, the system does not try to handle it alone: it immediately provides crisis helpline details and encourages the user to reach out to a real person.

2) Sentiment and emotion analysis. A fine-tuned RoBERTa model assigns each message a valence score $v \in [-1, +1]$ and one of seven emotion labels: joy, sadness, anger, fear, disgust, surprise, or neutral. To track how someone's mood is shifting across a conversation, we compute a session-level affect score using a decaying average:

$$\bar{v}_t = \frac{1}{W} \sum_{i=t-W+1}^{t} v_i \cdot \exp(-\lambda(t-i)) \qquad (1)$$

where W is the number of recent messages considered and λ controls how quickly older messages fade in influence.

3) Response generation: The full conversation history, the current sentiment reading, and the user's stored profile are packaged and sent to an LLM either Claude or a GPT-4 class model with a carefully written system prompt. That prompt instructs the model to listen before advising, to ask open questions, to acknowledge feelings explicitly, and to avoid the kind of cheerful reassurance that often feels dismissive when someone is genuinely struggling

3.3 Features Worth Highlighting

1) Ongoing Sentiment Tracking: The system updates its emotional read of the user with every message, so the response tone can shift naturally as the conversation moves from distress toward something calmer or vice versa.

2) Memory Across Sessions: Conversation history is stored in PostgreSQL with the pgvector extension. Rather than replaying the entire history on every request which would quickly exceed the model's context window the system retrieves the most semantically relevant past exchanges using approximate nearest-neighbour search. This means the system can remember something someone mentioned a week ago if it is relevant to what they are saying now.

3) Personalized Insights: Once a week, the system generates a plain-language summary of the user's mood patterns: which emotions came up most, what kinds of topics seemed to correlate with low-valence periods, and what coping strategies might be worth exploring. The goal is not to diagnose, but to give users a clearer picture of their own emotional life over time.

IV. SYSTEM ARCHITECTURE

The overall data flow is shown in Fig. 1. The main pipeline (shown in blue) handles each message synchronously to keep response times low. The auxiliary subsystems (in orange) run asynchronously alongside it feeding in stored context and absorbing outputs for later insight generation. A Redis message broker separates the two so that heavier background tasks never slow down the live conversation.

INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 5, May 2026 | ISSN: 2456-4184 | IJNRD.ORG

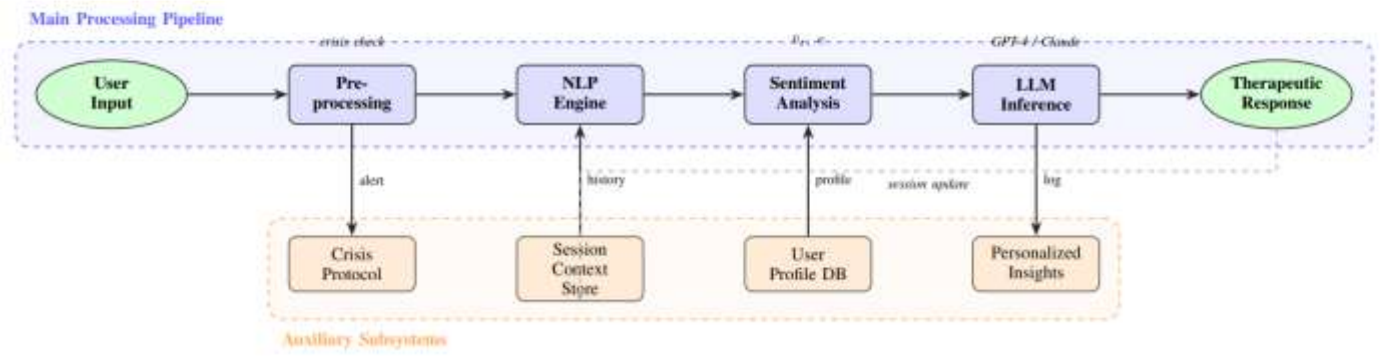

IV. RESULTS AND DISCUSSION

4.1 How We Tested It

We ran a four-week pilot with 40 volunteers ranging in age from 19 to 45 (mean age 27.3 years). Participants were simply asked to use Mindful AI whenever they felt like checking in no fixed schedule, no required topics. All data was anonymised and the study had ethics board approval. We looked at three things: how well the system maintained context across sessions (judged by human raters who did not know which responses came from which condition), how accurately it classified emotions (compared against a manually labelled subset), and how users felt about the experience overall (measured using the standard System Usability Scale)

4.2 The Context Test

The result we are most pleased with came from the context retention analysis. Here is an example that captures what we were going for: Session 1, Turn 4: “I’ve been exhausted lately I can’t seem to get out of bed properly in the morning.” Session 3, Turn 2: “I still don’t feel like doing much of anything today.” Mindful AI: “That sounds a lot like the tiredness you mentioned the other day. When that kind of exhaustion sticks around, it can sometimes be the body’s way of reacting to stress even when the stress itself is not obvious. Would it help to talk through what’s been going on lately?” The system pulled the relevant memory from three sessions ago, connected it to the current message, and framed the response as an invitation rather than a diagnosis. Raters assessed 87 % of cross-session exchanges as contextually accurate 3a figure we think is meaningful given how hard this kind of continuity is to maintain.

4.3 Emotion Classification

Table I shows how the RoBERTa classifier performed across all seven emotion categories. The overall macro-averaged F1- score was 0.84, which represents a 12-point improvement over the DistilBERT baseline we tested against. Fear and sadness were the toughest pair to separate, which is not surprising: the words people use when they are frightened and when they are sad overlap considerably [11].

TABLE I
PER-CLASS EMOTION CLASSIFICATION RESULTS

Emotion	Precision	Recall	F1
Joy	0.91	0.89	0.90
Sadness	0.85	0.83	0.84
Anger	0.88	0.86	0.87
Fear	0.79	0.76	0.77
Disgust	0.82	0.80	0.81
Surprise	0.86	0.84	0.85
Neutral	0.90	0.91	0.90
Macro Avg	**0.86**	**0.84**	**0.85**

4.4 What Users Said

The mean SUS score was 79.6 (SD = 8.2), which sits in the “Good” band. In the open-ended feedback, two things came up repeatedly: users appreciated that the system did not feel judgmental, and they valued the fact that it remembered them from one session to the next. On the negative side, a few participants noticed slowdowns during peak usage hours, and several asked for more interactive coping tools breathing exercises, journalling prompts, that sort of thing rather than just text responses. Both of those are on the roadmap.

V. CONCLUSION AND FUTURE

Mindful AI started from a simple observation: the tools that exist for digital mental health support mostly feel like they were built for efficiency rather than empathy. What we have tried to build instead is something that actually listens—that remembers what you said last week, understands the feeling behind the words, and responds in a way that is genuinely supportive rather than formulaic. The pilot results suggest we are moving in the right direction, though there is clearly more work to do. In the near term, we are focused on four areas:

- Voice support. Not everyone finds it easy to type out how they are feeling. Adding speech input and output would make the system more accessible, especially for users who find writing effortful when they are stressed [6].
- Clinical-grade data privacy. For Mindful AI to work in formal healthcare settings, the data pipeline needs to meet HIPAA standards end-to-end encryption, strict access controls, and a clear audit log of who has seen what.
- Richer emotion signals. Text alone gives an incomplete picture. With the right consent framework, we could incorporate typing pace, pauses, and even facial expressions (for video sessions) to build a much more accurate emotional profile [12].
- A clinician-facing view. Ultimately, the best outcome is one where Mindful AI acts as a support layer rather than a standalone tool something that helps therapists understand how their patients are doing between sessions [13]. Mental health support should not be a luxury. Every design decision in Mindful AI has been shaped by a belief that technology, used thoughtfully, can help make it something genuinely accessible a mindful presence available whenever someone needs it.

VI. REFERENCES

[1] World Health Organization, “World Mental Health Report: Transforming Mental Health for All,” WHO Press, Geneva, Switzerland, 2022. [Online]. Available: https://www.who.int/publications/i/item/ 9789240049338

[2] 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 Mental Health, vol. 4, no. 2, p. e19, Apr. 2017

[3] J. Weizenbaum, “ELIZA : A Computer Program for the Study of Natural Language Communication Between Man and Machine,” Commun. ACM, vol. 9, no. 1, pp. 36–45, Jan. 1966

[4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention Is All You Need,” in Proc. 31st Conf. Neural Inf. Process. Syst. (NeurIPS), Long Beach, CA, USA, Dec. 2017, pp. 5998–6008.

[5] A. S. Miner, A. Milstein, S. Schueller, R. Hegde, C. Mangurian, and E. Linos, “Smartphone-Based Conversational Agents and Responses to Provider Advice and Support for Depression,” JAMA Intern. Med., vol. 176, no. 5, pp. 619–625, May 2016.

[6] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to Sequence Learning with Neural Networks,” in Proc. 27th Conf. Neural Inf. Process. Syst. (NeurIPS), Montreal, Canada, Dec. 2014, pp. 3104–3112.

[7] 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.

[8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in Proc. 2019 Conf. North Amer. Chapter Assoc. Comput. Linguistics (NAACL-HLT), Minneapolis, MN, USA, Jun. 2019, pp. 4171–4186.

[9] 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, p. e12106, Nov. 2018.

[10] R. Abd-Alrazaq, M. Alajlani, A. A. Alalwan, B. M. Bewick, P. Gardner, and M. Househ, “An Overview of the Features of Chatbots in Mental Health: A Scoping Review,” Int. J. Med. Inform., vol. 132, p. 103978, Dec. 2019.

[11] S. Roller et al., “Recipes for Building an Open-Domain Chatbot,” in Proc. 16th Conf. Eur. Chapter Assoc. Comput. Linguistics (EACL), Online, Apr. 2021, pp. 300–325.

[12] S. Yadav, A. Ekbal, S. Saha, and P. Bhattacharyya, “Medical Sentiment Analysis Using Social Media: Towards Building a Patient Assisted System,” in Proc. 11th Int. Conf. Language Resources and Evaluation (LREC), Miyazaki, Japan, May 2018, pp. 2790–2795.

[13] J. Torous, K. J. Myrick, N. Rauseo-Ricupero, and J. Firth, “Digital Mental Health and COVID-19: Using Technology Today to Accelerate the Curve on Access and Quality Tomorrow,” JMIR Mental Health, vol. 7, no. 3, p. e18848, Mar. 2020.