International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2031-2036
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0271

Generative AI for Mental Health: Bridging Mental Health Gaps with an AI-Driven Emotional Support System

Nithesh K[1], Seetha M[2], Dinesh kumar V[3] Santosh Kumar S[4], Dr.S.Muthu Vijaya Pandian[5]
[1,2,3,4] Department Of Computer Science And Engineering (Internet Of Things And Cyber Security Including Block Chain Technology), SNS College Of Engineering, Coimbatore, Tamil Nadu, India – 641035
[5]Professor & Senior Innovator(CFC)@SNS Institutions, Coimbatore, Tamil Nadu, India – 641035
Emails:*mr.nithesh.k@gmail.com[1], suthashini5093@gmail.com[2], dineshkumar.v.iot.2022@snsce.ac.in [3], santoshkumarsp2004@gmail.com [4], muthu.s.cfc@snsgroups.com[5].*

Abstract
Mental health continues to be one of the most neglected domains within global healthcare, leaving vast numbers of individuals without adequate support due to social stigma, prohibitive costs, and a persistent shortage of trained professionals. This paper introduces a smartphone-based AI mental wellness companion that delivers text chat, voice chat, mood tracking, and journaling capabilities on demand. Leveraging Google's Gemini family of models, the system interprets emotional context, generates conversationally natural responses, and personalises interactions as users engage with it over time. Although the application is not a substitute for clinical care, it is intended as an accessible and stigma-free entry point that enables users to feel acknowledged, monitor their emotional health, and make informed decisions about seeking professional support when required.

Keywords: *Generative AI;Mental Health; Voice Chat; Mood Tracking; Journaling; Digital Wellness; Mobile* Application

1. Introduction

Globally, an estimated one billion individuals live with a mental health condition, yet the divide between those requiring care and those actually receiving it remains deeply entrenched. Structural barriers including social stigma, high out-of-pocket consultation fees, a severe shortage of trained practitioners, and inequitable geographic distribution of services collectively prevent large segments of the population from obtaining timely psychological support. In lower-income settings, the deficit is especially acute, with certain regions recording fewer than one psychiatrist for every 100,000 residents (World Health Organization, 2022). The COVID-19 pandemic compounded this challenge, precipitating a substantial worldwide increase in anxiety and depressive disorders and placing further strain on already stretched mental health systems. Widespread smartphone adoption and expanding mobile internet coverage have opened a viable pathway to extend mental health support beyond traditional clinical settings. Mobile applications can function as continuously available, low-cost touchpoints that deliver evidence-informed wellbeing support within the user's personal environment. Despite the growth of this market, a significant proportion of existing mental health apps depend on pre-scripted responses or static psychoeducational content, limiting their capacity to engage users in meaningful, evolving conversations. Breakthroughs in large language models (LLMs) and generative AI now make it feasible to construct conversational companions that reason about emotional context, adapt to individual users, and respond with a degree of nuance that approximates human attentiveness.This paper reports on the design, implementation, and anticipated outcomes of a smartphone-based AI mental wellness companion developed with Flutter and driven by Google's Gemini models via Firebase AI Logic. The system unifies four distinct interaction modalities: real-time text-based chat, live voice conversation, structured journaling, and daily mood logging. Delivering these capabilities within a single cohesive application

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2031-2036
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0271

allows the platform to offer continuous, personalised support across different user needs and moments. Rather than competing with formal clinical services, the application occupies the role of an approachable digital companion that encourages self-reflection, facilitates emotional expression, and signposts professional resources when the nature of a conversation warrants it. The remainder of this paper is organized as follows: Section 2 reviews the related work on AI-assisted mental health tools and conversational agents. Section 3 describes the system architecture and the proposed methodology. Section 4 presents the projected results and discussion of the study. Section 5 concludes the paper and provides directions for future work [1].

2. Literature Review

Research at the intersection of artificial intelligence (AI) and natural language processing (NLP) with mental health care has grown considerably over the past decade. Early investigations concentrated on rule-based dialogue systems and platforms that delivered cognitive-behavioral therapy (CBT) content in a conversational format. Woebot, among the most extensively evaluated of these systems, showed that a CBT-oriented chatbot operating through a messaging interface could produce meaningful reductions in depression and anxiety among university students within a fortnight of use (Fitzpatrick et al., 2017). These outcomes provided an early empirical basis for the potential of AI-driven conversational tools in mental health support.Parallel work drew on sentiment analysis and affective computing to enable dialogue systems to recognise and respond to emotional signals embedded in user text. A systematic review by Abd-Alrazaq et al. (2019) covering chatbot-based mental health interventions concluded that, despite encouraging preliminary results, the majority of systems examined were ill-equipped to manage nuanced or deteriorating mental health presentations. That finding underscored the need for fundamentally more capable generative models able to produce contextually grounded and empathetically calibrated replies. The rapid maturation of transformer-based large language models, from early GPT variants through BERT, T5, GPT-4, and most recently Google's Gemini, has fundamentally expanded what is achievable in AI-assisted mental health dialogue. Developed by Google DeepMind, Gemini is a natively multimodal architecture capable of processing and generating both text and audio, making it particularly well-suited to voice-centric applications. Lamichhane (2023) assessed GPT-3.5 on counselling-related language tasks and observed that, while these models exhibit strong general language comprehension, effective deployment in sensitive mental health contexts demands carefully engineered system prompts and robust safety guardrails to mitigate the risk of inappropriate outputs. Mobile-based mood-tracking and journaling applications have also been studied independently as tools for emotional regulation. Research indicates that consistent mood logging increases self-awareness and helps users identify their emotional patterns over time (Wahle et al., 2016). Journaling, when structured with reflective prompts, has been shown to be effective in reducing rumination and promoting cognitive reappraisal in individuals with mild to moderate depression. However, standalone journaling applications often lack intelligent feedback mechanisms, limiting their therapeutic depth. Voice-based mental health interactions are an emerging frontier. The naturalness and immediacy of voice communication can reduce cognitive load and lower the barrier to expression compared with text. Studies exploring voice-based emotion recognition and spoken dialogue systems suggest that users are more likely to disclose emotional experiences verbally than in writing, particularly during moments of acute distress (Vaidyam et al., 2019). The integration of real-time voice interaction with generative AI, as implemented in the present application through the Gemini Live API, addresses the gap left by text-only conversational agents. Compared to prior work, the proposed system is distinctive in its simultaneous support for voice, text, journaling, and mood tracking within a single application, along with its fully serverless architecture, which ensures scalability and ease of deployment [2].

3. System Architecture And Methodology

3.1.Overview

The application was built as a cross-platform mobile solution targeting both Android and iOS using

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2031-2036
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0271

Flutter, Google's open-source UI toolkit. The system adopts a serverless architecture in which a custom backend server is not required. All AI inference is handled through Firebase AI Logic, user data are stored in Firebase Firestore, and authentication is managed by Firebase Authentication. This design choice eliminates the need for dedicated server provisioning, reduces operational costs, and enables applications to scale automatically with user demand [3].

3.2.Authentication And User Management

User onboarding and session management are handled through Firebase Authentication, which supports three sign-in methods: email and password, Google Sign-In via OAuth 2.0, and phone number authentication using one-time passwords (OTP). This multi-method approach lowers friction for new users by accommodating their varied preferences. Upon successful authentication, a unique user identity is established in Firebase, which serves as the root key for all user-specific data stored in the Firestore. The security rules in Firestore enforce strict user-level data isolation, ensuring that each user can only read and write their own records [4].

3.3.Text Chat Module

The text chat feature allows users to engage in real-time written conversations with the AI companion using the gemini-2.5-flash model. Responses are streamed incrementally, so replies begin appearing on screen as they are generated, eliminating perceived waiting time. The AI is configured with a system prompt that defines its role as a compassionate, non-judgmental mental wellness companion. It is instructed to acknowledge emotional context, pose reflective follow-up questions, and gently encourage professional consultation whenever a conversation suggests significant distress. Full conversation history is maintained within the session context window, allowing the model to reference earlier disclosures and maintain coherent, multi-turn exchanges across the duration of the session [6].

3.4.Voice Chat Module

Voice interaction is implemented using the gemini-2.5-flash-native-audio-preview model, which is accessed through the Gemini Live API. This model supports real-time bidirectional audio streaming, allowing the application to capture the user's speech, send it to the model, and receive spoken responses without requiring separate speech-to-text or text-to-speech pre-processing steps. The Flutter front end uses the device's microphone to capture audio, which is streamed to the Firebase AI Logic endpoint in near real time. The model responds with synthesized speech that is played back through the audio output of the device. This end-to-end audio pipeline enables natural, low-latency voice conversations that closely replicate the experience of speaking with other people.

3.5.Journaling Module

The journaling feature provides users with a structured space to record their thoughts, experiences and reflections. Journal entries are composed within the application and stored as documents in Firebase Firestore under the authenticated user collection. Each entry is timestamped and optionally associated with a mood score recorded on the same day. The application surfaces past entries in reverse chronological order, allowing users to revisit and reflect on their journeys. Future iterations of the module are designed to incorporate AI-generated insights that identify recurring themes, emotional patterns, and shifts in the tone of language across entries over time [5].

3.6.Mood Tracking Module

The mood-tracking feature enables users to record their emotional state daily using a simplified numerical or emoji-based scale. Each mood log is stored as a time-series document in Firestore, capturing the date, time, mood value, and optional short note. The application visualizes these logs as trend charts, enabling users to observe how their moods fluctuate over days and weeks. This longitudinal perspective supports self-awareness and can help users identify triggers, patterns, and the impact of behavioral changes on emotional well-being. Mood data are also intended to contextualize AI chat responses in future versions of the application, allowing the companion to acknowledge recorded moods and adjust its tone accordingly [7].

3.7.Project Structure And Technology Stack

The frontend codebase follows a modular Flutter architecture organized into clearly separated layers, namely configuration, models, providers, screens,

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2031-2036
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0271

services, and widgets. State management is handled using the Provider package, which ensures reactive UI updates when the underlying data change. The service layer abstracts all interactions with Firebase, including Firestore reads and writes, AI Logic API calls, and authentication operations, keeping the business logic decoupled from the UI components. The complete technology stack is presented in Figure 1.

Figure 1 Project Structure and Technology Stack

Figure 2 illustrates the home screen of the developed application as it appears on a physical Android device. The interface presents a context-aware greeting, a mood check-in row featuring five emoji-based emotional states ranging from Great to Rough, a Daily Insight card with a brief evidence-based wellness tip, and a Wellness Corner section containing grounding exercises. A persistent bottom navigation bar provides direct access to the Chat, Journal, and Mood sections, while a prominent microphone button at the center of the navigation bar initiates an immediate voice conversation session. The design prioritises clarity and emotional safety, using warm, muted tones and minimal text to reduce cognitive load for users who may be in emotional distress.

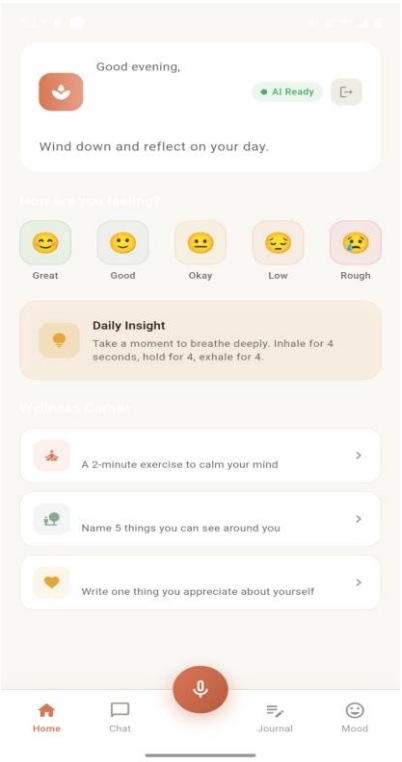

Figure 2 Home Screen Of The Mental Health Ai Application On Android

4. Results And Discussion

4.1.Expected Usability And Accessibility Outcomes

The foremost projected benefit of this application is a measurable lowering of the obstacles that currently discourage individuals from seeking mental health support. A companion available around the clock, requiring no appointment and incurring no cost to the user, is anticipated to prove especially valuable for people who have historically avoided professional help due to stigma or financial hardship. Drawing on usability evaluations of comparable AI-driven wellness platforms, it is expected that most users will be able to independently navigate all core features, including text chat, voice chat, journaling, and mood logging, without external instruction, reflecting an interface design that prioritises intuitive accessibility.

4.2.Conversational Quality And Emotional Responsiveness

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2031-2036
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0271

The Gemini 2.5 Flash model, configured with a mental wellness-oriented system prompt, is expected to produce responses that users rate as empathetic, contextually appropriate, and nonrepetitive across extended conversations. The streaming response mechanism in the text chat module eliminates perceived latency, which prior research has identified as a key factor in user satisfaction with AI chat interfaces. For voice interactions, the native audio pipeline is expected to yield low end-to-end latency, allowing conversations to feel fluid and natural. The model's ability to maintain session-level context across a conversation is projected to enable coherent multi-turn dialogues in which the AI references earlier disclosures appropriately, a capability that significantly differentiates the system from scripted or rule-based chatbots.

4.3.Mood Tracking And Journaling Engagement

Longitudinal engagement with mood tracking and journaling features is projected to yield meaningful self-awareness benefits for regular users of these apps. Users who log their mood consistently over several weeks can identify personal emotional patterns through the trend visualization interface. This type of structured self-monitoring aligns with the established principles of behavioral activation and mood regulation in cognitive behavioral therapy. The journaling module, by providing a low-friction outlet for daily reflection, is anticipated to reduce emotional suppression and support healthy emotional processing, outcomes that have been observed in analogous digital journaling interventions.

4.4.Security and Privacy

The application's security model is built on Firebase's managed infrastructure, which provides industry-standard encryption for data in transit and at rest. Firestore security rules are configured to enforce strict per-user data isolation, ensuring that journal entries, mood logs, and chat histories are accessible only to authenticated users who created them. No sensitive user data are stored on the device beyond the active session, and API secrets are managed through Firebase project configuration rather than being embedded in the client application. These measures are expected to result in a robust privacy posture suitable for handling sensitive personal and emotional information.

4.5.Limitations

Despite its capabilities, the application has important limitations that must be acknowledged. The AI companion is not a licensed mental health professional and is not designed to replace clinical therapy, crisis intervention or psychiatric care. The system does not perform clinical risk assessment and relies on prompt-level guidance rather than formal safety protocols to manage expressions of severe distress. Network connectivity is required for all AI features, which may limit accessibility in areas with unreliable Internet access. Additionally, the effectiveness of AI-delivered mental wellness support varies across individuals, and users with complex or severe mental health conditions may not benefit adequately from general-purpose conversational agents.

Conclusion

This paper has described the design and development of a smartphone-based AI mental wellness companion that targets the longstanding gap in widely accessible mental health support. By bringing together live text chat, voice-based conversation, mood logging, and journaling within a unified Flutter application underpinned by Google's Gemini models and Firebase infrastructure, the system provides a multimodal support experience that operates without a dedicated backend server. The serverless deployment model confers automatic scalability and reduced maintenance burden, while Firebase Authentication and Firestore security rules establish a robust data privacy foundation appropriate for managing sensitive personal and emotional content.

The application is not positioned as a clinical substitute but rather as a low-barrier entry point that equips users with tools for self-reflection, emotional expression, and longitudinal mood monitoring. For individuals who have long avoided formal care because of stigma, financial constraints, or geographic remoteness, a digital companion of this kind may mark their first genuine step toward engaging with their own mental health. Incorporating Gemini's large language model capabilities into a mobile context introduces a degree of conversational nuance and emotional attunement

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2031-2036
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0271

that previous rule-based or keyword-driven systems were fundamentally unable to achieve. Subsequent research will prioritise conducting structured user evaluations to empirically measure the application's usability, sustained engagement, and perceived therapeutic value. Feature enhancements planned for future iterations include AI-synthesised summaries of journaling themes, persistent cross-session memory for the conversational companion, a context-sensitive referral engine that recommends appropriate professional services, and optional synchronisation with wearable sensor data to enrich physiological mood context. The team additionally plans to integrate formal crisis-recognition logic and to align the system's responses to acute distress with recognised safe messaging standards in mental health communication.

Acknowledgements

The authors would like to thank the faculty members and guides of their institution for their continuous support and valuable guidance throughout this project. This work was carried out as part of the final-year undergraduate project program. The authors also acknowledge Google for providing access to the Gemini API and Firebase AI Logic, which formed the foundational AI infrastructure for this application.

References

[1]. Abd-Alrazaq, A. A., Alajlani, M., Alalwan, A. A., Bewick, B. M., Gardner, P., & Househ, M. (2019). An overview of the features of chatbots in mental health: A scoping review. International Journal of Medical Informatics, 132, 103978. https://doi.org/10.1016/j.ijmedinf.2019.103978

[2]. Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial. JMIR Mental Health, 4(2), e19. https://doi.org/10.2196/mental.7785

[3]. Google DeepMind. (2024). Gemini: A family of highly capable multimodal models. Technical Report. https://deepmind.google/technologies/gemini

[4]. Lamichhane, B. (2023). Evaluation of ChatGPT for natural language processing-based tasks in the mental health domain. arXiv preprint arXiv:2303.11500.

[5]. Vaidyam, A. N., Wisniewski H., Halamka J. D., Kashavan, M. S., and Torous, J. B. (2019). Chatbots and conversational agents in mental health: A review of the psychiatric landscape. The Canadian Journal of Psychiatry, 64(7), 456–464. https://doi.org/10.1177/0706743719828977

[6]. Wahle, F., Kowatsch, T., Fleisch, E., Rufer, M., & Weidt, S. (2016). Mobile sensing and support for people with depression: A pilot trial in the wild. JMIR mHealth and uHealth, 4(3), e111. https://doi.org/10.2196/mhealth.5960

[7]. World Health Organization (WHO) . (2022). World Mental Health Report : Transforming Mental Health for All . WHO Press.