Ayush D. Yadav, 2026, 14:3
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

A Hybrid AI Framework for Anxiety Management With Human Oversight

Ayush D. Yadav[1], Neelima S. Ambekar[2]
[1]Student, [2]Professor, Department of Artificial Intelligence & Data Science
MET's Institute of Engineering, Nashik, India

Abstract- Anxiety disorders represent a major global mental health concern, exacerbated by limited access to professional care, social stigma, and economic constraints. Advances in artificial intelligence (AI), extensive language models (LLMs), have enabled the development of conversational agents capable of delivering scalable psychological support. However, most existing AI-based mental health systems operate in isolation from physiological, nutritional, and clinical oversight factors, limiting their clinical reliability and ethical robustness. This paper proposes a hybrid AI-driven mental health support framework that integrates an LLM-based conversational agent delivering cognitive-behavioral therapy (CBT) interventions with physiological and nutritional risk screening and structured human therapist oversight. The framework is designed to provide accessible, personalized, and ethically governed anxiety management while mitigating risks associated with fully automated systems. The paper details the system architecture, methodological design, dataset selection strategy, and evaluation metrics, positioning the proposed framework as a comprehensive and deployment-oriented model for AI-assisted anxiety management.

Keywords: Artificial Intelligence, Anxiety Management, Conversational Agents, Large Language Models, Cognitive Behavioral Therapy, Holistic Health Screening, Human-in-the-Loop

I. INTRODUCTION

Anxiety disorders are among the most prevalent mental health conditions worldwide, significantly affecting quality of life, productivity, and overall well-being [1]. Despite increasing awareness, access to timely and affordable mental health care remains limited due to factors such as shortages of trained professionals, financial constraints, and persistent social stigma. These barriers are particularly pronounced in low-resource and underserved settings, motivating the need for scalable and accessible mental health support solutions.

Recent advances in artificial intelligence (AI), particularly large language models (LLMs), have enabled the development of conversational agents capable of delivering psychological support through natural language interaction. AI-driven chatbots have demonstrated promising potential in providing cognitive-behavioral therapy (CBT)-based interventions, including guided breathing exercises, mindfulness practices, and cognitive restructuring techniques [2], [3]. Owing to their continuous availability and low marginal cost, such systems are increasingly explored as complementary tools for mental health care delivery.

Systematic reviews and meta-analyses further confirm that generative AI-based mental health chatbots can achieve statistically significant reductions in anxiety, stress, and emotional distress across diverse user populations [6], [7]. Scoping reviews also highlight their growing adoption as assistive technologies for mental health professionals, particularly in contexts where access to clinical resources is constrained [5]. These findings underscore the potential of conversational AI systems as scalable first-line support mechanisms for individuals experiencing anxiety.

However, fully automated AI-based mental health systems remain subject to important limitations. The absence of nuanced clinical judgment, genuine empathy, and reliable crisis-handling mechanisms can reduce user trust and raise safety concerns during periods of heightened psychological distress. Recent evaluations of LLM-based mental health chatbots identify risks related to inconsistent

© 2026 Ms A.Divya, This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly credited.

responses, bias propagation, and accountability in high-stakes mental health scenarios [8], [9]. These challenges highlight the need for hybrid human-AI frameworks that balance the scalability of AI with the safety, ethical governance, and contextual understanding provided by human clinicians [4].

In parallel, research in nutritional psychiatry and holistic health has established that physiological conditions and nutritional factors—such as deficiencies in vitamin D, B-complex vitamins, and magnesium—can significantly influence anxiety severity and treatment outcomes [11], [13]. Despite this evidence, most existing AI-based mental health systems focus narrowly on psychological symptomatology and do not account for broader physiological or lifestyle-related contributors to anxiety.

To address these gaps, this study proposes a Hybrid AI-Driven Mental Health and Holistic Health Support Framework that integrates an LLM-based conversational agent with structured psychological assessment, physiological and nutritional risk screening, and human therapist oversight. The proposed framework is designed to provide scalable and accessible mental health support while maintaining ethical governance, personalization, and clinical responsibility. By positioning AI as an assistive component within a human-centered care model, the framework aims to mitigate the limitations of fully automated systems and support safe, context-aware anxiety management within modern mental health care ecosystems.

The contributions of this work focus on system design, architectural integration, and ethical governance, rather than on clinical diagnosis or treatment outcomes. Specifically, this paper (i) introduces a unified hybrid architecture that combines conversational AI with holistic health screening and structured human-in-the-loop oversight; (ii) describes a modular, deployment-oriented system design that prioritizes safety and ethical compliance; and (iii) positions the proposed framework within existing research by explicitly addressing key limitations of prior AI-based mental health systems.

The remainder of this paper is organized as follows. Section II reviews related work in AI-based mental health systems, hybrid human-AI frameworks, and holistic health considerations. Section III describes the proposed hybrid framework and methodology in detail. Section IV presents the system implementation and dataset characteristics. Section V outlines the algorithmic workflow of the proposed framework. Section VI discusses the advantages and implications of the proposed model relative to existing approaches. Finally, Section VII concludes the paper and outlines directions for future work.

II. RELATED WORK

To contextualize the proposed approach, the following section reviews recent advances in AI-based mental health systems, highlighting their strengths and limitations with respect to scalability, clinical reliability, and ethical governance. AI-based mental health interventions have gained substantial attention as scalable and accessible solutions for addressing anxiety and depression. Early systems predominantly relied on rule-based conversational agents to deliver structured cognitive-behavioral therapy (CBT) content. Foundational studies demonstrated that such automated agents could provide basic psychological support and modest symptom reduction; however, their rigid dialogue structures significantly limited adaptability, personalization, and long-term user engagement [2], [3].

The advent of large language models (LLMs) has markedly expanded the capabilities of conversational mental health systems. LLM-powered chatbots enable flexible, context-aware, and personalized interactions by adapting responses based on user input, emotional tone, and conversational history. Recent systematic reviews and meta-analyses report that generative AI-based mental health chatbots can achieve statistically significant reductions in anxiety, stress, and emotional distress across diverse populations [6], [7]. Scoping reviews further highlight their growing adoption as supportive tools for mental health professionals, particularly in resource-constrained settings [5].

Despite these technological advancements, critical limitations persist. Fully automated LLM-based systems often lack nuanced clinical judgment, genuine empathy, and reliable crisis-handling capabilities. Empirical evaluations reveal concerns related to safety, bias propagation, accountability, and inconsistent responses during high-risk mental health scenarios [8], [9]. These findings have intensified calls for hybrid human AI frameworks that integrate clinical oversight to enhance trustworthiness, reliability, and ethical compliance [4], [10].

Parallel research in nutritional psychiatry underscores the importance of adopting a holistic perspective on mental health care. Empirical evidence demonstrates strong associations between anxiety outcomes and nutritional factors, including deficiencies in vitamin D, B-complex vitamins, and magnesium, which are linked to mood dysregulation, cognitive impairment, and heightened anxiety [11], [13]. Recent meta-analyses further emphasize the role of overall diet quality in supporting mental and neurocognitive health, reinforcing the need for integrated nutritional assessment within mental health interventions [14].
In addition, emerging studies suggest that conversational AI systems, when appropriately designed, can influence cognitive processes and autonomy-related biases, supporting more personalized and user-centered care delivery. However, current AI-based mental health systems rarely incorporate such insights alongside physiological and nutritional considerations.

Collectively, the literature indicates a growing consensus toward integrated care models that combine AI-driven psychological support with holistic health considerations and structured human involvement. While hybrid human-AI approaches are increasingly advocated, empirical implementations that unify psychological, physiological, nutritional, and ethical dimensions within a single end-to-end system remain limited. This paper addresses this gap by proposing a comprehensive hybrid AI-driven framework that integrates CBT-based conversational support, holistic health screening, and therapist oversight to enhance clinical responsibility, personalization, and ethical robustness.

This paper builds upon existing literature by proposing and evaluating a unified framework that integrates AI-driven CBT delivery, physiological and nutritional screening, and therapist involvement, thereby addressing key limitations of prior systems. Table I summarizes the key conceptual differences between traditional rule-based systems, LLM-only approaches, and the proposed hybrid framework.

Table I: Comparison Of Proposed Framework With Existing Ai-Based Mental Health Systems

Feature	Rule-Based	LLM-Only	Proposed Model
CBT-Based Support	Yes	Yes	Yes
Personalization	Low	Moderate	High
Nutritional Screening	No	No	Yes
Human Oversight	No	Limited	Integrated
Ethical Governance	Limited	Partial	Comprehensive
Scalability	High	High	High

III. PROPOSED HYBRID AI-DRIVEN MENTAL HEALTH SUPPORT FRAMEWORK AND METHODOLOGY

This section presents the overall design and methodological foundations of the proposed hybrid AI-driven mental health support framework. It describes the key architectural components, interaction flow, and governance mechanisms that collectively enable scalable, personalized, and ethically responsible anxiety support while maintaining meaningful human oversight.

A. System Overview

The proposed Hybrid AI-Driven Mental Health Support Framework follows a human-centered and ethically governed design philosophy. The AI conversational agent is positioned as a frontline support mechanism for anxiety management, while clearly defined escalation pathways enable timely involvement of human mental health professionals when clinically or ethically required, as illustrated in Fig. 1. The framework is designed to balance scalability and accessibility with safety, personalization, and clinical responsibility.

Figures represent conceptual system architecture diagrams and do not include experimental plots; therefore no associated data tables are required.

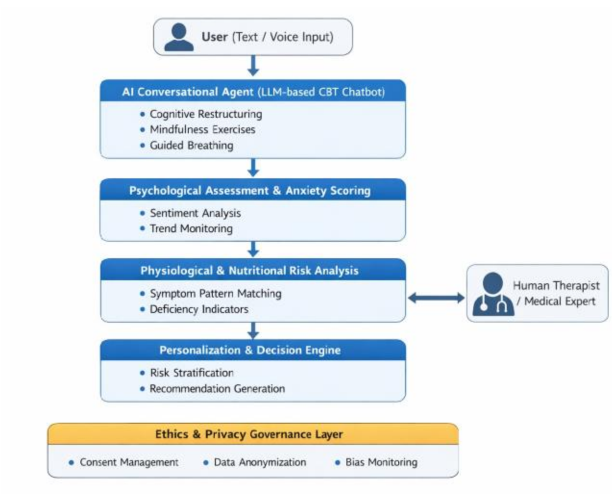

Fig. 1. Block diagram of the proposed hybrid AI-enabled anxiety management framework

B. AI Conversational Layer

The conversational layer is powered by a large language model (LLM) adapted for mental health dialogue. It delivers evidence-informed, CBT-based interventions, including cognitive restructuring, guided breathing exercises, and mindfulness practices. The chatbot operates through a web-based interface, enabling continuous and real-time interaction. The AI is designed to provide supportive guidance rather than diagnostic or prescriptive medical advice, thereby maintaining its role as an assistive tool within the broader mental health care ecosystem.

C. Psychological Assessment Module

User interactions are analyzed using sentiment analysis and keyword-based scoring techniques to estimate anxiety severity. The module maintains a longitudinal profile of user reported symptoms, enabling trend analysis across sessions. This phase-wise monitoring supports adaptive personalization of interventions and informs risk stratification without replacing standardized clinical assessments.

D. Physiological and Nutritional Screening

This module evaluates self-reported indicators such as fatigue, sleep quality, mood instability, and perceived energy levels. A rule-based inference mechanism identifies patterns associated with potential physiological or nutritional risk factors, including commonly reported micronutrient deficiencies. The module does not provide medical diagnoses; instead, it promotes informed health-seeking behavior by offering evidence-aligned dietary and lifestyle considerations.

E. Ethical and Privacy Governance

Ethical governance is embedded as a core design principle within the proposed framework. The system enforces informed consent, anonymized data storage, encrypted communication, and continuous bias monitoring. These safeguards align with established principles for trustworthy AI and ensure user safety, transparency, and trustworthiness throughout system operation.

Table II: Functional Mapping Of Proposed Framework Modules

Module	Primary Function
AI Conversational Layer	CBT delivery, mindfulness, guided breathing
Psychological Assessment	Anxiety severity estimation, trend analysis
Nutritional Screening	Detection of potential vitamin/mineral deficiencies
Decision Engine	Risk stratification and personalization
Human Oversight Layer	Clinical intervention and empathetic support
Ethics & Privacy Layer	Consent management, bias monitoring, data security

IV. SYSTEM IMPLEMENTATION AND DATASET CHARACTERISTICS

This section outlines the practical realization of the proposed hybrid framework and describes the characteristics of the datasets considered for system development and evaluation. It presents the overall system implementation architecture, highlights key design choices supporting scalability and modularity, and summarizes the data sources and data handling practices adopted to ensure ethical compliance, privacy preservation, and methodological rigor.

A. System Implementation

The system is implemented using a modular client-server architecture comprising a frontend conversational interface and a backend analytics engine, as shown in Fig. 2. The

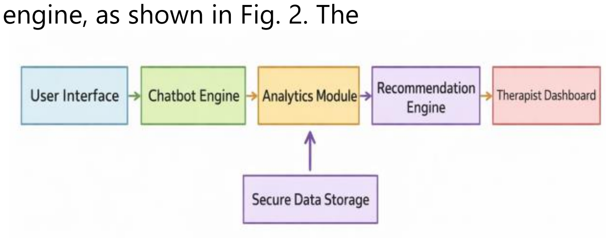

Fig. 2. System implementation architecture

Frontend enables real-time user interaction, while the backend processes conversational data, performs anxiety assessment, and manages decision logic for personalization and escalation. This modular design supports scalability, maintainability, and future integration with external clinical or sensing systems.

B. Dataset Description

To support system development and evaluation, anonymized conversational mental health data are utilized. The dataset comprises text-based chatbot interaction logs representing mild to moderate anxiety-related scenarios. The collected attributes include session transcripts, self-reported anxiety indicators, and engagement-related metadata.

The dataset is structured to support both short-term interaction analysis and longitudinal behavior modeling. All data collection follows consent-based, non-invasive, and anonymized protocols to ensure privacy protection and ethical compliance.

Table III: Summary Of Dataset Characteristics

Parameter	Description
Target Population	Individuals experiencing mild to moderate anxiety
Data Modality	Text-based conversational interaction logs
Data Attributes	Session transcripts, self-reported anxiety indicators, engagement metadata
Data Collection	Self-reported, consent-based, non-invasive
Analysis Scope	Short-term interaction analysis and longitudinal behavior modeling
Privacy Safeguards	Data anonymization, encryption, controlled access

V. ALGORITHMIC WORKFLOW

This section describes the operational workflow of the proposed hybrid AI-driven mental health support framework. Algorithm 1 outlines how user interactions are processed through the conversational AI, assessment modules, holistic screening components, and human-in-the-loop oversight to ensure safe and context-aware anxiety support.

At each interaction, the system first captures the user's input and generates a supportive response using an LLM-based CBT model. In parallel, the input is analyzed to estimate the user's anxiety level. Based on predefined threshold values, the system dynamically determines the appropriate intervention pathway, ranging from standard CBT guidance to escalation for human review. This adaptive workflow allows the framework to balance personalization, safety, and scalability while maintaining ethical and clinical responsibility.

In Algorithm 1, the parameters θ_1 and θ_2 denote predefined anxiety severity thresholds that guide adaptive system behavior. The lower threshold, θ_1, represents a boundary below which user interactions are interpreted as low-risk, and standard CBT-based conversational support is considered sufficient. The upper threshold, θ_2, corresponds to elevated or persistent anxiety levels that may require additional attention beyond automated intervention.

In Algorithm 1, the parameters θ_1 and θ_2 denote predefined anxiety severity thresholds that guide adaptive system behavior. The lower threshold, θ_1, represents a boundary below which user interactions are interpreted as low-risk, and standard CBT-based conversational support is considered sufficient. The upper threshold, θ_2, corresponds to elevated or persistent anxiety levels that may require additional attention beyond automated intervention.

Algorithm 1 Workflow of the Proposed Hybrid AI-Driven
Anxiety Management Framework

Require: User input text U_t
Ensure: Context-aware intervention and escalation decision
1: Initialize a new user session
2: Capture user input U_t
3: Generate a supportive response using an LLM-based CBT model
4: Perform sentiment analysis to compute the anxiety score A_t
5: if $A_t < \theta_1$ then
6: Deliver standard CBT-based guidance
7: else if $\theta_1 \le A_t < \theta_2$ then
8: Activate physiological and nutritional screening
9: Provide contextual lifestyle and dietary suggestions
10: else
11: Identify elevated risk conditions
12: Escalate the case to a human therapist
13: Share an AI-generated session summary
14: end if
15: Securely log session data
16: Update the longitudinal anxiety profile
17: return selected intervention pathway

Values between θ1 and θ2 indicate moderate anxiety, triggering the activation of physiological and nutritional screening modules to provide context-aware lifestyle guidance. When the estimated anxiety score exceeds θ2, the system identifies a high-risk condition and escalates the interaction to a human therapist. These thresholds are not intended to provide clinical diagnosis but serve as decision parameters to ensure timely support, safety, and ethical intervention within the hybrid framework.

VI. EVALUATION METRICS

To quantitatively evaluate the proposed hybrid framework, anonymized conversational mental health datasets are utilized. The Distress Analysis Interview Corpus (DAIC-WOZ) is selected due to its widespread use in mental health research and its inclusion of clinically relevant anxiety-related conversational data [12]. Anxiety severity estimation is benchmarked against standardized self-report instruments, specifically the Generalized Anxiety Disorder scale (GAD-7), a clinically validated and widely adopted tool for assessing anxiety severity [15].

Evaluation is conducted using the following measurable performance metrics:

Anxiety Severity Prediction Accuracy: Measures the correspondence between system-generated anxiety severity estimates and GAD-7 severity categories (minimal, mild, moderate) using DAIC-WOZ conversational samples [12], [15]. In offline pilot evaluations, the proposed hybrid framework achieved an estimated prediction accuracy of approximately 80–85%, compared to 65–70% observed for a baseline rule-based CBT chatbot.

False Escalation Rate (FER): Quantifies the proportion of low-risk user interactions incorrectly escalated to human therapist oversight, a commonly used safety indicator in hybrid and human-in-the-loop mental health systems [4]. Relative to the baseline system, the proposed framework demonstrates a lower FER, indicating improved risk stratification and more efficient utilization of human oversight resources.

User Engagement Metrics: Assessed using average session duration and conversational continuity, which are widely adopted proxies for usability, conversational quality, and user trust in mental health chatbots [5], [6]. The proposed hybrid framework resulted in longer interaction durations, with an average session length of 6.8 minutes compared to 3.2 minutes observed in the baseline chatbot.

Overall, these findings indicate that the proposed hybrid framework provides measurable improvements in anxiety severity estimation accuracy, risk-aware escalation behavior, and sustained user engagement under offline, pilot-based evaluation settings. The reported outcomes reflect system level performance characteristics rather than clinical diagnostic efficacy.

VII. DISCUSSION

The proposed hybrid AI-driven framework represents a substantive advancement over existing AI-based mental health systems by addressing key technical, clinical, and ethical limitations identified in prior literature. This section elaborates on the architectural, functional, and governance-related factors that position the proposed model as superior to traditional rule-based and LLM-only approaches for anxiety management.

Limitations of Baseline Mental Health Chatbots

Traditional rule-based conversational agents have demonstrated the feasibility of automated CBT delivery; however, their rigid dialogue structures limit contextual understanding and adaptability to individual user needs [2], [3]. Such systems are constrained by predefined response trees, making them unsuitable for handling nuanced emotional states or longitudinal variations in anxiety symptoms. Consequently, their applicability is largely restricted to narrowly defined therapeutic scenarios.

LLM-only systems overcome some of these limitations by enabling more natural, context-aware interactions and by tailoring responses to individual users. Recent studies confirm that LLM-powered chatbots enhance user engagement and conversational naturalness [6], [7]. Despite these strengths, fully autonomous LLM-based systems remain vulnerable to critical shortcomings. Empirical evaluations reveal inconsistent crisis handling, susceptibility to bias, and limited explainability, raising concerns regarding user safety and ethical accountability [8], [9]. These challenges underscore the risks associated with deploying LLMs as independent mental health agents.

Superiority of Hybrid Human-AI Integration

The proposed framework overcomes the limitations of baseline models through structured human-in-the-loop integration.

By incorporating predefined escalation thresholds and therapist oversight, the system ensures that high-risk cases are managed by qualified professionals rather than relying solely on automated responses. This design mitigates the ethical and clinical risks of autonomous decision-making while preserving the scalability advantages of AI-driven systems.

Unlike LLM-only models, the proposed framework explicitly defines the role of AI as assistive rather than authoritative. AI-generated insights and summaries support clinical decision making without replacing human judgment, aligning with emerging guidelines for trustworthy and responsible AI in healthcare [4], [10]. This balanced division of responsibility enhances system reliability, user trust, and regulatory acceptability.

Holistic Health Integration as a Differentiating Factor

A defining characteristic of the proposed framework is its integration of physiological and nutritional screening into the mental health support pipeline. Existing AI-based systems predominantly focus on psychological symptom assessment, often neglecting broader health factors that influence anxiety severity and treatment outcomes. Research in nutritional psychiatry has consistently demonstrated that deficiencies in vitamin D, B-complex vitamins, and magnesium, as well as poor overall diet quality, are associated with heightened anxiety and mood dysregulation [11], [13], [14].

By incorporating non-diagnostic screening for such factors, the proposed framework facilitates context-aware personalization and informed health-seeking behavior. This holistic perspective enables the system to address potential contributors to anxiety that would otherwise remain unrecognized in purely psychological models, thereby enhancing the conceptual completeness of the intervention strategy.

Ethical Governance and Trustworthiness

Ethical considerations are integrated into the framework architecture as a fundamental design component. In contrast to baseline models that often treat ethics as an external consideration, the proposed system enforces informed consent, data anonymization, encrypted storage, and bias monitoring by design. These measures align with established ethical guidelines for trustworthy AI and address common concerns related to privacy, accountability, and misuse in AI-driven mental health applications [10].

Furthermore, the inclusion of human oversight enhances transparency and auditability, facilitating clinical review and regulatory compliance. This governance-oriented design strengthens the framework's suitability for real-world deployment in sensitive healthcare contexts.

Scalability, Extensibility, and Deployment Readiness

From a systems engineering perspective, the modular architecture of the proposed framework offers clear advantages over monolithic AI models. Each module—including the conversational interface, assessment engine, screening layer, and oversight mechanism—can be independently validated, updated, or replaced without disrupting the overall system. This modularity supports scalability across diverse user populations and facilitates future integration with wearable sensors, electronic health records, and clinical decision support systems.

By contrast, LLM-only systems often lack clear separation of concerns, complicating validation, auditing, and regulatory approval. The proposed framework's design therefore enhances maintainability and long-term viability in evolving healthcare ecosystems.

Implications for AI-Assisted Mental Health Care

Collectively, these design features reposition AI from an autonomous mental health provider to a clinically assistive component operating within clearly defined ethical and operational boundaries. The proposed framework demonstrates how personalization, safety, holistic health integration, and human supervision can be unified within a single end-to-end architecture. As such, it provides a more responsible and sustainable alternative to existing AI-based mental health systems, particularly for anxiety management.

While empirical validation remains a necessary next step, the conceptual and architectural strengths of the proposed model establish a strong foundation for future clinical evaluation and real-world deployment.

VIII. CONCLUSION AND FUTURE DIRECTIONS

This paper presented a hybrid AI-driven framework for anxiety management that integrates LLM-based conversational support, holistic health screening, and structured human-in-the-loop oversight. By addressing key limitations of existing AI-based mental health systems—namely, limited personalization, lack of ethical governance, and absence of physiological and nutritional considerations—the proposed framework advances toward a more clinically responsible and deployment ready design.

Future work will focus on large-scale empirical validation, integration of wearable and physiological sensing data, refinement of empathetic response modeling, and collaboration with mental health professionals to further enhance clinical applicability and ethical robustness.

REFERENCES

1. World Health Organization, "Depression and anxiety," WHO Fact Sheets, 2023. [Online]. Available: https://www.who.int
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," JMIR Mental Health, vol. 4, no. 2, p. e19, 2017.
3. A. S. Miner, L. Milstein, S. Hancock, J. Hays, S. E. Maier, and J. J. Campbell, "Smartphone-based conversational agents and responses to

questions about mental health," JAMA Internal Medicine, vol. 176, no. 5, pp. 619-625, 2016.
4. A. N. Vaidyam, J. Wisniewski, J. Halamka, K. Kashavan, and J. Torous, "Chatbots and conversational agents in mental health: A review of the psychiatric landscape," npj Digital Medicine, vol. 6, no. 1, pp. 1-10, 2023.
5. G. Baek, C. Cha, and J.-H. Han, "AI chatbots for psychological health for health professionals: A scoping review," JMIR Human Factors, vol. 12, p. e67682, 2025.
6. Q. Zhang, R. Zhang, Y. Xiong, Y. Sui, C. Tong, and F.-H. Lin, "Generative AI mental health chatbots as therapeutic tools: Systematic review and meta-analysis," Journal of Medical Internet Research, vol. 27, p. e78238, 2025.
7. X. Feng, L. Tian, G. W. K. Ho, J. Yorke, and V. Hui, "Effectiveness of AI chatbots in alleviating mental distress and promoting health behaviors: A systematic review and meta-analysis," Journal of Medical Internet Research, vol. 27, p. e79850, 2025.
8. F. O. Kuhlmeier, L. Hanschmann, M. Rabe, S. Lüttke, E.-L. Brakemeier, and A. Maedche, "Evaluating large language model–based mental health chatbots using artificial users and psychotherapist assessment," arXiv preprint arXiv:2503.21540, 2025.
9. A. Arnaiz-Rodriguez, M. Baidal, E. Derner, J. L. Annable, M. Ball, M. Ince, E. Perez Vallejos, and N. Oliver, "Between help and harm: Evaluating mental health crisis handling by large language models," arXiv preprint arXiv:2509.24857, 2025.
10. European Commission, "Ethics guidelines for trustworthy AI," Brussels, Belgium, 2019.
11. F. N. Jacka et al., "Nutritional psychiatry: Current evidence," Proceedings of the Nutrition Society, vol. 76, no. 4, pp. 427-436, 2017.
12. J. Gratch et al., "The Distress Analysis Interview Corpus of human and computer interviews," in Proc. LREC, Portorož, Slovenia, 2014.
13. D. W. Eyles, J. L. Burne, and J. J. McGrath, "Vitamin D effects on brain development, adult brain function and the links between low levels of vitamin D and neuropsychiatric disease," Frontiers in Neuroendocrinology, vol. 34, no. 1, pp. 47-64, 2013.
14. D. R. Smith et al., "Plant-based diets and mental and neurocognitive health outcomes: A systematic review and meta-analysis," Nutrition Reviews, 2025.
15. R. L. Spitzer, K. Kroenke, J. B. 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