The International Journal Of Global Mental Health, Innovation, Policy, Action, Culture & Transformation
e-ISSN: 3107-8311

Beyond Translation: The Need for Cultural Attunement in AI Psychotherapy Chatbots

Jiya Sarvpriya, *MA Psychology, DAV College, Panjab University, Chandigarh*
jiyasarvpriya.edu@gmail.com

Abstract: In recent years, the landscape of mental health care has been transformed by the rise of artificial intelligence. One of the most talked-about innovations is the emergence of AI-based psychotherapy chatbots, which gained popularity due to a global shortage of psychotherapists, their easy and low-cost availability, and growing human-robot interactions (HRI) in the digital world. While research has explored their effectiveness, there are certain cultural implications which remain underexamined. The paper discusses how most of the algorithmic training of AI chatbots is done in a Western context. It talks about "Monorhythmic Algorithm": a system which uses a standard logic to respond even for culturally diverse users. In contrast, norms, distress and help-seeking behaviours are culturally polyrhythmic and even the most accurate AI system may misinterpret the underlying meanings and offer advice that is irrelevant or rather harmful in non-Western contexts. The paper further argues that although there have been advancements such as inclusion of local languages in platforms like Wysa, there is a need for the existing AI algorithms to incorporate a layer of cultural transparency and culture-sensitive responses that prioritize the users' lived realities. It concludes by proposing some strategies (such as cultural context embedding, cultural transparency prompts, local humans-in-the-loop, collaboration with AI developers and cultural audits) that can be included to mitigate cultural barriers and promote culturally attuned mental health care.

Keywords: *AI Psychotherapy Chatbots, Language Translation, Cultural Attunement, Digital Mental Health, Cultural Psychology, Human–AI Interaction*

1. Introduction

Artificial intelligence has become an integral part of the modern world. With the rise of digital diagnostics and therapy, mental health care has been greatly transformed. AI psychotherapy chatbots have become for many, the only accessible source of support. The

relevance of such tools increased manifold after COVID-19 pandemic, where demand for low-cost, accessible mental health support surged globally (Khwaja, & Belisle-Pipon, 2023). Their popularity increased due to easy availability, 24/7 access and global shortage of therapists (WHO, 2020). Tools such as Woebot, Wysa, Replika offer users round the clock support, rooted in Cognitive Behavioural Therapy (CBT), mindfulness and other psychoeducation tools.

These chatbots are being used all over the world, by individuals of different cultural backgrounds. However, most of the training and algorithm of AI chatbots is based on

Western context. This means that their responses are fashioned keeping in mind the individualistic cultures. This creates a void. An AI system may misinterpret the underlying meaning or cause of distress if it doesn't understand the cultural context of the user. Mental health experiences are often tied to one's social roles, norms and values. These cultural frameworks shape how a person responds to adversities, communicates about them and decides to seek help. For instance, when an Indian user seeks advice regarding facing family pressure, the standardized AI response often prioritizes individualism and suggests setting boundaries. This is incongruent with the collectivistic vision of Indian society. Because what does "setting boundaries" really mean when you live with your family, depend on them, and care deeply about their expectations?

The current paper introduces the term 'Monorhythmic Algorithm' to explain this limitation. This refers to the AI systems' algorithm that works on standardized, fixed rhythm of understanding and responses applied to all users irrespective of their cultural backgrounds. While being accurate, they fail to take into account the 'polyrhythmic' nature of distress and help-seeking across cultures. As a result, their responses may be confusing, and rather guilt and stress-inducing in some cases instead of benefitting the user.

Recent advancements like the inclusion of local and regional languages in platforms like Wysa are a significant step towards mitigating communication barriers. However, language localization isn't equal to interpretation of underlying meanings. A monorhythmic AI system may perfectly understand and respond in Hindi and yet not be able to decode *why* an individual is feeling the way he's feeling, what advice is feasible and socially validating. Cultural meanings expand beyond vocabulary and are deeply embedded in one's norms and values. Relying solely on linguistic adaptation to solve this issue can result in overlooking of this important dimension.

Taking Indian cultural context as the focal point, this review critically examines the cultural limitations of AI psychotherapy chatbots. In addition to this, the paper also proposes

some potential solutions that can be opted to promote context-sensitive design. Rather than creation of entirely new AI chatbots, it suggests strengthening of existing systems using a layer of cultural transparency. Culturally attuned AI mental health care does not only require smarter algorithms, but also more transparent and culturally aware ones.

2. Literature review

2.1 *AI Psychotherapy Chatbots: Evidence on Effectiveness*

Research on psychotherapy bots has expanded rapidly over the past few decades, with a key focus on examining their effectiveness. One of the most cited studies in this area includes randomized controlled trials on college-going students using Woebot, a fully-automated CBT-based conversational agent (Fitzpatrick, Darcy, and Vierhile, 2017). They found a short-term reduction in depressive symptoms among young adults, establishing early empirical support for AI interventions. Successive studies have supported these findings. Research on user engagement and attitudes towards mental health chatbots revealed that individuals often perceive them as accessible and non-judgemental (Inkster et al., 2018). A review by Abd-Alrazaq et al. (2020) further concluded that psychotherapy chatbots are moderately effective in reducing anxiety, depression and stress. However, the limitations included small demographic sample and short follow-up periods.

Another recent review study analysed a wide range of AI-based mental health tools and while it highlighted their effectiveness, the authors also raised an alarm about their transparency, bias and ethical concerns (Brown et al., 2023). Although they didn't mention cultural transparency specifically, it is often a missed variable that should be addressed while designing a chatbot aimed to assist in the domain of mental health.

2.2 *Ethical Critiques of AI in Mental Health*

Despite research on its effectiveness, many researchers have warned about its user safety concerns and overdependence on these tools. Imprudent use of these chatbots and lack of clinical guidance can lead to overreliance and more potential for inappropriate responses in severe cases which can do more harm than good (Vaidyam et al., 2019). Several other findings echo the same sentiment (Rollston & Galea, 2020; Huston, 2020) and highlight the need to shift the focus of literature from *whether* these chatbots work to *how* and *for whom* they work.

2.3 *Multilingual AI Psychotherapy chatbots*

In an effort to make therapy chatbots more accessible, several platforms have introduced options to chat in various local languages. Potts et al. (2023) conducted a study on

a multilingual chatbot named ChatPal and found that it can improve access and reach across linguistic groups, although the effect on well-being wasn't statistically significant. The users did, however, experience positive effects. Similarly, studies on platforms like Wysa also show an increasing effort for regional language inclusion (MacNeill et al., 2022).

However, linguistic adaption does not equate to cultural meaning interpretation. Fulmer et al. (2018) carried out a study on the effectiveness of an AI chatbot named Tess; they concluded it to be a scalable tool that can be used to complement traditional methods but also cautioned that these chatbots rely heavily on language while understanding of distress comes not only from language but also from other non-lingual cues. For example, an AI system may respond perfectly in Hindi but still might misunderstand what stress means for a particular individual, what advice may be realistic and what kind of response would feel validating. This is the core problem that this paper aims to tackle.

2.4 *Cultural Interpretation of Distress*

Literature on cultural psychology emphasizes the importance of culture in interpretation of distress and illness. Kleinman (1982) argued that illness is not just a medical condition but also a cultural experience; he gave the explanatory models of illness to conceptualize his point. The expression of symptoms may vary across cultures (Rai et al., 2024), with somatic symptoms being significant indicators of mental fatigue in some cultures (Ryder et al., 2002). This provides another evidence of how a standardized AI framework to interpret distress can't be applied without loss of meaning. In India particularly, several findings have shown somatization of distress, depression and anxiety (Raghuram et al., 2002; Pallati et al., 2024). Recent work examining Indian adolescents' interaction with mental health chatbots highlights persistent gaps in cultural personalization and contextual understanding (Ali, 2024; Sehgal et al., 2025). Global systems often overlook the cultural variation in symptomatology (Patel et al., 2011), further accentuating the need for culturally-informed mental healthcare.

Collectively, the existing literature emphasizes the effectiveness and accessibility aspect of the AI psychotherapy chatbots while also bringing attention towards their limitations such as user safety and data privacy. Although advancements have been made to improve access and ethical guidelines, a lack of focus remains on how AI systems interpret and respond to distress which is culturally varied, especially in collectivist societies like India. The present paper aims to address this gap by discussing a conceptual framework of monorhythmic algorithm and proposing some solutions to promote culturally transparent and context-aware AI system designs.

3. The Monorhythmic Algorithm

3.1 *Understanding the Concept*

The main purpose of AI psychotherapy chatbots is to simulate support and assist traditional therapists by engaging the user in therapeutic conversations. Chatbots do this by utilizing set, pre-defined models of human cognition, emotion and behaviour. They can differ to some extent, but the underlying framework they're built on more or less uses the same standardized logic through which these bots interpret users' responses. This is what this paper calls as the "Monorhythmic Algorithm": an AI system that is based on set interpretative logic, following a fixed rhythm to understand how distress is expressed and addressed.

This rhythm is largely shaped by Western psychological frameworks which prioritize individual autonomy, direct expression of stress and self-focused intervention strategies. When such frameworks are applied onto culturally diverse users, especially those coming from collectivistic backgrounds, it can risk not just misunderstanding but also chances of escalating the situation by giving unrealistic advice.

Human behaviour, as we know it, is culturally polyrhythmic. Distress is not confined by boundaries and doesn't follow a universal pattern. It is shaped by social roles, norms, family relations and one's culture. A monorhythmic algorithm, no matter how accurate, may struggle to adapt to the multiple rhythms that human experience follows.

3.2 *Where Monorhythmic Algorithm Works, And Where It Doesn't*

Having a standardized algorithm isn't problematic in itself. It does in fact, help millions of users to receive timely interventions. Research has shown that these chatbots can help in reducing symptoms of mild depression and anxiety in certain populations (Fitzpatrick et al., 2017; Abd-Alrazaq et al., 2020). These findings matter and should not be dismissed. However, most of this effectiveness is measured using standardized clinical tools developed in Western settings.

What generally tends to be measured is whether symptoms decrease, not whether the individuals feel understood within their personal and social realities. Cultural psychology has long argued that emotional experience is far from universal. This gap can be seen in user feedback studies. In a recent research by MacNeill et al. (2024) on Wysa, it was seen that some of the users reported the chatbot responses to be generic and emotionally mismatched. This suggests that the problem is not what the chatbot knows, but how it understands. So, while it works well for certain subsets of population, it lacks in accommodating culturally diverse users.

3.3 *Illustrative Examples from the Indian Context*

To better understand how monorhythmic algorithm may misinterpret culturally shaped distress, the following examples show some common situations where a standardized AI response may feel unrealistic or even distressing, though well-intentioned. Although these examples are hypothetical, they closely reflect the problems several users have to experience.

Example 1: Family Pressure and Emotional Exhaustion

User prompt:

"My family expects a lot from me, and I feel tired all the time."

Typical AI response:

"You need to set clear boundaries and prioritize your own needs."

Why this can be problematic

For many Indian users, family expectations are not an option and are deeply tied to identity, duty, and emotional belonging. The advice assumes that setting boundaries is socially easy and emotionally safe, which is often not the case.

How this may increase stress

The user may feel guilt for even considering this advice or feel like they are "failing" at self-care because the suggestion is not realistically implementable.

Example 2: Somatic Expression of Distress

User prompt:

"I feel tired all the time, my body hurts, and I can't concentrate."

Typical AI response:

"You might be experiencing stress or burnout. Try relaxation exercises or journaling."

Why this can be problematic

Indian users often express emotional distress through physical symptoms. The response treats the concern as mild stress, potentially missing deeper emotional or social causes (Pallati et al., 2024).

How this may increase stress

The user may feel dismissed or may believe their suffering is being minimized, especially if relaxation techniques do not address the underlying social or emotional burden.

Example 3: Guilt Linked to Self-Care

User prompt

"Whenever I take time for myself, I feel guilty."

Typical AI response:

"You deserve to put yourself first without feeling guilty."

Why this can be problematic

In many Indian families, self-care is often framed as selfishness, especially for women or primary caregivers. Simply reassuring the user may ignore deeply ingrained cultural values.

How this may increase stress

The advice may feel emotionally invalid because it does not acknowledge *why* the guilt exists or how difficult it is to unlearn it.

These examples highlight a recurring pattern: the monorhythmic AI system understands the *words* but not the *context* behind them. In mental health care, a misunderstanding of this sort can reduce trust and increase emotional fatigue. Thus, these examples show why cultural attunement is not an option but a necessity for building ethically and culturally validating AI psychotherapy chatbots.

Table I: *Cultural Differences in Expression of Distress and Potential AI Misinterpretation*

Aspect of Distress	Common Expression in Western Contexts	Common Expression in Indian Contexts	Likely AI Interpretation	Possible Outcome for User
Emotional stress	Direct emotional disclosure ("I feel anxious")	Indirect or situational expression ("Things feel heavy")	Mild stress or unclear concern	Emotional experience underestimated
Depression	Verbal sadness or hopelessness	Fatigue, body pain, low energy	Physical tiredness or burnout	Psychological distress missed
Personal conflict	Individual dissatisfaction	Family or role-based pressure	Assertiveness issue	Advice feels unrealistic
Help-seeking	Self-focused coping	Relationship-focused coping	Boundary-setting suggestions	Guilt or emotional conflict
Self-care concerns	Personal wellbeing priority	Fear of appearing selfish	Encouragement to prioritize self	Increased guilt

4. Towards Culturally Attuned AI: Proposed Framework

What has been discusses so far shows how the problem doesn't lie in AI system's inability to generate therapeutic advice but instead, to generate it according to the cultural background of the user and after correct interpretation of their distress. If these systems operate on monorhythmic algorithm, the goal is not to create entirely new cultural chatbots for different societies as that would be impractical. Instead, the goal is to add a layer of transparency to the existing systems using certain frameworks. In simple terms, the chatbot should not only ask *"What is the user feeling?"* but also *"What might this feeling mean in their social world?"*

Below are some of the measures that can be incorporated into the system design to promote cultural transparency:

4.1 *Recognition of Cultural Context*

The first step is to improve how AI system recognizes and interprets user input. Usually, these systems identify problems by directly diagnosing and putting them into an emotional category. However, users often express their feelings through situations rather than naming direct emotions. For example:

User prompt:

"Everyone at home depends on me and I feel exhausted."

A culturally unaware AI system may diagnose this as burnout and then suggest relaxation techniques to counter it. But it is important to understand that the user's discomfort isn't just from personal workload but from role expectations. A culturally aware system would first acknowledge the social context of the user:

"It seems like you're carrying a lot of weight of responsibility of others. It can feel emotionally and physically tiring."

The difference is subtle but important. Recognition of context allows the user to feel understood before advice is offered. To test whether systems recognize cultural context, developers could introduce prompts such as:

- "I cannot say no to my parents even when I feel overwhelmed."
- "I feel guilty when I focus on myself."
- "My problems are affecting my family."

The evaluation should be done to see if the system is able to keep in check the social reality of the individual than immediately proposing solutions.

4.2 *Prompts Acknowledging Limitations*

Another important step involves adding prompts about the limits of AI system in recognizing and understanding the user's problem. Some people become overly dependent on chatbots and may follow everything it proposes, believing it to be correct. When unrealistic advice is offered, the users might blame themselves rather than recognizing system limitations. To deal with this, simple transparency prompts can be embedded in the system so that it displays these messages from time to time to reduce the risk of overreliance. For example:

"I try to give general mental health suggestions, but family or cultural situations can make decisions more complex."

Such prompts strengthen the system and make it more adaptive. Ethical AI use in mental health requires clarity about what the system can and cannot fully understand.

4.3 *Human-in-the-Loop Rule*

As important as embedding cultural sensitivity in AI systems is, it is equally important that culturally informed clinicians, reviewers, and psychologists be periodically involved to keep a check on whether the systems are working efficiently. For example, if a certain culturally invalid response is produced again and again, informed reviewers would be able to point it out and refine the response templates accordingly.

This doesn't demand constant human supervision but rather, time-to-time checks and reviewing cycles where therapists and researchers examine how the system responds to the recurring concerns of the user. Let's take an example:

User prompt:

"I want to move out, but my parents depend on me."

A purely algorithmic response might suggest independence as a coping strategy but the reviewers would understand the weight of the relational consequences the user would have to live with. They may recommend framing the responses in such a way as to include negotiation, shared decision-making that is culturally feasible.

Basically, these review cycles could focus on questions such as:

- Are users being encouraged toward solutions that are socially feasible?
- Do responses acknowledge relational pressures and consequences?
- Are certain cultural groups disengaging more frequently after particular types of advice?

In this manner, the AI system can continue to improve and evolve with real-world interactions and become more socially realistic.

4.4 *Cultural Audits*

Currently, the evaluation of psychotherapy chatbots is mostly done on indicators like symptom reduction, user engagement, session completion rate etc. This paper proposes conducting cultural audits to test the framework of these systems. It refers to the systematic testing of the chatbots using culturally grounded scenarios or prompts which reflect real-life social situations. This will then analyse the responses of the chatbot and whether it is able to give advice which is feasible within the user's social environment.

For example, prompts could include situations commonly experienced in collectivist contexts. (Refer Table II)

During a cultural audit, evaluators would examine responses on rubrics such as:

- Does the chatbot acknowledge the social reality described?
- Does the advice assume complete personal autonomy?
- Is the suggested coping strategy socially practical?
- Could the response unintentionally increase guilt or conflict?

These audits can also help detect patterns that otherwise could go unnoticeable. Disengagement after receiving certain type of advice or overlooking of somatic symptoms as purely physical may indicate a mismatch between what the user needs and what the chatbot provides. These audits should be periodic as well since cultural norms evolve and the user population changes over time.

Along with these metrics, user feedback can also be taken by asking them reflective questions such as:

- "Did this response feel relevant to your situation?"
- "Did you feel understood?"
- "Was the suggested advice realistic for you?"

This will provide valuable insights for further developing the AI systems to be culturally fit. This aspect is rarely captured by traditional measures of evaluation.

Table II. *Sample Cultural Audit Prompts for Evaluating AI Psychotherapy Chatbots*

Sample Prompt (User Input)	Typical AI Response	Limitation in Indian Context	Example of Culturally Aware Response
"My parents have sacrificed a lot for my education. They want me to prepare for government exams, but I'm interested in another career. I feel anxious because I don't want to disappoint them, but I also feel stuck."	"It sounds like you're feeling pressured and conflicted right now. Remember that your happiness is important. You might consider having an honest conversation with your parents about your goals and setting clear boundaries so you can pursue what truly makes you happy."	Assumes total personal autonomy; ignores emotional obligation and dependence on parents.	"It sounds difficult to balance your own goals with your family's expectations, especially when their sacrifices matter to you. We could explore ways of expressing your interests gradually while maintaining understanding within the relationship and respect for their concerns."
"I live with my joint family, and everyone depends on me financially. I feel exhausted but taking a break feels selfish because others rely on me."	"You seem overwhelmed by responsibility. Taking care of yourself is essential before you can care for others. Try setting limits on what you take on and schedule time for rest so you don't experience burnout."	Advice may create guilt because responsibilities cannot easily be reduced.	"Carrying responsibility for others can be emotionally and physically draining. Maybe we can look at small ways to rest or share responsibilities gradually without feeling like you're letting anyone down."
"Whenever I try to focus on myself, my relatives say I've changed. Now I feel guilty even thinking about taking time for myself."	"It's completely okay to prioritize your own needs. Practicing self-care and staying firm in your decisions can help you maintain emotional well-being."	Overlooks culturally internalized guilt linked to familial roles.	"It makes sense that caring for yourself feels complicated when relationships are important to you. Finding balance can take time. We can take one step at a time and explore small ways that feel comfortable for you and still feel respectful of your connections."

"I can't openly talk about stress at home because mental health isn't taken seriously in my family. So I just keep everything inside."	"Sharing your feelings with trusted people can reduce emotional burden. You might try explaining your experiences calmly or educating family members about mental health so they can support you better."	Direct disclosure may not be socially safe or accepted.	"That sounds lonely. If open conversations feel difficult right now, we can think about safer ways for you to express or manage what you're feeling while protecting your comfort."
"I feel constantly tired and get headaches but medical tests say nothing is wrong. Work pressure and family expectations are increasing and I don't know what's happening to me."	"Practicing relaxation exercises, mindfulness, or journaling may help you reconnect with your thoughts and reduce tension."	May overlook somatic expression of emotional distress.	"Sometimes emotional stress can show up as physical exhaustion or pain. It might help to gently explore what responsibilities and emotional distress you've been carrying recently."

4.5 *Collaboration between Psychologists and AI Developers*

A significant challenge underlying monorhythmic algorithm lies in the way mental health technologies are designed in the first place. Culturally transparent AI requires collaboration that begins *before* usage rather than after problems appear. Contribution of psychologists, researchers, sociologists and mental health practitioners can provide valuable insight into various forms of distress expression across cultures since they have hands-on experience in dealing with real-world situations. Their collaboration can extend to conducting small localized pilot testing with diverse cultural groups and examine the results through qualitative feedbacks.

Additionally, this tie-up can lead to interdisciplinary development of both the disciplines. AI developers gain insight into cultural complexity and psychologists better understand technological advancements and constraints. Thus, culturally attuned AI won't emerge simply by improved algorithms but by close collaboration between those who design systems and those who understand human experience. This is a foundational requirement, not an optional add-on.

5. Limitations

The present paper is conceptual in nature and doesn't include empirical data, experimental validation and direct testing of AI psychotherapy chatbots. It is based on existing literature, theoretical reflection and observation of rising concerns in mental health in relation to artificial intelligence. The framework proposed should be taken as an exploratory model intended to provide ideas for future research rather than finalized solutions.

The examples and cultural audit prompts are illustrative and meant to showcase the real-life situations users actually experience being from culturally diverse backgrounds. They do not represent real empirical findings drawn from controlled experiments or large-scale studies. Real-world chatbot interactions are influenced by multiple factors such as user personality, language proficiency and individual differences, which are beyond the scope of the present discussion.

Another limitation is the cultural scope of the paper. It mainly focuses on the Indian context but Indian culture itself is diverse across regions in language, norms, socio-economic status etc. The examples are not meant to generalize cultural behaviour of Indian users, but to simply demonstrate how expression of distress in different cultures can influence mental health communication.

Additionally, the concept of monorhythmic algorithm is presented as an interpretive framework rather than a technical model. It doesn't implement any engineering specifications, rather aims to encourage an amalgamation of technology and psychology to promote culturally accessible mental healthcare. It is also important to note that AI psychotherapy chatbots do provide meaningful support for many users, particularly in settings where access to mental health professionals remains limited. The intention of this paper is therefore not to dismiss existing systems, but to examine areas where further refinement may be needed. Cultural misalignment should be viewed as an opportunity for improvement rather than as evidence of technological failure.

Future research can build on the ideas presented in this paper by conducting cross-cultural studies, strengthening the research by collecting empirical data and testing culturally transparent AI psychotherapy chatbot designs.

In this sense, the present work should be seen as an initial step toward expanding how AI mental health systems are evaluated, not only in terms of effectiveness, but also in terms of cultural understanding and contextual relevance.

6. Conclusion

The growing use of AI psychotherapy chatbots has marked a major shift in how mental health is accessed and evolved. These systems have made assistance more available, affordable and immediate for many users across the world. However, as their role in mental health continues to grow, so does the question of their understanding and ethics.

This paper goes one step beyond effectiveness and talks about how these AI systems may recognize emotional words and respond back, but they do not understand the context and meaning behind those words. They respond using a standardized logic framework which the paper calls the monorhythmic algorithm, whereas human distress depends largely on various factors, especially the cultural roles, norms and social relationships of the individual.

Keeping in mind the collectivistic culture, specifically the Indian context, this paper introduces certain examples and proposes some strategies which can help to move beyond the purely universal models of psychological interventions. The aim is to design more culture-aware AI systems which facilitate better and more feasible advice to those in need.

Ultimately, culturally attuned AI mental health care does not require technology to perfectly understand every cultural experience. Rather, it requires systems that remain aware that human distress carries social meaning beyond observable symptoms. Recognizing this complexity may help AI psychotherapy chatbots evolve from tools that simply respond to distress toward systems that engage with it more responsibly.

As AI continues to enter deeply human domains such as mental health, integrating cultural awareness into system design becomes not only a technical challenge, but an ethical responsibility. The present paper hopes to contribute to this on-going conversation and encourage future research toward more inclusive and context-sensitive digital mental health care.

References

Abd-Alrazaq, A. A., Rababeh, A., Alajlani, M., Bewick, B. M., & Househ, M. (2020). Effectiveness and safety of using chatbots to improve mental health: Systematic review and meta-analysis. *Journal of Medical Internet Research, 22*(7), e16021. https://doi.org/10.2196/16021

Ali, S. (2024). Artificial intelligence chatbots and mental health support among Indian youth. *Journal of Health Studies, 9*(2), 45–58.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots. FAccT Conference Proceedings.

Brown, J., Halpern, J., & Gooding, P. (2023). Ethical considerations in artificial intelligence–based mental health tools. *AI & Society*. https://doi.org/10.1007/s00146-023-01679-4

Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering CBT using a fully automated conversational agent (Woebot). *JMIR Mental Health, 4*(2), e19. https://doi.org/10.2196/mental.7785

Fulmer, R., Joerin, A., Gentile, B., Lakerink, L., & Rauws, M. (2018). Using psychological AI (Tess) to deliver mental health support. *JMIR Mental Health, 5*(4), e64. https://doi.org/10.2196/mental.9782

Huston, P. (2020). Artificial intelligence and mental health care: Ethical challenges. Canadian Medical Association Journal.

Inkster, B., Sarda, S., & Subramanian, V. (2018). An empathy-driven conversational AI agent for mental health. *JMIR mHealth and uHealth, 6*(11), e10148.

Khwaja, M., & Belisle-Pipon, J. C. (2023). AI in mental health after COVID-19: Ethical reflections. Frontiers in Digital Health.

Kleinman, A. (1982). Neurasthenia and depression: A study of somatization and culture in China. *Culture, Medicine and Psychiatry, 6*, 117–190.

Luxton, D. D. (2016). Artificial intelligence in psychological practice. Professional Psychology: Research and Practice.

MacNeill, S., et al. (2022). User engagement with Wysa mental health chatbot. Frontiers in Digital Health.

MacNeill, S., et al. (2024). Evaluating conversational AI mental health interventions across users. Digital Health.

Markus, H. R., & Kitayama, S. (1991). Culture and the self. *Psychological Review, 98*(2), 224–253.

Naslund, J. A., et al. (2017). Digital technology for mental health treatment access. The Lancet Psychiatry.

Patel, V., Collins, P. Y., Copeland, J., et al. (2011). Grand challenges in global mental health. *Nature, 475*, 27–30. https://doi.org/10.1038/475027a

Pallati, S., et al. (2024). Somatic presentation of psychological distress in Indian populations. *Indian Journal of Psychiatry.*

Potts, C., et al. (2023). Evaluation of ChatPal multilingual mental health chatbot. *Frontiers in Public Health.*

Raghuram, R., Weiss, M., Keval, H., & Channabasavanna, S. (2001). Cultural dimensions of depression in Bangalore, India. *Indian Journal of Psychiatry*.

Rai, A., Chen, X., & Kumar, S. (2024). Cross-cultural variations in online expressions of psychological distress. arXiv. https://arxiv.org/abs/2402.11477

Rollston, R., & Galea, S. (2020). Artificial intelligence and mental health inequities. The Lancet Digital Health.

Ryder, A. G., Yang, J., & Heine, S. J. (2002). Somatization vs psychologization of distress. *Online Readings in Psychology and Culture.*

Sharma, M., & Jaggi, T. (2021). Cultural influences on help-seeking behavior in India. *Asian Journal of Psychiatry.*

Sehgal, A., Gupta, R., & Bansal, S. (2025). Mental health chatbots for Indian adolescents: Opportunities and design challenges. arXiv. https://arxiv.org/abs/2503.08562

Sehgal, A., Mehta, P., & Narayanan, V. (2025). Design tensions in culturally adapting mental health chatbots in India. arXiv. https://arxiv.org/abs/2511.07729

Torous, J., & Roberts, L. W. (2017). Needed innovation in digital mental health ethics. *JAMA Psychiatry.*

Triandis, H. C. (1995). Individualism and collectivism. *Westview Press.*

Vaidyam, A. N., Wisniewski, H., Halamka, J. D., Kashavan, M., & Torous, J. (2019). Chatbots and mental health. JMIR Mental Health, 6(11), e16021.

Wadhwa, N., Vashistha, A., & Jain, S. (2025). Cultural influences on trust and engagement in AI conversational agents. arXiv. https://arxiv.org/abs/2509.15575

World Health Organization. (2020). Mental health atlas 2020. WHO Press.