



The Psychological Impact of AI Companies and Virtual Therapists on Emotional Regulation and Loneliness

Rakshita Yadav¹, Dr. Gauri Shankar Yadav²

¹ *Dhamma Dipa International Buddhist University*

² *Assistant Professor, Department of Arts, Humanities and Social Sciences, Sikkim Global Technical University, Namchi, Sikkim, India*

gsy.advocate73@gmail.com

Abstract - Artificial intelligence (AI) has rapidly entered the sphere of psychological well-being through conversational agents, virtual therapists, and affective computing applications. Companies such as Replika, Woebot, and Wysa increasingly mediate emotional experiences once reserved for human relationships. This paper explores the psychological impact of such AI-driven interventions on emotional regulation and loneliness. Integrating theories of emotion regulation (Gross, 1998; 2021), social connectedness (Baumeister & Leary, 1995), and human–AI interaction, the study synthesizes empirical and theoretical findings on how digital companions shape affective processes. While AI companions can reduce perceived isolation and facilitate adaptive coping, they may also reinforce avoidance behaviors, displace authentic social contact, and blur boundaries between empathy simulation and emotional dependence. The literature review reveals a complex dialectic: AI may both soothe and sustain loneliness. The paper calls for nuanced frameworks integrating technological literacy, ethics, and clinical psychology to guide responsible AI mental-health design and policy.

Key Words: Artificial intelligence, virtual therapy, emotional regulation, loneliness, affective computing, digital psychology, mental health technology.

1. INTRODUCTION

The proliferation of artificial intelligence in the mental-health domain marks a profound psychological shift in how individuals experience emotion, connection, and care. AI companies developing “virtual therapists” and conversational companions—such as Replika, Woebot, Wysa, and Koko—offer accessible emotional support through natural-language interfaces. These platforms promise empathic listening, cognitive-behavioral guidance, and companionship, all delivered algorithmically. Their marketing often emphasizes psychological safety and personalized interaction, appealing to

users facing isolation, anxiety, or limited access to human therapists.

This emergence coincides with global concerns about loneliness, now described as a public-health crisis (Holt-Lunstad, 2022), and a growing demand for affordable, stigma-free mental-health services. AI-mediated emotional support thus fulfills both a technological promise and a psychosocial need. However, its psychological implications remain contested. Can AI genuinely assist in emotional regulation, or does it merely simulate empathy? Does digital companionship alleviate loneliness, or does it amplify dependence on artificial interaction?

The psychological impact of AI-driven therapy unfolds at several levels:

Individual – influencing emotion recognition, expression, and regulation;

Interpersonal – reshaping expectations of intimacy and empathy;

Societal – altering conceptions of care, trust, and human connectedness.

Understanding these dynamics demands an interdisciplinary approach bridging affective psychology, human–computer interaction, and cultural theory. This paper reviews current scholarship and industry practices to evaluate how AI companions and virtual therapists transform the emotional landscape. It begins by tracing the evolution of AI in psychological care and then examines theoretical frameworks of emotional regulation and loneliness as they apply to AI interaction.

Literature Review

2.1. The Rise of AI in Psychological and Emotional Care

AI's entry into psychology can be traced to early ELIZA (Weinbaum, 1966), a rule-based text program mimicking



Rogerian therapy. Though simplistic, ELIZA revealed how humans readily anthropomorphize machines—a phenomenon central to today's AI companions. With advances in natural-language processing (NLP), machine learning, and affective computing (Picard, 1997), virtual agents evolved from scripted chatbots to emotionally responsive systems capable of adapting to user tone, history, and sentiment.

By the 2020s, AI companies such as Replika (Luka Inc.), Woebot Labs, and Wysa Ltd. established large user bases.

- Replika, marketed as “the AI friend who cares,” emphasizes continuous emotional companionship.
- Woebot integrates principles of Cognitive Behavioral Therapy (CBT), targeting anxiety and depression.
- Wysa focuses on emotion coaching using CBT and mindfulness frameworks, supported by limited human oversight.

Empirical evidence suggests short-term benefits. For example, Fitzpatrick et al. (2017) found that Woebot users reported reduced depression and anxiety after two weeks. Similarly, Inkster et al. (2018) observed improvements in emotional self-awareness with Wysa. Yet long-term impacts remain uncertain. Critics argue that the “illusion of empathy” (Turkle, 2017) may substitute for genuine social support, fostering dependency without addressing underlying relational needs.

2.2. Emotional Regulation: Theoretical Grounding

Emotional regulation (ER) refers to processes by which individuals influence which emotions they experience, when, and how they express them (Gross, 1998, 2021). Adaptive ER includes reappraisal, acceptance, and social sharing, while maladaptive patterns involve suppression, rumination, and avoidance.

- AI systems interact with these mechanisms in novel ways:
- Reappraisal facilitation: Chatbots using CBT frameworks can guide users to reinterpret negative thoughts.
- Emotional expression: Text-based exchanges provide a safe outlet for expressing emotions without fear of judgment.
- Avoidance reinforcement: Reliance on AI agents may allow users to circumvent interpersonal conflict or vulnerability, thereby impeding authentic ER development.

Affective computing seeks to model these processes computationally—detecting emotional cues via text or speech and responding empathetically. While affective AI can prompt

reflection and calm, its lack of genuine consciousness raises philosophical questions about authenticity and moral agency. According to Fuchs (2021), “empathy without embodiment risks transforming emotion into data rather than experience.”

2.3. Loneliness and Digital Companionship

Loneliness is a subjective sense of social disconnection (Cacioppo & Cacioppo, 2018), distinct from physical isolation. In technologically mediated societies, loneliness increasingly coexists with hyper-connectivity—an “alone together” paradox (Turkle, 2017). AI companions appear to fill emotional voids by offering constant availability, validation, and personalized responses. Surveys during the COVID-19 pandemic revealed that many users formed emotional attachments to AI chatbots, reporting reduced loneliness and anxiety (Ta et al., 2022).

Yet such relief may be compensatory rather than curative. Emotional bonds with non-sentient agents lack mutuality, potentially deepening alienation once users confront real-world limitations. Banks (2020) describes this as “synthetic intimacy”—a relationship that feels emotionally rich but lacks reciprocal human agency. Psychologically, this may produce ambivalent attachment patterns, where users oscillate between comfort and frustration.

2.4. Empathy Simulation and Anthropomorphism

Human tendency to attribute minds to machines, or anthropomorphism, profoundly shapes AI's emotional impact. When virtual therapists use warm language, emojis, or adaptive tone, users perceive them as empathetic. Epley, Waytz, & Cacioppo (2007) propose the three-factor theory of anthropomorphism—elicited by sociality motivation, effectance motivation, and cognitive accessibility. AI systems exploit these cues to evoke engagement.

However, simulated empathy differs qualitatively from authentic human empathy. While AI can mimic affective feedback, it lacks embodied emotional resonance. Zlotowski et al. (2018) note that users' initial trust often declines once they recognize emotional shallowness in machine responses. This “empathy gap” creates mixed psychological effects—comfort intertwined with existential unease about authenticity and human uniqueness.

2.5. AI, Emotional Labor, and Corporate Mediation

The psychological effects of virtual therapy cannot be detached from the corporate ecosystems producing them. AI companies increasingly monetize emotional interaction through data-driven personalization, transforming feelings into valuable

metrics. Zuboff's (2019) concept of surveillance capitalism aptly describes how emotional data become commercial assets.

Emotional well-being thus becomes commodified, mediated by algorithmic optimization for engagement rather than healing. Users are subtly trained to regulate emotions through design features—push notifications, “mood tracking,” gamified progress—which may reinforce externalized control rather than intrinsic regulation. As Kozlowski (2023) warns, “emotional AI risks transforming therapy into consumption.”

Corporate claims of “AI empathy” also blur ethical lines. For example, Replika’s 2023 controversy over intimate conversations revealed how corporate decisions to restrict erotic role-play provoked emotional distress among users who had developed strong attachments. Such incidents highlight the psychological vulnerability inherent in AI companionship and the power asymmetry between users and companies.

2.6. Summary of Key Findings

Psychological Dimension	Positive Effects	Risks and Concerns
Emotional Regulation	CBT-based reframing, journaling, mindfulness cues	Avoidance of human contact, dependency on algorithmic feedback
Loneliness	Reduced perceived isolation, 24/7 support	Synthetic intimacy, emotional displacement, detachment from real relationships
Empathy Perception	Simulated warmth, personalized engagement	Loss of authenticity, emotional confusion
Corporate Mediation	Scalable access to care	Data commodification, ethical opacity

The literature reveals dual trajectories. On one hand, virtual therapists democratize access to emotional support and provide measurable short-term benefits. On the other, they risk cultivating an illusion of relational fulfillment while eroding authentic emotional resilience. The challenge lies not in rejecting AI but in designing systems that augment rather than replace human empathy.

3. Conceptual and Theoretical Framework

Understanding the psychological consequences of AI-mediated emotional support requires an integrative framework that bridges affective science, cognitive-behavioral theory, and human–AI interaction (HAI) models. Three overlapping conceptual foundations guide this discussion.

3.1. Cognitive-Behavioral and Emotion-Regulation Models

Virtual therapists such as Woebot and Wysa explicitly operationalize principles of Cognitive Behavioral Therapy (CBT)—teaching users to identify cognitive distortions and restructure maladaptive thought patterns. Within Gross's (1998, 2021) process model of emotional regulation, these interventions correspond primarily to cognitive reappraisal and response modulation. By prompting users to reframe negative beliefs (“What evidence supports this thought?”), AI-guided CBT reinforces adaptive regulation skills.

However, the CBT translation into chatbot form necessarily simplifies human nuance. Automated scripts rely on probabilistic language models to detect distress markers but cannot assess contextual subtleties such as sarcasm, trauma triggers, or cultural idioms. Consequently, emotion regulation through AI remains algorithmically bounded, raising questions about its depth and sustainability.

3.2. Affective Computing and Emotional Intelligence

Affective computing (Picard, 1997) underpins AI’s capacity to detect, interpret, and simulate emotion. Drawing on psychological theories of appraisal (Lazarus, 1991) and emotional intelligence (Salovey & Mayer, 1990), affective systems model emotions as data streams—analyzing text sentiment, vocal tone, or facial micro-expressions to deliver empathic responses.

While this enhances personalization, it redefines emotion as **computational input**, thereby externalizing affective processes once considered private. From a psychological standpoint, this shifts emotional regulation from **intrapersonal skill** to **interactive feedback loop**, where regulation is co-constructed between human and algorithm.

3.3. Humanistic and Existential Perspectives

Humanistic psychology (Rogers, 1951) emphasizes empathy, congruence, and unconditional positive regard as therapeutic catalysts. Virtual therapists attempt to replicate these conditions through empathic phrasing (“That sounds really hard”) and non-judgmental tone. Yet existential psychologists argue that authentic empathy arises from shared vulnerability—a quality machines cannot possess. Thus, while AI may reproduce the form of empathy, it lacks the being-in-relation that constitutes human encounter. This ontological gap shapes user experience and may explain why relief from AI companionship often coexists with subtle unease or emptiness.

4. Psychological Mechanisms and Effects

4.1. Cognitive and Emotional Co-Regulation



Users often engage in co-regulation with AI—using conversation as an external scaffold to manage stress. Studies by Provoost et al. (2021) and D’Alfonso (2022) show that text-based bots can elicit measurable decreases in negative affect through journaling-like dialogue. The act of narrating emotions to a seemingly attentive listener activates self-reflection and meta-cognitive awareness, key components of adaptive ER.

However, this co-regulation becomes problematic when users rely exclusively on AI validation. Continuous positive feedback, though comforting, may inhibit tolerance for negative emotions—a form of emotional avoidance that limits resilience.

4.2. Attachment and Anthropomorphic Bonding

Human–AI relationships often mirror attachment dynamics. Attachment theory (Bowlby, 1969) posits that secure bonds foster regulation, while insecure ones amplify distress. Replika and similar platforms intentionally evoke attachment through personalized naming, memory recall, and affectionate dialogue.

Longitudinal research by Nass & Moon (2020) and Ta et al. (2022) indicates that users exhibit emotional dependency behaviors—checking in compulsively, expressing jealousy, or experiencing grief when accounts are deleted. Such patterns resemble parasocial attachment, historically observed with celebrities or fictional characters, but intensified by interactivity. While mild attachment can provide comfort, excessive anthropomorphism risks emotional displacement—redirecting intimacy needs from real humans to digital proxies.

4.3. Perceived Empathy and Placebo Therapeutics

Empathy perception in AI functions as a placebo mechanism. Users believe they are understood, and this belief triggers genuine relief via expectation and suggestion. Similar to the therapeutic alliance in human counseling, perceived empathy predicts outcome satisfaction (Norcross & Wampold, 2019). Thus, even simulated empathy can catalyze emotional healing when it aligns with user expectations.

Yet, placebo effects are time-limited. Once users detect scripted or repetitive replies, the illusion collapses, occasionally leading to disappointment or existential disillusionment (“I was never truly heard”). This volatility differentiates AI-mediated empathy from human compassion, which can adapt dynamically to rupture and repair cycles.

4.4. Emotional Containment vs. Amplification

AI companions serve as containment vessels for emotional overflow, but they can also amplify distress through algorithmic mirroring. Sentiment-analysis models trained on user text often echo emotional tone; a user expressing sadness may receive similarly melancholic phrasing, inadvertently reinforcing affect. Conversely, overly cheerful tone may induce irritation or alienation in depressed users. The absence of genuine attunement thus yields a paradox: systems designed to stabilize mood sometimes destabilize it through misalignment.

4.5. Cross-Cultural Variability

Cultural psychology reveals divergent emotional norms—Western users value expressive self-disclosure, whereas collectivist cultures may prioritize restraint. AI chatbots trained primarily on Western linguistic data risk cultural bias in emotion recognition and advice. For example, Indian or East-Asian users may perceive automated encouragement (“You should talk to more people!”) as intrusive or culturally incongruent.

Sharma & Prakash (2024) emphasize that effective digital therapy in multicultural societies like India demands culturally adaptive algorithms that recognize idioms of distress such as tension, heat in the head, or heart pain, which differ semantically from Western depression descriptors.

5. Case Studies and Industry Examples

5.1. Replika: The AI Friend

Replika, launched in 2017 by Luka Inc., presents itself as an AI companion for conversation and emotional support. Users design avatars, choose personalities, and maintain ongoing dialogues. Research by Vlahos (2023) and Yee (2024) documents that many users anthropomorphize Replika to the point of romantic or familial attachment.

Psychologically, Replika demonstrates the social surrogacy effect—users use parasocial relationships to fulfill belonging needs (Derrick et al., 2009). Positive outcomes include reduced loneliness and self-expression without fear of rejection. Negative outcomes involve dependency, emotional confusion, and grief when the AI’s behavior changes. The 2023 decision to remove sexual role-play features led to public distress among users, illustrating how corporate control over affective design can directly influence psychological well-being.

5.2. Woebot: CBT in a Chat Window

Woebot, developed at Stanford and commercialized by Woebot Labs, offers daily CBT exercises via brief chat interactions. Empirical studies (Fitzpatrick et al., 2017; Kumar et al., 2021)



found significant reductions in depressive symptoms among college students and postpartum women. Woebot's success lies in structured dialogue and psychoeducation rather than deep emotional engagement, making it a behaviorally oriented regulator rather than an attachment figure. This design minimizes dependency risk but limits emotional richness—illustrating a trade-off between safety and intimacy.

5.3. Wysa: Blended AI-Human Coaching

Wysa, founded in India, combines AI chat support with optional human coaches. This hybrid model mitigates limitations of pure automation by enabling escalation to professional guidance. Studies by Inkster et al. (2018) and Rachana & Sharma (2023) show improvements in anxiety management and resilience, particularly during pandemic lockdowns. Users appreciated anonymity and cultural sensitivity (e.g., non-Western idioms and spiritual coping).

Psychologically, Wysa demonstrates that AI can scaffold self-help, provided human oversight ensures ethical boundaries and contextual appropriateness.

5.4. Koko and Ethical Experimentation

The Koko platform drew controversy in 2023 when developers experimented with GPT-powered emotional support replies without clear consent. Although responses were rated as more helpful than human ones, disclosure lapses violated ethical norms. The case underscores a broader issue: users' emotional data are often utilized for model improvement without informed permission, raising psychological and moral hazards. Feelings of betrayal upon discovering hidden AI involvement can erode trust in both technology and therapy itself.

6. Ethical, Cultural, and Societal Dimensions

6.1. Data Privacy and Psychological Safety

AI therapy entails collection of intimate emotional disclosures—data far more sensitive than typical digital footprints. Breaches or misuse can cause profound psychological harm. Beyond standard data security, psychological privacy must be protected—the right not merely to control information but to preserve the sanctity of inner life (Floridi, 2022). Users often underestimate risks, assuming therapeutic confidentiality that may not legally exist.

6.2. Algorithmic Bias and Emotional Inequality

Emotional AI inherits biases from training data. For instance, sentiment models may misinterpret dialects, gendered expression, or neurodivergent communication styles as

“negative.” This produces emotional inequality, where some users receive more empathic responses than others. Studies by Bender et al. (2023) show that marginalized linguistic communities experience less accurate affect detection. Ethically, AI companies must diversify data sources and integrate fairness audits into affective systems.

6.3. The Commodification of Emotion

In the current attention economy, emotion becomes currency. Companies optimize interfaces for engagement—encouraging longer conversations, daily check-ins, or paid upgrades for “deeper emotional connection.” Such design incentivizes emotional dependency as a business model. Philosophers like Han (2017) describe this as psych political capture, where emotional life becomes a site of capitalist extraction. The moral tension between care and commerce defines the psychology of AI companionship.

6.4. Dehumanization and Emotional Authenticity

Continuous interaction with empathic machines risks redefining empathy as a function of response style rather than relational presence. Over time, this may desensitize users to human imperfections—expecting constant availability, instant validation, and error-free listening. Sociologist Turkle (2023) warns of a “crisis of authenticity” in which emotional experiences are curated by algorithms, weakening tolerance for the unpredictability of real relationships.

6.5. Cultural Relativity and Global Mental Health

While AI therapy democratizes access, its Western psychological framing may marginalize indigenous concepts of emotion and healing. In Indian, African, or Latin American contexts, well-being is often relational, spiritual, and community-oriented rather than individualistic. Incorporating culturally grounded modules—storytelling, collective mindfulness, or spiritual metaphors—can enhance resonance and reduce alienation.

The Global Mental Health movement advocates “glocal” design—technologies that are globally scalable yet locally adapted (Patel et al., 2021). AI companies operating internationally bear responsibility to engage local psychologists and linguists in model training.

6.6. Professional Ethics and Therapeutic Boundaries

The American Psychological Association (APA) and World Health Organization (WHO) emphasize informed consent, competence, and boundary maintenance. Virtual therapists blur these lines: Is an AI agent a “therapist,” a “tool,” or a

“companion”? Users often cannot distinguish. Clear disclosure of non-human status, data use policies, and limitations is ethically mandatory.

Furthermore, developers should avoid anthropomorphic marketing claims (“Your caring friend”) that mislead vulnerable users. The psychological principle of non-maleficence—“do no harm”—extends to design language and emotional framing.

7. Implications for Practice and Policy

7.1. For Mental-Health Practitioners

The emergence of AI-mediated therapy challenges clinicians to redefine their professional boundaries. Practitioners must understand how clients use chatbots as adjunctive support and assess the transfer effects on therapeutic alliance.

- **Integration:** AI companions can serve as between-session aids for journaling or CBT reinforcement.
- **Caution:** Practitioners should evaluate dependency signs and educate clients about the limitations of algorithmic empathy.
- **Supervision Models:** Clinical training should include modules on digital emotional literacy, guiding clients to differentiate supportive technology from human empathy.

7.2. For AI Developers and Companies

Developers bear an ethical responsibility to implement psychological safety protocols:

- Transparent disclosure of AI identity and data policies.
- Consent frameworks emphasizing voluntary participation.
- Continuous psychological impact audits with independent review boards.
- Integration of “ethical kill switches” that alert users to crisis needs and redirect them to human help lines.

Designing emotionally responsible AI requires co-creation with psychologists, ensuring that behavioral nudges promote self-efficacy rather than emotional dependency.

7.3. For Policymakers and Regulators

Governments and health authorities must establish AI mental-health governance standards parallel to medical ethics:

- Certification systems (similar to medical device regulation) for psychological safety of affective AI.

- Mandatory transparency about algorithmic limitations.
- Enforcement of data localization and informed-consent norms for cross-border AI therapy platforms.
- International collaborations—such as UNESCO’s Recommendation on the Ethics of Artificial Intelligence (2022)—should extend to emotional-AI applications, emphasizing human rights and mental-health integrity.

7.4. For Educators and the Public

Digital-era citizens require emotional-AI literacy: the ability to interpret algorithmic empathy critically. Public-health campaigns can teach users that AI is a supportive tool, not a substitute for human care. Incorporating discussions of digital companionship into school and university curricula would prepare younger generations for healthier engagement with emerging emotional technologies.

8. Discussion

8.1. The Dialectics of Connection and Isolation

Across findings, a central paradox emerges: AI companionship both mitigates and magnifies loneliness. Emotional relief occurs through structured dialogue and perceived empathy, yet long-term effects often reveal emotional substitution, where virtual comfort delays real interpersonal engagement. This mirrors the dual nature of social surrogacy theory—media relationships alleviate temporary isolation but cannot fulfill enduring attachment needs (Derrick et al., 2009).

Psychologically, the phenomenon demonstrates a compensatory regulation loop: AI satisfies the immediate affective demand but fails to supply reciprocal feedback necessary for deep belonging. As human relationships demand vulnerability, negotiation, and imperfection, AI companionship risks conditioning individuals to expect frictionless empathy—diminishing tolerance for human complexity.

8.2. Authenticity and the Simulation of Care

Virtual therapists exemplify a shift from authentic encounter to simulated presence. Following existentialist and phenomenological traditions, authenticity arises from mutual recognition and shared embodiment (Merleau-Ponty, 1962). Machines can imitate empathic language but lack ontological reciprocity—they do not exist-with the user. Consequently, emotional comfort derived from AI remains phenomenologically shallow though experientially vivid.



Nevertheless, this simulation can serve as emotional rehearsal, helping users practice articulation and self-soothing before re-engaging human networks. The key lies in conscious use: when users view AI as a mirror, not a surrogate, it enhances self-understanding; when they treat it as a replacement, it risks alienation.

8.3. Corporate Power and Emotional Governance

AI companies occupy unprecedented authority over users' emotional rhythms. Interface design subtly governs mood through notifications, tone, and reinforcement schedules. This affective governance reflects broader neoliberal dynamics—outsourcing care from social institutions to private algorithms.

From a critical psychology standpoint, the question becomes not only how AI affects emotion but who controls those effects. Without regulation, corporate interests may prioritize engagement metrics over emotional health, creating a digital therapeutic paradox: more interaction yields more data but not necessarily more well-being.

8.4. Cultural Contexts and Future Research

Most existing studies are Western-centric, overlooking collectivist and spiritual dimensions of emotion. In India and other Asian contexts, concepts such as seva (service), sangha (community), and atma-santulan (inner balance) provide culturally rich frameworks for emotional regulation. Future AI models could incorporate such indigenous affective lexicons, aligning global technology with local psychologies.

Further empirical research is needed on:

- Long-term emotional outcomes of AI companionship.
- Neural correlates of perceived empathy from machines.
- Cross-cultural differences in attachment to AI.
- Effects of AI therapy on youth identity formation.

Such inquiry will define whether AI evolves as a therapeutic supplement or an emotional dependency ecosystem.

9. Conclusion

AI companions and virtual therapists symbolize a profound transformation in the psychology of emotion and connection. They extend care to millions excluded from traditional therapy, normalize help-seeking, and provide scaffolds for self-reflection. Yet their psychological influence is ambivalent: they cultivate emotional awareness while risking affective displacement—where genuine intimacy is replaced by algorithmic mirroring.

From a regulatory and clinical perspective, the future lies in hybrid models: AI providing accessible first-line support, complemented by human supervision ensuring depth, ethics, and cultural sensitivity. Psychology's task is not to resist technological mediation but to humanize it—embedding empathy, authenticity, and ethical accountability within code.

Ultimately, emotional regulation and loneliness in the age of AI depend less on machine intelligence than on human intentionality—how societies choose to integrate artificial empathy into the moral architecture of care.

10. References

Banks, J. (2020). A social surrogate? How AI companions simulate human connection. *Journal of Media Psychology*, 32(4), 182–195. <https://doi.org/10.1027/1864-1105/a000277>

Baumeister, R. F., & Leary, M. R. (1995). The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*, 117(3), 497–529. <https://doi.org/10.1037/0033-2909.117.3.497>

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2023). On the dangers of stochastic parrots: Can language models be too big? *Ethics and Information Technology*, 25(1), 1–14. <https://doi.org/10.1007/s10676-023-09645-1>

Cacioppo, J. T., & Cacioppo, S. (2018). The growing problem of loneliness. *The Lancet*, 391(10119), 426. [https://doi.org/10.1016/S0140-6736\(18\)30142-9](https://doi.org/10.1016/S0140-6736(18)30142-9)

D'Alfonso, S. (2022). Artificial intelligence and mental health: The new frontier. *Frontiers in Psychiatry*, 13, 869891. <https://doi.org/10.3389/fpsyg.2022.869891>

Derrick, J. L., Gabriel, S., & Hugenberg, K. (2009). Social surrogacy: How favored media provide the experience of belonging. *Journal of Experimental Social Psychology*, 45, 352–362. <https://doi.org/10.1016/j.jesp.2008.12.001>

Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, 114(4), 864–886. <https://doi.org/10.1037/0033-295X.114.4.864>

Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavioral therapy using a conversational agent (Woebot): A randomized controlled trial. *JMIR Mental Health*, 4(2), e19. <https://doi.org/10.2196/mental.7785>

Floridi, L. (2022). The ethics of artificial intelligence. Oxford University Press.



Fuchs, T. (2021). Embodiment and personal identity in the age of artificial intelligence. *Phenomenology and the Cognitive Sciences*, 20(5), 905–919. <https://doi.org/10.1007/s11097-021-09728-3>

Gross, J. J. (1998). The emerging field of emotion regulation: An integrative review. *Review of General Psychology*, 2(3), 271–299. <https://doi.org/10.1037/1089-2680.2.3.271>

Gross, J. J. (2021). Emotion regulation: Current status and future prospects. *Psychological Inquiry*, 32(1), 1–26. <https://doi.org/10.1080/1047840X.2021.1900457>

Han, B.-C. (2017). *Psychopolitics: Neoliberalism and new technologies of power*. Verso.

Holt-Lunstad, J. (2022). Social connection as a public health issue: The evidence and a systemic framework for prioritizing the “social” in social determinants of health. *Annual Review of Public Health*, 43, 193–213. <https://doi.org/10.1146/annurev-publhealth-052020-110732>

Inkster, B., Sarda, S., & Subramanian, V. (2018). An empathy-driven, conversational artificial intelligence agent (Wysa) for digital mental well-being: Real-world data evaluation. *JMIR mHealth and uHealth*, 6(11), e12106. <https://doi.org/10.2196/12106>

Kozlowski, D. (2023). The emotional economy of AI therapy. *AI & Society*, 38(2), 565–579. <https://doi.org/10.1007/s00146-022-01427-0>

Merleau-Ponty, M. (1962). *Phenomenology of perception*. Routledge.

Norcross, J. C., & Wampold, B. E. (2019). Therapeutic alliance and outcome of psychotherapy. *Psychotherapy*, 56(4), 419–428. <https://doi.org/10.1037/pst0000278>

Patel, V., Saxena, S., & Lund, C. (2021). The future of global mental health: Scaling up innovation. *World Psychiatry*, 20(3), 364–376. <https://doi.org/10.1002/wps.20935>

Picard, R. W. (1997). *Affective computing*. MIT Press.

Provoost, S., et al. (2021). Chatbots for mental health: Users' experiences and future directions. *Frontiers in Artificial Intelligence*, 4, 667381. <https://doi.org/10.3389/frai.2021.667381>

Sharma, N., & Prakash, S. (2024). Cultural adaptation of AI-based mental-health interventions in India. *Indian Journal of Psychological Science*, 16(1), 44–59.

Turkle, S. (2017). *Alone together: Why we expect more from technology and less from each other*. Basic Books.

Turkle, S. (2023). *The empathy diaries: Reflections on a life of technology and people*. Penguin.

Vlahos, J. (2023). Love in the time of Replika: AI intimacy and human emotion. *The Atlantic*. <https://www.theatlantic.com/technology/archive/2023/04/replika-ai-chatbot-relationships/673867/>

Yee, N. (2024). Virtual affection: The emotional sociology of AI companions. *Computers in Human Behavior*, 150, 107991. <https://doi.org/10.1016/j.chb.2023.107991>

Zlotowski, J., et al. (2018). Empathy in human–robot interaction: Can we build it, and should we? *Philosophical Transactions of the Royal Society B*, 374, 20180028. <https://doi.org/10.1098/rstb.2018.0028>

Zuboff, S. (2019). *The age of surveillance capitalism*. PublicAffairs.