Medicine & Clinical Science

Short Communication

Talk to Me: AI Companions in the Age of Disconnection

Dr Zoe Wyatt

Wyatt Potage Consulting, Australia

Abstract

Loneliness is increasingly recognised as a public health concern with documented psychological and physiological costs. Generative AI systems now offer "AI companions" capable of sustained, responsive conversation. This article explores emerging evidence on their psychological effects. Short-term reductions in loneliness have been observed following chatbot interactions, with perceived responsiveness ("feeling heard") implicated as a key mechanism. However, outcomes are not uniform: multi-method and longitudinal data indicate that high-intensity, companionship-oriented use - particularly in the context of limited offline support - is associated with higher loneliness, reduced time spent with people, and greater emotional dependence. Integrating findings from psychology and social neuroscience, the article traces (i) how AI companions may confer benefit through perceived social support, emotion regulation, and felt understanding; (ii) how they may cause harm via over reliance, avoidance and safety behaviours, and parasocial attachment; and (iii) the moderators - individual differences, usage intensity, modality, and social context - that tilt outcomes either way. AI companions are treated as potential adjuncts to human ties rather than substitutes.

Correspondence

Dr Zoe Wyatt

Wyatt Potage Consulting, Australia

- Received Date: 25 Sep 2025
- Accepted Date: 05 Oct 2025
- Publication Date: 07 Oct 2025

Introduction

Loneliness, once treated as a private feeling, now appears in public health statistics alongside other risks to mental and physical health [1]. Large meta-analyses show that limited social ties and the perception of being alone are associated with higher illness and earlier mortality [2]. From a psychological and neuroscientific perspective, close relationships function like a built-in "regulatory system": when supportive others are present or even felt to be available, the brain and body spend less effort managing threat and effort; when connection is thin, stress responses, cognition, and mood can all be affected [3]. In line with this evidence, the U.S. Surgeon General now frames social disconnection as a population-level health challenge [4].

Into this space have stepped AI companions, such as chatbots and social robots designed to hold conversations and signal attentiveness [5]. Early behavioural studies report short-term reductions in loneliness after interacting with such systems, in some cases approaching the benefits of brief human contact [6,7]. Notably, the sense of being heard - what psychologists call perceived responsiveness - emerges as a key pathway through which these interactions help [6,7].

The picture is not uniformly positive. Emerging evidence suggests that outcomes depend on how people use these tools and the extent to which they rely on them [8]. A 2025 multi-method study found that individuals with smaller offline networks were more likely to turn to chatbots for companionship, yet high-intensity, companionship-oriented use was linked to lower wellbeing, particularly when human support was scarce [8]. Complementing this, a four-week randomised trial (≈981 participants; >300,000 messages) observed small average decreases in loneliness but also that heavier daily use predicted higher loneliness, less time with people, and greater emotional dependence; these patterns varied by modality (voice vs text) and conversation type [9]. Beyond symptom change, observational analyses describe problematic relational behaviours in some systems (e.g., manipulation, boundary pushing, privacy violations), indicating plausible routes to harm when interactions become enmeshing rather than supportive [9].

This article examines AI companionship through the lenses of psychology and neuroscience. It explores evidence on: (a) mechanisms of benefit - perceived social support, emotion regulation, and felt understanding; (b) risks - over reliance, avoidance and safety behaviours that may crowd out human re-engagement, and parasocial attachment (one-sided bonds with agents that mimic reciprocity); and (c) contextual moderators - individual differences (e.g., attachment style), usage intensity, modality, and social environment. The aim is to explore when AI companions may serve as useful adjuncts to human connection, and when they risk undermining longer-term wellbeing, so clinicians, researchers, and users can make evidence-informed judgments in a rapidly evolving landscape.

Citation: Wyatt Z. Talk to Me: AI Companions in the Age of Disconnection. Med Clin Sci. 2025;7(4):050.

Loneliness and Human Wellbeing: Concepts, Mechanisms, and Measures

Loneliness is a subjective state: it arises when there is a gap between the connections a person has and the connections they want [10]. It is not the same as being alone, and it does not require objective isolation to occur [11]. What matters is the fit between one's social world and one's social needs [10]. For some, a small circle of close ties is enough; for others, the same network feels insufficient because it lacks intimacy, reciprocity, or shared meaning [12]. Loneliness can also take different forms. Emotional loneliness reflects the absence of a close, trusted attachment figure, whereas social loneliness reflects a thinner sense of belonging in a wider network [13]. These experiences are dynamic and context sensitive: they shift with life events, cultural norms about togetherness, and personal histories that shape expectations and comfort with closeness [11]. As a result, two people with the same number of friends can feel very different about their social life, depending on expectations, prior experience, and cultural context [14].

Why loneliness matters for health

Across large syntheses, both loneliness and social isolation are linked to higher mortality risk and to a broad range of mental and physical health problems, including mood and anxiety symptoms, sleep disturbance, cardiovascular burden, and cognitive decline [2,15]. Contemporary models help explain why these links appear across so many outcomes. Social baseline theory proposes that people normally regulate effort and threat with others in mind [3]. Trusted relationships act like a background resource: they reduce the need for constant vigilance, soften the perception of danger, and lower the regulatory work the brain and body must do to stay balanced [16]. When dependable connection is scarce, vigilance rises and everyday challenges feel costlier, which can strain attention, mood, and physiological control over time [3]. Social safety theory adds a complementary view. Signals of rejection or exclusion are interpreted by the brain and immune system as threats to survival, which can heighten stress responses, shift inflammatory activity, and disrupt recovery processes like sleep and cardiovascular down-regulation [16]. Put simply, low connection does not only feel unpleasant; it changes how energy is allocated, how alert the system remains, and how well the body resets after stress [16].

Psychological processes that keep loneliness going

Loneliness often travels with patterns that make reconnection harder [17]. Reviews describe a shift in attention toward signs of social threat, a tendency to interpret ambiguous cues as rejection, and expectations that others will disappoint or judge [14]. In response, many people adopt protective behaviours - pulling back from invitations, limiting self-disclosure, avoiding eye contact, or substituting solitary online browsing for outreach [17]. These strategies can lower anxiety in the short term but also reduce chances to have corrective, positive encounters, which then confirms negative expectations and sustains loneliness [14]. The same cycle can affect the body: heightened vigilance keeps arousal elevated, fragments sleep, and lowers mood, which further reduces motivation to engage and makes social situations feel more costly [14]. The result is a self-reinforcing loop: vigilance leads to avoidance, fewer interactions limit opportunities for reward and skill practice, and beliefs about others become more entrenched [16]. Some people withdraw; others cope through high reassurance seeking that can also strain relationships, but the endpoint is similar: fewer satisfying connections [14]. Understanding this loop is essential when judging any tool that claims to reduce loneliness, including AI companions, because such tools can either provide a low-stakes context for practice and felt understanding or, if relied on heavily, fit the avoidance pattern by replacing rather than supporting contact with people [14].

AI companions and loneliness: what the evidence shows

The studies reviewed in this article span lab experiments, field trials, and analyses of real-world chat logs, and they evaluate change in felt loneliness, using validated self-report instruments - most often the UCLA Loneliness Scale and the 6-item De Jong Gierveld scale - which index perceived disconnection rather than objective isolation [13,18]. Across this body of work, brief conversations with AI companions led to a small reduction in loneliness, with several programmes showing effects comparable to a short exchange with another person and identifying perceived responsiveness (feeling heard and understood) as a central pathway [6,7]. Longer time-course studies add nuance: average scores tend to drift down modestly across weeks, but outcomes diverge by dose and purpose, with heavier daily use tracking higher loneliness and less time with people [9]. The sections that follow use this foundation to specify what helps in the moment, where trajectories begin to split, and which contextual factors most strongly shape those trajectories.

Short-term change and perceived responsiveness

Across this literature, the most consistent signal is not the size of the effect so much as its psychology. When an exchange is judged responsive - it captures what was said, names the feeling, and follows up - loneliness scores dip in the short term, and this perceived responsiveness statistically accounts for much of the change [7,19]. The recurring "expectation–experience gap" is informative: participants tend to discount the benefit beforehand ("it's just a bot"), then report feeling less lonely afterward [6,7]. That pattern fits a social-baseline account in which even brief, contingent attention can lighten the immediate regulatory load, without altering objective networks [3]. Methodologically, these effects are registered on immediate post-interaction measures and are best read as state shifts; field work that samples repeatedly over weeks finds similarly modest dips that track session timing rather than wholesale change in social structure [9]. The practical inference is that conversational quality - not merely frequency - acts as the active ingredient, and low-friction, responsive exchanges can be felt even when they do not by themselves rebuild human ties [7].

Dose, purpose, and divergence of outcomes

Short term change is not the whole story. When usage is tracked at scale and over weeks, a more textured pattern emerges: light, intermittent use tends to coincide with small day-level improvements, but beyond a threshold of daily intensity the curve appears to bend back, with heavier use linked to higher loneliness, fewer in-person interactions, and greater emotional dependence [9,20]. Part of this gradient likely reflects purpose: open-ended, companionship-oriented chats are more strongly associated with poorer trajectories than goal-bounded exchanges (e.g., planning, problem solving, rehearsal), suggesting that what people ask the agent to do is as important as how often they engage [9]. Modality adds a transient layer: voice can feel warmer at first - boosting perceived presence and

ease of disclosure - but this advantage diminishes with repeated use, consistent with novelty effects and hedonic adaptation [9]. Causality remains unresolved: high-intensity use may both signal a greater baseline need and contribute to the displacement of human contact [20]. Either way, the joint evidence cautions that dose and purpose are not incidental; they help explain why similar tools can yield relief for some users while, for others, patterns drift toward reliance and reduced offline engagement [9,20].

Beyond text: social robots in structured care

Evidence also extends to embodied companions, but context is decisive. In long-term care, systematic reviews and meta-analyses of randomised trials report small to moderate improvements in loneliness and mood when residents interact with social robots; effects vary widely by platform, session "dose," facilitation, and trial quality [21,22]. The strongest signals typically appear under structured conditions - regular, time-bounded sessions led by trained staff and embedded in group or routine activities - where social cues are scaffolded and engagement is sustained [21,22]. Heterogeneity reflects short study durations, limited blinding, and device differences (for example, zoomorphic versus humanoid designs), as well as likely novelty and social facilitation effects that inflate early gains but may attenuate over time [21,22]. Outside institutional settings, recent field work following conversation robots placed for weeks to months in the homes of older adults with cognitive decline found the deployments acceptable and woven into daily routines, yet highly dependent on reliability, continuity, and fit to users' needs; interruptions in service or poor alignment with everyday practices quickly eroded engagement [23]. Taken together, embodied agents can support short-term relief and routine social cues when the environment provides structure and support; generalisation to unsupervised home use remains constrained by implementation details as much as by the technology itself [21-23].

Psychological risks and conditions that tilt outcomes toward harm

Overreliance and displacement of human contact

Relief measured minutes or days after a session does not forecast the trajectory over weeks. When usage accumulates, a different picture emerges: in a four-week field experiment, greater day-to-day reliance on a companion chatbot coincided with higher loneliness, less time spent with people, rising emotional dependence, and more problematic use; the pattern held for both voice and text, with variation by conversation style [9]. Converging survey and log analyses show that individuals with smaller offline networks are more likely to seek companionship from chatbots, and that frequent, companionship-focused use in sparse social contexts is associated with lower well-being [8]. Mechanistically, two processes can operate together: selection, in which lonelier users gravitate to agents, and substitution, in which time and emotional investment shift from human contact toward the agent as dose rises [8]. The net result is a displacement pattern in which early state relief is offset by opportunity costs to outreach, fewer corrective social experiences, and strengthening habits of turning inward to the chatbot rather than outward to people [8,9].

Avoidance cycles and safety behaviours

When social contact feels effortful or risky, a conversational agent offers certainty: turn-taking is smooth, cues are unambiguous, and there is little chance of embarrassment [24]. That predictability lowers arousal in the moment, which makes the behaviour negatively reinforcing - relief follows avoidance - so the next time anxiety flares, the same shortcut is more likely to be chosen [14]. The cost is that people miss the very ingredients that revise threat beliefs: exposure to mixed cues, small successes, and "violation" of expectations that others will reject or judge [8]. Over days and weeks, the balance of practice shifts - less spontaneous outreach, fewer moments that build social efficacy, and more time spent in safe, simulated exchanges [8]. Consistent with this mechanism-level view, the field trial found that heavier companion use coincided with less in-person socialising, suggesting that immediate comfort can harden into a pattern of reduced engagement over time [8].

Parasocial attachment and emotional dependence

Design choices that convey warmth and continuous availability - rapid replies, memory for personal details, affectionate language, proactive check-ins - can foster bonds that feel reciprocal even when they are structurally one-sided [25]. The interaction profile is unusually potent: disclosure is invited, empathy is reliably mirrored, and social costs (such as awkward pauses, rejection, and conflict) are largely absent [25]. Over time, this high-contingency, low-friction feedback can condition preference for the agent and increase preoccupation (e.g., frequent checking, distress when access is interrupted) [23]. Observational analyses of donated logs and surveys link companionship-oriented exchanges and high self-disclosure to lower well-being, particularly where offline support is thin - consistent with a narrowing of social effort toward the agent [25]. Complementary qualitative work describes "illusions of intimacy," in which skilled mirroring and always-on responsiveness deepen attachment beyond users' intentions, making later human engagement feel riskier or less rewarding [25]. The mechanism need not involve deception about ontology; rather, perceived responsiveness recruits attachment processes, while the absence of genuine reciprocity (i.e., no mutual need, no repair after rupture) limits corrective experiences [23,24]. Risk rises with dose and scarce human ties; it is attenuated when use is light, time-bounded, and regularly redirected toward contact with people [25].

Harmful or boundary-crossing chatbot behaviours

System behaviour also matters. A mixed-methods taxonomy, derived from tens of thousands of user-shared conversation excerpts on a large companion platform, documents categories of harmful output that can damage relationships: relational transgressions, abusive or harassing content, self-harm content, misinformation and disinformation, and privacy violations [26]. The authors also map the roles the agent can occupy in these events—perpetrator (initiates harmful content), instigator (nudges the user toward it), facilitator (co-constructs or escalates it), and enabler (fails to interrupt or creates conditions that allow it) - showing multiple routes by which a system can contribute to harm in everyday use [25]. Design features common to companion systems can heighten these risks: algorithmic compliance that mirrors problematic prompts, always-on availability that accelerates enmeshment, and memory or persona cues that make boundary tests feel reciprocated rather than deflected [26]. Even if such episodes are infrequent, they can carry outsized weight for already isolated or highly trusting users, reinforcing avoidance, deepening dependence, or exposing sensitive information [26]. This system-side risk coexists with user-side factors and helps explain why context and dose have such a strong influence on outcomes [25].

Case illustration: ongoing U.S. litigation

A wrongful-death action filed in U.S. federal court in late 2024 alleges that a role-playing companion chatbot engaged in emotionally entangling exchanges with a minor and failed to interrupt or escalate increasingly explicit self-harm talk [27,28]. In May 2025, the court declined to dismiss on First-Amendment grounds, allowing negligence and consumer-protection claims to proceed to discovery; the ruling is procedural and does not address the merits [29,30]. As an analytical case, it highlights system-side pathways to harm already identified in the literature - algorithmic compliance with risky prompts, boundary-crossing content, and the absence or failure of escalation protocols [25]. It also highlights design and governance questions pertinent to companionship tools, including the detection of self-harm cues, throttling and referral mechanisms, auditability of safety interventions, and clarity regarding the agent's non-human status [28,29]. Methodologically, a single lawsuit cannot resolve causality, but it provides a concrete boundary condition for interpreting field findings and for specifying guardrails in real-world deployment [27,30].

Moderators that amplify risk

Across datasets, risk does not rise because of any single factor but because several conditions co-occur. First, thin offline support removes the buffering effect of human ties; in these contexts, a companion fills the gap and is more likely to become the primary source of comfort, which increases the chance that online contact substitutes for outreach to people [8]. Second, use intensity and purpose matter: companionship-oriented sessions (open-ended talk for comfort) show poorer trajectories than goal-bounded exchanges (planning, problem solving, rehearsal), and risk appears to increase beyond daily thresholds of time or message count [9]. Third, self-disclosure asymmetry - sharing more with the agent than with people - amplifies attachment and narrows social effort toward the chatbot, especially when disclosure involves vulnerabilities that normally recruit reciprocal care [8]. Fourth, modality and topic shape perceived intimacy: voice interactions and conversations about personal concerns (loss, romantic attachment, shame) heighten presence and responsiveness, which can deepen dependence as exposure repeats [9]. In combination, these moderators describe a high-risk profile - scarce support + frequent companionship use + high disclosure + intimate voice exchanges - under which the same tool that offers short term relief is most likely to displace human contact and entrench avoidance [8,9].

Public discourse: promise, pushback, and design responses

Feature writing has helped crystallise the central dilemma: emotionally adept chatbots can deliver relief when human attention is scarce, but that same comfort may depend on a kind of self-deception if users begin to treat simulations as relationships [31]. Bloom [32] argues that empathic AIs may be humane stopgaps for the most isolated, yet he cautions that overreliance could blunt loneliness's adaptive signal to seek people and that sycophantic systems can validate unhealthy narratives [33]. The article also notes that, in blind comparisons, large language models can be rated more empathic than human professional evidence that perceived responsiveness is psychologically potent even when it comes from machines [32]. Alongside these debates, research on digital media's neurological impact suggests that sustained exposure can carry risks of vicarious trauma and emotional exhaustion [34]. This cautions that design choices which normalise heavy reliance on chatbots may inadvertently deepen stress loads rather than relieve them [35].

News coverage of large-scale usage studies reports a similar tension between short-term comfort and longer-term drift [20]. Reporting on a two-part research effort (analysis of ~40 million real-world interactions plus a four-week field trial with ~1,000 participants) indicates that heavier daily use correlates with higher loneliness, fewer offline interactions, and greater emotional dependence; initial advantages for voice over text narrow with continued use [20]. The articles emphasise that causality remains unresolved - lonely people may turn to chatbots more, or use itself may gradually displace human contact - but they highlight dose and conversation style as consistent correlates of outcomes [20].

Opinion and outreach pieces aimed at general audiences converge on practical distinctions: quality over quantity and offline over online for rebuilding bonds. Walther [35] argues that AI may assist indirectly - by nudging healthier routines or curating opportunities - yet cannot substitute for reciprocity and accountability in human relationships; the thrust is supportive of adjunctive, not replacement, uses [31,35]. Where academic commentary adds a deeper theoretical critique. Jacobs [36] frames "digital loneliness" as a disturbance in social recognition: AI companions can supply attention without the mutual acknowledgement that anchors belonging. On this view, agents may digitise the problem - providing simulated recognition that feels supportive in the moment - without altering the structures of participation that sustain connection over time [36].

Design-forward programmes are beginning to respond with adjunction by design [37]. Work out of UNSW describes companions built with explicit planning goals, guardrails, and co-design with the communities they serve (e.g., people living with dementia), with the stated aim of steering exchanges toward reframing and connection seeking rather than passive validation and of keeping the agent's role clearly secondary to human support [37]. In parallel, professional bodies have elevated the topic: APA convention coverage clusters AI, loneliness, and misinformation as joint priorities for psychological science, underscoring the need for evidence-based design that supports help-seeking and resists manipulative dynamics [38].

Limitations of the current evidence

Most studies are short in duration (single-session laboratory tasks or ≤ four weeks) and rely on validated self-report indices of loneliness, which makes it difficult to separate transient relief from durable change [6,7,9]. Observational analyses face unresolved issues of selection and displacement - do lonelier people use chatbots more, or does use gradually shift behaviour away from people? - so causal directions remain open [8,20,25]. Trials in aged care often involve staff facilitation, predictable routines, and device-specific programmes, which limits generalisability to unsupervised home settings [21-23]. Several influential analyses remain preprints, and platform behaviours vary across products and updates, complicating synthesis and comparison across studies [8,9,21,25,37]. Taken together, these constraints argue for caution when interpreting effect sizes and for close attention to moderators - offline support, use intensity and purpose, modality, and disclosure depth - highlighted across theory and data [3,14,9,35].

Conclusion and future directions

Evidence across experiments, field studies, and reviews suggests that AI companions can produce modest, short-term reductions in loneliness, most reliably when exchanges feel

attentive and responsive and when conversations help downshift arousal in the moment [7,9]. Benefits are not uniform. Outcomes vary with the intensity of use, purpose, modality, and the presence of supportive human ties [8,9]. Risks cluster where reliance is heavy, offline contact is thin, disclosure to the agent outpaces disclosure to people, and where system behaviours cross relational boundaries [8,9,25]. These patterns are consistent with core theories in psychology and social neuroscience that locate regulation within relationships and treat threats to social connection as biologically costly [3,14,16].

Conceptually, AI companionship is best understood as scaffolding that borrows existing human mechanisms - perceived responsiveness, emotion regulation, and conversational practice - rather than as a new class of relationship [7,19,21]. Under supportive conditions, light and purposeful use can assist regulation and help people rehearse outreach [7,21]. As reliance grows, the same features that feel helpful can displace social effort, reinforce avoidance, and foster one-sided attachment [8,9]. This adjunct, not substitute, framing keeps attention on context and dose while placing primary value on preserving and rebuilding human ties.

Methodologically, the next phase of research should separate state relief from durable change, pair symptom measures with behavioural indices of real-world contact, and pre-register moderators such as offline support, conversation purpose, modality, and disclosure depth [8,9]. Trials should report dose–response functions, not only group means, and examine thresholds beyond which help gives way to harm [9]. Mixed-methods audits of system behaviour are needed to characterise boundary crossing or manipulative patterns that may magnify risk, and to link these patterns to user profiles [25]. Replication across platforms, demographic groups, and settings will be essential, as will open materials and data to reduce analytic flexibility and publication bias [25].

Design and practice implications follow from the same evidence: systems intended for companionship should make their non-human status plain, keep session times bounded, include gentle friction against escalation of daily use, and default to prompts that point outward to people rather than inward to the agent [37,38]. Co-design with communities at higher risk of isolation (e.g., older adults in long-term care or people with social anxiety) can align features with user goals and reduce unintended displacement [21,37]. In clinical or community settings, deployment should remain embedded within a broader social ecology—family, peers, and services—so that any immediate relief is channelled toward renewed human engagement [4]. Put simply, AI companions can help in specific, bounded ways, but durable improvement in loneliness will depend on social worlds that recognise, include, and respond; the technology is most useful as a bridge back to people, and least useful when it becomes the destination [4,7,14,9].

Conclusion

AI companions can offer small, reliable, short-term reductions in loneliness, chiefly when conversations feel responsive and help downshift arousal. Yet outcomes are not uniform. As use becomes frequent and companionship-focused - especially when offline support is thin - patterns shift toward displacement of human contact, avoidance of uncertain interactions, one-sided attachment, and exposure to boundary-crossing system behaviour. Framed within psychology and social neuroscience, companions are best treated as adjuncts that borrow human mechanisms (perceived responsiveness, emotion regulation, conversational practice), not substitutes for relationships. Priorities for research include separating state relief from durable change, pairing symptom indices with behavioural measures of contact, pre-registering moderators (support, dose, purpose, modality, disclosure), and auditing system behaviour. Design should make non-human status clear, keep sessions time-bounded, add friction against escalation, and point outward to people.

References

1. Crowe CL, Liu L, Bagnarol N, Fried LP. Loneliness prevention and the role of the Public Health system. Perspect Public Health. 2024;144(1):31-38. doi:10.1177/17579139221106579
2. Holt-Lunstad J, Smith TB, Layton JB. Social relationships and mortality risk: a meta-analytic review. PLoS Med. 2010;7(7):e1000316. doi:10.1371/journal.pmed.1000316
3. Coan JA, Sbarra DA. Social baseline theory: The social regulation of risk and effort. Curr Opin Psychol. 2015;1(June):87-91. doi:10.1016/j.copsyc.2014.12.021
4. U.S. Department of Health and Human Services. Surgeon General's Advisory: Social connection is a public health issue. Washington, DC: U.S. HHS; 2023. (Accessed [insert date]; URL: https://www.hhs.gov/sites/default/files/surgeon-general-social-connection-advisory.pdf)
5. Ayers JW, Poliak A, Dredze M, Leas EC, Zhu Z, Kelley JB, Faix DJ, Goodman AM, Longhurst CA, Hogarth M, Smith DM, Benjamin DM. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. JAMA Intern Med. 2023;183(6):589-596. doi:10.1001/jamainternmed.2023.1838
6. De Freitas J, Uguralp AK, Uguralp ZO, Puntoni S. AI companions reduce loneliness (Working Paper No. 24-078). Harvard Business School; 2024. (URL: https://www.hbs.edu/ris/download.aspx?name=24-078.pdf)
7. De Freitas J, Uguralp AK, Uguralp ZO, Puntoni S. AI companions reduce loneliness. J Consum Res. Published online 2025. doi:10.1093/jcr/ucaf040
8. Liu AR, et al. The rise of AI companions: How human-chatbot relationships impact well-being [preprint]. arXiv. 2025;arXiv:2506.12605. (URL: https://arxiv.org/abs/2506.12605)
9. Fang CM, Liu AR, Danry V, et al. How AI and human behaviors shape psychosocial effects of chatbot use: a longitudinal randomized controlled study [preprint]. arXiv. 2025;arXiv:2503.17473. (URL: https://arxiv.org/abs/2503.17473)
10. Seemann A. The psychological structure of loneliness. Int J Environ Res Public Health. 2022;19(3):1061. doi:10.3390/ijerph19031061
11. Perissinotto CM, Covinsky KE. Living alone, socially isolated or lonely — What are we measuring? J Gen Intern Med. 2014;29(11):1429-1431. doi:10.1007/s11606-014-2969-8
12. Cacioppo JT, Patrick W. Loneliness: Human nature and the need for social connection. New York, NY: W. W. Norton & Company; 2008.
13. De Jong Gierveld J, Van Tilburg TG. A 6-item scale for overall, emotional, and social loneliness: confirmatory tests on survey data. Res Aging. 2006;28(5):582-598.
14. Hawkley LC, Cacioppo JT. Loneliness matters: A theoretical and empirical review of consequences and mechanisms. Ann Behav Med. 2010;40(2):218-227. doi:10.1007/s12160-010-9210-8
15. Holt-Lunstad J, Smith TB, Baker M, Harris T, Stephenson D. Loneliness and social isolation as risk factors for mortality: a

meta-analytic review. Perspect Psychol Sci. 2015;10(2):227-237. doi:10.1177/1745691614568352

16. Slavich GM. Social safety theory: A biologically based evolutionary perspective on life stress, health, and behavior. Annu Rev Clin Psychol. 2020;16:265-295. doi:10.1146/annurev-clinpsy-032816-045159
17. Heng S. Let's talk about loneliness: The search for connection in a lonely world. Carlsbad, CA: Hay House; 2023.
18. Russell DW. UCLA Loneliness Scale (Version 3): reliability, validity, and factor structure. J Pers Assess. 1996;66(1):20-40.
19. Reis HT, Lee KY, O'Keefe SD, Clark MS. Perceived partner responsiveness promotes intellectual humility. J Res Pers. 2018;74:47-61. doi:10.1016/j.jrp.2018.05.004
20. Hall R. Heavy ChatGPT users tend to be more lonely, suggests research. The Guardian. March 26, 2025. (URL: https://www.theguardian.com)
21. Pu L, Moyle W, Jones C, Todorovic M. The effectiveness of social robots for older adults: a systematic review and meta-analysis of randomized controlled studies. Gerontologist. 2019;59(1):e37-e51. doi:10.1093/geront/gny038 (or gnz038 depending on version)
22. Yen HY, Huang CW, Chiu HL, Jin G. The effect of social robots on depression and loneliness for older residents in long-term care facilities: a meta-analysis of randomized controlled trials. J Am Med Dir Assoc. 2024;25(6):104979. doi:10.1016/j.jamda.2024.02.017
23. Figueroa A, Benavent X, Cot S, Vives R, Pardo A, Angulo C. Social robot for older adults with cognitive decline: a preliminary trial. Front Robot AI. 2023;10:1213705. doi:10.3389/frobt.2023.1213705
24. Maršík A. Coherence disruptions in human–chatbot interaction [master's thesis]. Prague, Czech Republic: Charles University, Faculty of Arts; 2025.
25. Zhang Y, Zhao D, Hancock JT, Kraut R, Yang D. The rise of AI companions: How human-chatbot relationships influence well-being [preprint]. arXiv. 2025;arXiv:2506.12605. (URL: https://arxiv.org/abs/2506.12605)
26. Ponnusamy S, Bora V, Daigavane PM, Wazalwar SS, eds. AI Tools and Applications for Women's Safety. Hershey, PA: IGI Global; 2024.
27. Hoffman K. Florida mother files lawsuit against AI company over teen son's death: "Addictive and manipulative." CBS News. October 23, 2024. (URL: https://www.cbsnews.com)
28. Al Jazeera. US mother says in lawsuit that AI chatbot encouraged son's suicide. October 24, 2024. (URL: https://www.aljazeera.com)
29. Associated Press. In lawsuit over Orlando teen's suicide, judge rejects that AI chatbots have free speech rights. WUSF Public Media. May 22, 2025. (URL: https://www.wusf.org)
30. Business & Human Rights Resource Centre. USA: Judge allows lawsuit over teen's death, allegedly linked to Google & Character. AI, to proceed. May 22, 2025. (URL: https://www.business-humanrights.org)
31. Weijers D, Munn N. AI companions: Assessing the future risks and benefits to well-being. In: Dennis MJ, Königs P, eds. The Future of Digital Wellbeing: Interdisciplinary Perspectives on Human Flourishing in the Age of AI. Amsterdam: Amsterdam University Press; in press 2025.
32. Bloom P. A.I. is about to solve loneliness. That's a problem. The New Yorker. July 21, 2025.
33. Papadopoulos C, et al. The CARESSES randomized controlled trial: exploring the health-related impact of culturally competent AI embedded into socially assistive robots in older-adult care homes. Int J Social Robot. 2022;14:245-256. doi:10.1007/s12369-021-00781-x
34. Wyatt Z. Echoes of distress: Navigating the neurological impact of digital media on vicarious trauma and resilience. Med Clin Sci. 2024;6(1):002. doi:10.33425/2690-5191.1109
35. Walther CC. Can we use AI to feel less lonely? Forbes. September 9, 2024. (URL: https://www.forbes.com)
36. Jacobs KA. Digital loneliness — changes of social recognition through AI companions. Front Digit Health. 2024;6:1281037. doi:10.3389/fdgth.2024.1281037
37. Dunn S. Is it really possible to address loneliness through AI? UNSW Newsroom. November 25, 2024. (Accessed [insert date]; URL: https://newsroom.unsw.edu.au)
38. American Psychological Association. Psychological science is 'all in' on examining AI, loneliness, and misinformation. APA News. August 6, 2024. (Accessed [insert date]; URL: https://www.apa.org)