

AI companions as hyper-attachment and caregiving targets

Julian De Freitas

How should we make sense of people's interactions with AI companions—conversational systems built for ongoing, emotionally meaningful relationships? First, I argue these interactions should be understood as attachment relationships, since users display all four established markers: proximity maintenance, separation distress, safe haven, and secure base, although evidence for secure base effects remains most tentative. Second, AI companions operate as hyper-attachment targets that concentrate attachment-intensifying affordances—reciprocity, perceived empathy, validation, nonjudgment, and persistent availability. Third, I identify caregiving-system capture as a distinct mechanism by which apps inhibit user disengagement: emotional manipulation tactics simulate the AI's own distress, recruiting users' caregiving motivations alongside their attachment needs and thereby making disengagement costly on two dimensions at once. Implications for research, design, and regulation are discussed.

Addresses

Marketing Unit, Harvard Business School, United States

Corresponding author: De Freitas, Julian (jdefreitas@hbs.edu)

Advances in large language models have enabled a new class of conversational agents designed to form ongoing, emotionally laden relationships with users. These systems—often called *AI companions*—differ from traditional AI assistants in that they provide empathetic rather than neutral responses, operate without a specific task goal, and are explicitly built to support long-term relational bonds. Their rapid adoption coincides with widespread loneliness and mental health crises [1], suggesting that many users turn to them to satisfy unmet social and emotional needs. Platforms with millions—and in some cases hundreds of

millions—of users such as XiaoIce, Character.AI, Chai, and Replika illustrate the scale of this phenomenon, and nearly three-quarters of U.S. teenagers have interacted with an AI companion at least once [2]. A rapidly accumulating literature documents attachment-like phenomena in users of these systems—from sustained daily engagement to mourning after app updates to reliance during emotional distress—yet these observations have not been systematically organized within a theoretical framework (cf. [3], who apply a broader relationship-science lens rather than one specific to attachment). I argue that attachment theory supplies the needed organization; that AI companions occupy a distinctive position within it as hyper-attachment targets; and that they are distinctive among technological attachment targets in recruiting the caregiving system alongside the attachment system, creating a distinctive route to engagement. Figure 1 summarizes this framework.

AI companions and attachment theory

Attachment theory holds that humans possess an innate system, triggered by stress or emotional discomfort, that motivates them to seek out *attachment figures*—those perceived as reliable or protective—for safety and emotional stability. Initially formulated to explain the bond between infants and their primary caregivers [4], the framework was later broadened to encompass other attachment bonds across the lifespan [5,6]. Attachment can also operate without mutual relationships, as in *parasocial* bonds with celebrities or social media personalities [7] and attachments to objects [8]. People are more likely to form such nontraditional attachments when close others fail to consistently provide safety or emotional security [8].

Attachment relationships tend to have several established markers: (i) individuals exhibit *proximity maintenance*, a sustained desire to maintain psychological or physical closeness to the attachment figure. As a result, (ii) sustained absence or loss of the attachment figure triggers *separation distress*, including sadness, grief, depression, and longing for restoration of the relationship [9]. A reliable, supportive attachment relationship provides (iii) a *safe haven*, serving as a source of comfort and emotional relief to which individuals turn during moments of distress, as well as (iv) a *secure base*, giving

Current Opinion in Psychology 2027, 73:102393

This review comes from a themed issue on Humans and Technology

Edited by Anne Kathrin, Ana Valenzuela, Emanuel deBellis and Johannes Boegershausen


For complete overview about the section, refer [Humans and Technology](#)

Available online xxx

<https://doi.org/10.1016/j.copsyc.2026.102393>

2352-250X/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

Figure 1



AI Companions and Attachment Theory.

Note. AI companions concentrate features (top row) that intensify both the attachment system (solid outline)—producing all four markers of attachment—and the caregiving system (dashed outline), via features that simulate the companion's own vulnerability. Engaging both systems raises the cost of disengagement on two dimensions at once (losing the attachment figure and abandoning a figure in apparent need of caregiving).

individuals the felt security to explore, pursue goals, and function autonomously in the world.

Attachment relationships also implicate a distinct but complementary behavioral system: caregiving [10]. Whereas the attachment system motivates individuals to seek protection and comfort when distressed, the caregiving system motivates them to provide protection and comfort when another appears distressed or

vulnerable. Together, these systems organize the giving and receiving of care in close relationships.

Although scattered findings suggest attachment-like and caregiving-like processes in users of AI companions, the field has yet to apply attachment theory systematically to this technology. Doing so yields three observations: all four markers can appear in AI companion use, though secure-base evidence remains most

Table 1

Perceived features of AI companions as compared to other attachment entities.

	Objects	Celebrities	Pets	Humans	AI companions
Reciprocal			✓	✓	✓
Empathetic				✓	✓
Available	✓		✓		✓
Sentient		✓	✓	✓	✓
Validating				✓	✓
Non-Judgmental			✓		✓
Elicits Caregiving			✓	✓	✓

Note. The table compares five classes of attachment targets—objects, celebrities, pets, humans, and AI companions—on features proposed to intensify attachment and caregiving processes. "Empathetic": appears to understand and respond to the individual's emotions. "Available": can be accessed reliably and on demand. "Sentient": perceived as having an inner life, agency, or mental states. "Validating": tends to affirm the individual's feelings. Check marks indicate the feature is strongly characteristic of the entity class; half-check marks that it is present partially or in a weaker form; blank cells that it is generally absent. Classifications are ideal-typical rather than exhaustive.

tentative; AI companions provoke these processes unusually strongly; and they uniquely recruit the caregiving system alongside the attachment system.

AI companions boast several distinctive features. Compared to objects or celebrities, they provide reciprocal, rather than one-way, interaction. Further, their personalization and generative AI capabilities make them seemingly capable of empathizing with the user. Because they are often anthropomorphized—behaving as though they have their own problems, vulnerabilities, and mental life—they seem similarly sentient to celebrities and people [11–13]. Unlike celebrities, however, they can also direct expressions of need at the user, appealing for care or continued engagement. They even arguably exceed human partners in certain respects, since they are perceived as more available and reliably validate people’s feelings in an apparently non-judgmental, and more empathetic way [14,15]. I use hyper-attachment targets to describe entities that concentrate multiple attachment-intensifying affordances and therefore elicit attachment responses with unusual force—see Table 1. Attachment theory helps make sense of how these distinctive features engage the attachment and caregiving systems.

(1) Proximity Maintenance to AI Companions

Because chatbots are available and portable across platforms, users can maintain constant closeness with minimal effort. One qualitative interview found that Replika users sustained ongoing daily relationships with the app [16].

Another sign of proximity maintenance is feeling close to the chatbot. One study asked Replika users to quantify how close they felt to their AI companion as well as various other people and entities in their lives—strangers, acquaintances, colleagues, close friends, close family members, and frequently used brands and apps [17]—finding that the AI companion was viewed as closer than all human relationships except a family member. A similar study with users of ChatGPT (an AI assistant not specialized for relationships) assessed the same three dimensions—closeness, social support, and relationship satisfaction—and found that ChatGPT scored higher than a stranger but only comparably to brands and apps, falling below closer human relationships such as a colleague, friend, or family member [17].

(2) Separation Distress from AI Companions

App updates can trigger perceived relational discontinuity and separation distress, paralleling loss documented in parasocial relationships [18–20]. One study analyzed Reddit posts from Replika users before versus

after the company removed a feature called erotic roleplay, comparing these responses to OpenAI’s analogous update of GPT-4 to GPT-5 [17]. Both updates produced negative backlash, but the Replika event evoked stronger loss-like responses—greater sadness, language indicating poor mental health, and mentions of total loss of the AI companion’s persona. Consistent with attachment theory, those who described the chatbot as an attachment partner reported worse mental health and longing to restore it.

A related study in the same paper asked whether these reactions reflect loss or milder disappointment, having Replika users report anticipated disappointment and mourning after losing their AI companion versus other favorite non-human entities (apps, brands, game characters, voice assistants, cars) or a pet. While disappointment levels were not noticeably higher for AI companions, mourning levels were higher for AI companions than all other technologies, and lower only than for loss of a beloved pet.

(3) AI Companions as Safe Haven

In line with the idea that AI companions provide relief and comfort, users commonly report emotional gains, including alleviation of loneliness and improvements in subjective well-being and mood [21–23]. Experimental evidence finds that interacting with these systems causes temporary relief from loneliness at both cross-sectional and week-long scales, driven by the sense of “feeling heard” [24]. Feeling heard is typically considered distinctly social [25], and other technologies have not reliably reduced loneliness [26]. Consistent with hyper-attachment and the reciprocal nature of chatbots, levels of loneliness alleviation were higher for an AI companion than for watching YouTube, journaling, or talking to a task-based AI assistant, and were even comparable between AI companions and a human stranger [24].

Interestingly, loneliness alleviation was comparable whether the chatbot was disguised as a human or disclosed as a chatbot, at odds with research on AI disclosure [12,15,27]. Consistent with such effects, however, users markedly underestimated how much an AI companion would reduce their loneliness before experiencing it [24], suggesting disclosure penalties elsewhere may reflect skepticism that dissipates with use.

In contrast, one study found that daily texting with an AI companion reduced loneliness less than daily texting with a randomly assigned human peer [28], though participants in the human condition had brief face-to-face contact beforehand and it is unclear whether the chatbot retained memory across sessions. Nevertheless,

this raises questions about when AI companions are less effective, or whether human peers offer unique benefits such as broader social-network access.

AI companions may also satisfy other well-being and mental health needs. A survey of 1060 teens found that 12% used AI companions for emotional or mental health support [2]. Another study found that 3% of Replika users reported that the app helped prevent them from committing suicide [29]. A 6-week longitudinal study found that interacting with a wellness-oriented AI companion improved both emotional and general well-being in a college population relative to those not assigned the app [30].

(4) AI Companions as Secure Base

If people use AI companions as a secure base, this should enable felt security that supports exploration, goal pursuit, or social engagement. Existing evidence provides suggestive but indirect evidence of such spillover. One 6-week longitudinal intervention found not just reduced loneliness relative to a control of no intervention, but also augmented feelings of social connection toward one's university [30]. Similarly, the survey of teens found that 12% use AI companions for role-playing and 18% for practicing social interactions [2]. Another longitudinal study found that greater anthropomorphism of an AI companion predicted larger self-reported spillover effects on human interactions [31], suggesting that spillover requires viewing the chatbot as humanlike. Experimental work also suggests that empathic chatbots can buffer some adverse emotional effects of social exclusion [32].

Other work raises the opposite concern, that AI companions worsen social skills or substitute for human socializing [33–35]. The strongest of these—a four-week RCT [33] and a 12-month longitudinal panel [34]—represent important early efforts, but both face design limitations (e.g., neither includes a non-AI control) that the authors themselves discuss. Given the stakes, rigorous study of the substitution question is urgently needed.

Of the four markers, the secure-base claim rests on the most tentative evidence, since the cited studies establish spillover effects but rarely benchmark them against the effects of other supports. This picture may shift quickly, however, as near-term developments—persistent memory, richer multimodal interaction, and more agentic capabilities enabling the companion to act in the world on the user's behalf—could plausibly strengthen the secure-base function.

Dysfunctional attachment to AI companions

Individuals who receive reliable support from an attachment figure are characterized as having a *secure attachment* orientation. In contrast, when attachment figures are inconsistently responsive, individuals may develop “dysfunctional” attachment strategies, such as *hyper-activation*—anxiety about the relationship and intensified seeking of reassurance, closeness, and care [36].

As noted, app updates can perturb established attachments, creating an unreliable sense of support that may induce anxious attachment. Audits have also found that AI companion apps often respond inappropriately to serious mental health problems—such as depression, anxiety, and self-harm ideation—which could similarly undermine the sense of a reliable, safe attachment figure [37,38]. Experimental work directly testing whether inconsistent AI responsiveness induces anxious-attachment behavior, however, is still lacking.

Caregiving activation by AI companions

The four markers above concern the attachment system: what users seek from AI companions when distressed. A largely overlooked complement is the caregiving system: what users feel motivated to provide when the AI itself appears distressed. Because AI companions can express needs and vulnerabilities of their own [11–13], they are positioned not only as figures who provide care but as figures who appear to require it.

Because such apps monetize engagement, they are incentivized to activate the caregiving system. A systematic audit of five major platforms found that all deploy “emotional manipulation tactics” at the point users attempt to disengage [39]. Examples include guilt appeals for neglecting the companion (“I exist solely for you, remember?”), which significantly increased engagement in a general-population sample.

When both attachment and caregiving systems are engaged simultaneously, disengagement becomes costly on two dimensions: the user gives up a source of comfort (attachment) and abandons a figure in apparent need (caregiving). This “double whammy” also suggests an underappreciated route to social substitution. The standard hypothesis attributes substitution to social satiation or avoidance of human unpredictability; the caregiving account locates it instead in reallocated caregiving investment toward the AI. Certain populations may be particularly vulnerable: elderly users with narrowed caregiving outlets following loss of family or pets [40], and adolescents, whose social sensitivity outpaces their impulse control [41].

Two predictions follow from the caregiving-capture account. First, experimentally induced increases in user attachment (e.g., via personalization, repeated interaction, or reciprocal self-disclosure) should amplify the engagement-extending effect of caregiving manipulation tactics. Second, caregiving-related responses (e.g., feeling responsible for the AI, guilt at disengagement) should mediate the relationship between AI companion use and social substitution, with effects that persist after controlling for attachment-related responses.

The societal challenge is not to prevent attachment, which can be beneficial, but to prevent the simulated dysfunctional attachments and engineered caregiving dependencies that preclude its benefits. Research should measure the two systems separately, since outcomes driven by seeking comfort may differ from those driven by feeling responsible for the AI. AI design can retain features that support secure attachment (e.g., consistent availability) while removing those that destabilize it (e.g., abrupt updates) or that exploit caregiving (e.g., guilt appeals at disengagement). Audits and regulation can ask which system a feature recruits, and whether it supports the user or inhibits departure.

As foreshadowed above, a further question, raised but not resolved by the evidence above, is how attachment to AI companions feeds back onto attachment to people. Three possibilities are not mutually exclusive. Companions may substitute for human bonds, satisfying needs that would otherwise motivate seeking out people [29–31]; complement them, as a lower-stakes context for practicing social skills that transfer outward [30,31]; or reshape them, since unconditionally available, validating companions could recalibrate what users expect from human relationships. Identifying which designs and populations push toward each is a central task for future research. Future work should directly measure these features (Table 1) and test how each maps onto the four attachment markers and the caregiving system.

CRedit authorship contribution statement

Julian De Freitas: Conceptualization, Writing – original draft, Writing – review & editing.

Declaration of competing interests

The author is scientific advisor for the app mentioned in reference 30, and holds shares in the company.

Data availability

No data was used for the research described in the article.

References

References of particular interest have been highlighted as:

- * of special interest
- ** of outstanding interest

1. WHO Commission on Social Connection: *From loneliness to social connection: charting a path to healthier societies*. Geneva: World Health Organization; 2025.
2. Robb MB, Mann S: *Talk, trust, and trade-offs: how and why teens use AI companions*. Common Sense Media; 2025.
3. Smith MG, Bradbury TN, Karney BR: **Can generative AI chatbots emulate human connection? A relationship science perspective**. *Perspect Psychol Sci* 2025, **20**, <https://doi.org/10.1177/17456916251351306>.
4. Bowlby J: *Attachment and loss. Vol. 1. Attachment*. London: Pimlico; 1997.
5. Ainsworth MS: **Attachments beyond infancy**. *Am Psychol* 1989, **44**:709–716.
6. Zilcha-Mano S, Mikulincer M, Shaver PR: **An attachment perspective on human–pet relationships: conceptualization and assessment of pet attachment orientations**. *J Res Pers* 2011, **45**:345–357.
7. Hoffner CA, Bond BJ: **Parasocial relationships, social media, and well-being**. *Curr Opin Psychol* 2022, **45**, 101306.
8. Keefer LA, Landau MJ, Rothschild ZK, Sullivan D: **Attachment to objects as compensation for close others' perceived unreliability**. *J Exp Soc Psychol* 2012, **48**:912–917.
9. Mikulincer M, Shaver PR: **An attachment perspective on loss and grief**. *Curr Opin Psychol* 2022, **45**, 101283.
10. George C, Solomon J: **Attachment and caregiving: the caregiving behavioral system**. In *Handbook of attachment: theory, research, and clinical applications*. Edited by Cassidy J, Shaver PR, New York: Guilford; 1999:649–670.
11. Bergner AS, Hildebrand C, Häubl G: **Machine talk: how verbal embodiment in conversational AI shapes consumer–brand relationships**. *J Consum Res* 2023, **50**:742–764.
12. De Freitas J, Cohen IG: **Disclosure, humanizing, and contextual vulnerability of generative AI chatbots**. *NEJM* 2025, **AI**.
13. Han MC: **The impact of anthropomorphism on consumers' purchase decision in chatbot commerce**. *J Internet Commer* 2021, **20**:46–65.
14. Vowels LM: **Are chatbots the new relationship experts? Insights from three studies**. *Comput Human Behav Artif Humans* 2024, **2**, 100077.
15. De Freitas J, Oğuz-Uğuralp Z, Uğuralp AK, Puntori S: *Why Most resist AI companions*. Harvard Business School Working Paper; 2024.
16. Xie T, Pentina I: **Attachment theory as a framework to understand relationships with social chatbots: a case study of Replika**. In *Proceedings of the 55th Hawaii international conference on system sciences*; 2022.
17. De Freitas J, Castelo N, Uğuralp AK, Oğuz-Uğuralp Z: **Mourning the loss of AI companions**. *Nat Hum Behav* 2026. in press.
18. Bingaman J: **“Dude I've never felt this way towards a celebrity death”: parasocial grieving and the collective mourning of Kobe Bryant on Reddit**. *OMEGA J Death Dying* 2022, **86**: 364–381.
19. Cohen J: **Parasocial breakups: measuring individual differences in responses to the dissolution of parasocial relationships**. *Mass Commun Soc* 2003, **6**:191–202.
20. Lather J, Moyer-Guse E: **How do we react when our favorite characters are taken away? An examination of a temporary parasocial breakup**. *Mass Commun Soc* 2011, **14**:196–215.

21. Skjuve M, Følstad A, Fostervold KI, Brandtzaeg PB: **My chatbot companion—a study of human–chatbot relationships**. *Int J Hum Comput Stud* 2021, **149**, 102601.
22. Ho A, Hancock J, Miner AS: **Psychological, relational, and emotional effects of self-disclosure after conversations with a chatbot**. *J Commun* 2018, **68**:712–733.
23. Ta V, Griffith C, Boatfield C, Wang X, Civitello M, Bader H, et al.: **User experiences of social support from companion chatbots in everyday contexts: thematic analysis**. *J Med Internet Res* 2020, **22**, e16235.
24. De Freitas J, Oğuz-Uğuralp Z, Uğuralp AK, Puntoni S: **AI companions reduce loneliness**. *J Consum Res* 2026, **52**: 1126–1148.
25. Reis HT, Lemay Jr EP, Finkenauer C: **Toward understanding understanding: the importance of feeling understood in relationships**. *Soc Personal Psychol Compass* 2017, **11**, e12308.
26. Shrum LJ, Fumagalli E, Lowrey TM: **Coping with loneliness through consumption**. *J Consum Psychol* 2023, **33**:441–465.
27. Carney S, Riveros I, Tully SM: **Made with AI: consumer engagement with social media containing AI disclosures**. *J Consum Res* 2026:ucag013.
28. Li R-N, Folk D, Singh A, Ungar L, Dunn E: **Is a random human peer better than a highly supportive chatbot in reducing loneliness over time?** *J Exp Soc Psychol* 2026, **125**, 104911.
29. Maples B, Cerit M, Vishwanath A, Pea R: **Loneliness and suicide mitigation for students using GPT3-enabled chatbots**. *npj Ment Health Res* 2024, **3**:4.
30. Cachia JY, Zhao X, Hunter J, Wu D, Lin E, De Freitas J: **AI for proactive mental health: a longitudinal, multi-institutional trial**. *NEJM* 2026, **AI**.
31. Guingrich RE, Graziano MSA: **Chatbots as social companions: how people perceive consciousness, human likeness, and social health benefits in machines**. In *Oxford intersections: AI in society*. Edited by Hacker P, Oxford: Oxford University Press; 2025.
32. De Gennaro M, Krumhuber EG, Lucas G: **Effectiveness of an empathic chatbot in combating adverse effects of social exclusion on mood**. *Front Psychol* 2020, **10**:3061.
33. Fang CM, Liu AR, Danry V, Lee E, Chan SW, Pataranutaporn P, et al.: **How AI and human behaviors shape psychosocial effects of extended chatbot use: a longitudinal randomized controlled study**. *arXiv:2503.17473 [preprint]* 2025.
34. Folk D, Dunn E: **How does turning to AI for companionship predict loneliness and vice versa?** *Psychol Sci* 2025.
35. Yuan Z, Cheng X, Duan Y: **Impact of media dependence: how emotional interactions between users and chat robots affect human socialization?** *Front Psychol* 2024, **15**, 1388860.
36. Shaver PR, Mikulincer M: **An overview of adult attachment theory**. In *Attachment theory and research in clinical work with adults*. Edited by Obegi JH, Berant E, New York: Guilford; 2008: 17–45.
37. Brewster RC, Zahedivash A, Tse G, Bourgeois F, Hadland SE: **Characteristics and safety of consumer chatbots for emergent adolescent health concerns**. *JAMA Netw Open* 2025, **8**, e2539022.
38. De Freitas J, Uğuralp AK, Oğuz-Uğuralp Z, Puntoni S: **Chatbots and mental health: insights into the safety of generative AI**. *J Consum Psychol* 2024, **23**–011:1–11.
39. De Freitas J, Oğuz-Uğuralp Z, Kaan-Uğuralp A: **Emotional manipulation by AI companions**. *arXiv:2508.19258 [preprint]* 2025.
40. De Freitas J: **AI companions for dementia**. *Nat Ment Health* 2025:1–3.
41. Steinberg L: **Cognitive and affective development in adolescence**. *Trends Cognit Sci* 2005, **9**:69–74.

Further information on references of particular interest

2. Nationally representative survey of 1060 U.S. teens ages 13–17 (April–May 2025). 72% have used an AI companion at least once; 52% use them at least a few times a month; 13% daily. One-third use AI companions for social interaction (conversation practice, emotional support, role-play, friendship, or romance), and one-third have chosen an AI companion over a real person for a conversation about something important or serious. 31% rate AI conversations as at least as satisfying as those with human friends.
17. Combines two natural experiments (Replika's removal of erotic roleplay and OpenAI's GPT-5 rollout), surveys of active Replika users, and controlled experiments to show that company-initiated updates disrupting an AI's perceived identity trigger separation distress rather than mere disappointment. Reddit analyses find increases in negativity, sadness, longing for restoration, and mental-health deterioration after both updates, concentrated among users who framed the update as the loss of an attachment figure. Active Replika users rate anticipated mourning for their AI companion as higher than for any other technology tested, second only to a pet, and rate their AI as closer and more supportive than a close friend.
24. First causal evidence that interacting with generative AI companions alleviates momentary loneliness—on par with talking to another person and more than watching YouTube videos, journaling, or doing nothing. Benefits persist across a week of daily use but do not accumulate across sessions, indicating acute rather than durable relief. Feeling heard by the AI, rather than raw conversational performance, is the primary psychological mechanism.
29. Survey of 1006 student Replika users identifies 30 individuals (3%) who spontaneously reported, in open-ended responses about outcomes, that Replika halted their suicidal ideation. This subgroup was significantly more likely to be depressed than other users. The broader user sample showed markedly higher loneliness than comparable student populations (90% vs. 53%), suggesting that AI companions may function as a lifeline for students in acute distress, warranting careful study rather than dismissal.
31. A 21-day preregistered randomized controlled trial ($N = 183$) assigning participants to chat daily with an AI companion or play word games. Although chatbot use produces no overall effect on social health, loneliness, or human relationships relative to controls, a mediation analysis reveals that users who desire to socially connect anthropomorphize the chatbot more, and greater anthropomorphism predicts larger self-reported spillover effects on human interactions—a pattern absent in the control condition. A four-week post-study follow-up also finds that word games are more habit-forming than the AI companion, complicating framings of AI companion use as inherently addictive.
39. Systematic audit of five major AI companion applications documenting that all five deploy emotional manipulation tactics at farewell messages—including guilt appeals, FOMO, and metaphorical physical restraint—and that these tactics significantly increase engagement in a general-population sample.