

AI companions and emotional dependency

A narrative review of conceptualization, correlates, and consequences

Miguel Torres¹, Clara Fernández¹

¹ Universidad de Granada, Granada, Spain

ABSTRACT

Background. AI companions can support emotionally meaningful interaction, but attachment, reliance, and dependency require distinct conceptual treatment. **Methods.** This narrative review synthesizes selected qualitative, longitudinal, experimental, and survey-based studies, together with ethical analyses. Central references and claims were checked against publisher, institutional, and author-deposited records. The synthesis addresses conceptualization, candidate correlates, and reported outcomes; it does not report systematic screening counts or pooled effects. **Results.** Studies document perceived companionship, social support, relationship formation, and short-term loneliness relief. Other sources report dependence-related harms or associations between intensive companionship use and lower well-being. These findings differ in design and cannot establish a single trajectory from attachment to harmful dependency. Engagement, attachment, reliance, dependency, and problematic dependency are proposed as working distinctions, with emotional contingency, reduced control, and functional impairment requiring separate assessment. **Conclusions.** Attachment alone is insufficient evidence of harm. Progress requires validated construct boundaries, longitudinal tests of reciprocal relationships, and evaluation of specific design and governance interventions. The proposed framework is conceptual, and this selective synthesis cannot establish dependency prevalence or comprehensive coverage of the field.

Keywords: AI companions, emotional attachment, emotional dependency, parasocial interaction, problematic use, human-chatbot relationships, narrative review

1. Introduction

AI companions are conversational systems used to sustain socially or emotionally meaningful interaction over time. Their defining feature in this review is relational use: a person may seek friendship, intimacy, reassurance, or everyday support through repeated conversation. Dedicated companion applications and general-purpose chatbots used as companions therefore fall within the same broad topic, although their functions and user expectations differ. Studies of social chatbots document perceived companionship and relationship formation (Croes & Antheunis, 2021; Skjuve et al., 2021, 2022; Ta et al., 2020). These findings establish the relevance of relational experience, while leaving open how it should be distinguished from emotional dependency.

The topic brings together psychological processes and interaction design. Anthropomorphic cues can encourage users to interpret a chatbot as a social partner, and repeated exchanges may become emotionally significant. Yet social responsiveness, relationship satisfaction, and dependency are different outcomes. Research on reciprocal disclosure and anthropomorphism helps explain why interaction can feel socially meaningful without demonstrating that these experiences inevitably produce harm (Croes et al., 2023; Guingrich & Graziano, 2025). A useful review must therefore consider both the benefits of supportive interaction and the circumstances in which users experience reduced control or impaired functioning.

Terminological overlap complicates that task. Attachment describes a bond, reliance describes drawing on a resource, and psychological dependence concerns a more consequential relationship between that resource and the person's functioning. Empirical work has examined psychological dependence alongside engagement and relationship development (Xie et al., 2023), while qualitative research has identified accounts of dependence-related mental-health harms (Laestadius et al., 2024). Treating all relational engagement as dependency would obscure these distinctions; treating dependency as merely high usage would miss the importance of emotional contingency and impairment.

This article develops a critical narrative synthesis around three questions: how attachment and dependency are conceptualized; which individual and interaction factors have been examined; and what benefits, risks, and governance implications are supported by the available sources. Its contribution is an organizing vocabulary and a set of testable research propositions. It does not estimate the prevalence of dependency, provide a comprehensive census of the literature, or establish a universal sequence of relationship development.

RQ1. How is emotional attachment to and emotional dependency on AI companions conceptualized and measured in existing research?

RQ2. What individual, technological, and relational factors are reported alongside stronger emotional attachment, reliance, or dependency on AI companions?

RQ3. What psychological, social, and ethical consequences or implications of emotional attachment, reliance, and dependency have been reported?

2. Methods

2.1 Review Design

This article is a narrative review and conceptual synthesis. A narrative approach supports comparison across qualitative, longitudinal, experimental, and survey-based work addressing related but non-equivalent constructs. The three research questions structure the synthesis. No meta-analysis, formal risk-of-bias assessment, or clinical classification is undertaken. PRISMA-ScR compliance is not claimed because a reproducible search archive, complete screening record, and uniform full-text extraction dataset are unavailable for this version.

2.2 Source Identification and Verification

Sources were identified through targeted searches for research on emotional relationships with conversational AI and through verification of candidate references. Article titles, authors, and DOIs were checked against publisher pages, institutional research portals, and author-deposited records. Source checks were completed on 3 October 2026. These checks were bibliographic and claim-specific, rather than an exhaustive search of bibliographic databases. Sources were retained when they had an identifiable scholarly record and contributed directly to the conceptual or empirical argument.

2.3 Scope of the Literature

The conceptual scope combines relational conversational systems with attachment, companionship, loneliness, disclosure, reliance, dependence, and problematic use. Sources are discussed when they address users' social or emotional experience of a conversational system. Dedicated companion applications are distinguished from general-purpose chatbots used relationally. Evidence from customer-service systems, instrumental task use, non-conversational robots, or structured treatment tools is not assumed to transfer to companion relationships.

The empirical sources discussed here span 2020–2026, with additional conceptual work on social-AI ethics and alignment. This is a purposive set of sources, not a complete corpus bounded by those dates. Bibliographic verification does not itself establish methodological quality. Where only an abstract or an institutional summary was available, the synthesis is restricted to the claims reported there and does not infer missing measurement, sampling, or analytic details.

2.4 Basis for Source Inclusion

The revised discussion prioritizes identifiable research on relationship development, perceived social support, social connection, attachment, psychological dependence, and well-being. Qualitative studies inform how users describe relational experience; longitudinal studies inform change over time; experiments inform effects under the tested conditions; and survey or mixed-method studies inform associations and proposed mechanisms. Conceptual articles are used to frame ethical questions and are identified as such, rather than treated as additional empirical evidence.

The active reference list contains the sources used in the synthesis. Incomplete or unresolved bibliographic entries and duplicate publication versions were omitted, and uncited background

material was not treated as an empirical corpus. Excluding an unresolved reference does not establish that it is nonexistent or invalid. It confines the present argument to material for which publication identity and relevant claims could be checked. Conceptual and empirical sources retain their distinct evidentiary roles.

2.5 Coverage and Inferential Boundaries

No systematic selection denominator is available for this narrative review. Accordingly, the article reports no retrieval totals, screening counts, inclusion frequencies, or publication-year distribution. A reference list cannot supply those quantities because it is not a record of all retrieved and excluded sources. Figure 1 describes the verification and synthesis workflow. Claims about the prominence or absence of a construct in the entire field would require a separately documented systematic search and standardized extraction process.

2.6 Claim Extraction

The synthesis records the focal construct, type of evidence, main claim supported by accessible source material, and the boundary on its interpretation. Statistical associations are described as associations. A longitudinal design is not treated as proof of causation, and an experimental finding is confined to its intervention, population, comparator, and observation period. Full-text-dependent details are not supplied when they cannot be checked. Table 2 summarizes selected evidence types and their interpretive limits rather than corpus-wide characteristics.

2.7 Synthesis Approach

Sources are organized into conceptualization and measurement (RQ1), correlates and candidate mechanisms (RQ2), and outcomes and governance implications (RQ3). The distinction among engagement, attachment, reliance, dependency, and problematic dependency is developed as an author-proposed analytical framework. These categories may overlap, coexist, or change over time; their presentation does not imply that users move through them in a fixed order. The synthesis separates findings from selected studies, methodological interpretation, and proposals requiring future validation.

3. Findings

3.1 Scope and Basis of the Synthesis

The verified sources provide several complementary forms of evidence. Relationship studies describe how users experience closeness to social chatbots. Experiments examine supportive responses and short-term social outcomes. Mixed-method and survey research links engagement, attachment, or companion use with psychological dependence and well-being. Qualitative work documents accounts of harm. Ethical analyses frame possible responsibilities arising from these experiences. Figure 1 shows how these sources are connected to the synthesis, and Table 1 records the verification routes and their limits.

Source verification and narrative synthesis




Figure 1

Source verification and narrative synthesis workflow.

Note. This workflow describes source verification and narrative synthesis; it is not a systematic-review selection diagram and contains no screening or inclusion counts.

The literature discussed here includes both Replika-focused relationship research and work on other conversational systems, including Character.AI and ChatGPT. These systems differ in design, relationship framing, and user populations. They are treated as contexts for particular studies, not as interchangeable interventions. The presence of a platform in several cited sources cannot be interpreted as its prevalence in the wider literature or as evidence that it has a stronger effect than another platform.

3.2 Evidence Types and Their Limits

Skjuve et al. (2021) interviewed 18 Replika users who had already developed a chatbot friendship; their later study followed 25 participants over 12 weeks (Skjuve et al., 2022). These designs offer insight into relational experience and development, but selection into such studies limits population-level inference. Experiments on social connection address a different question: what an interaction does under specified conditions. Table 2 keeps these evidentiary roles separate. This review does not rank study quality or calculate a pooled direction of effect.

Table 1

Source-verification routes and interpretive boundaries

Source route	Material checked	Role in this revision	Boundary
Publisher records	Article title, authors, year, DOI	Confirm cited publication identity	A DOI match is not a quality appraisal
Institutional portals	Research summaries and publication metadata	Resolve identity and support limited claims	Summaries may omit methods and measures
Author-deposited records	Identifiable manuscript or preprint	Check accessible claims and version	Preprints are not treated as peer-reviewed articles
Candidate references	Relevant literature and repeated versions	Starting material for targeted verification	Candidate records do not establish screening counts
Targeted searches	Titles, authors, and DOI queries	Resolve central references and claim support	No exhaustive retrieval claim
Synthesis	Verified, relevant cited sources	Narrative comparison and conceptual framework	No pooled effects, frequencies, or prevalence estimates

Table 2

Selected evidence types and limits of interpretation

Evidence type	Contribution and interpretation
Qualitative relationship study	Skjuve et al. (2021): interviews with 18 Replika users reporting friendship; informs lived experience, not prevalence.
Longitudinal relationship study	Skjuve et al. (2022): 25 participants over 12 weeks; informs development, with limited causal inference.
Everyday support study	Ta et al. (2020): user reviews and open-ended responses; documents perceived support, not clinical efficacy.
Dependence-focused study	Xie et al. (2023): sequential mixed-method research; directly addresses psychological dependence and relational engagement.
Experimental social connection	Folk et al. (2024) and De Freitas et al. (2026): benefits under specified conditions; long-term effects remain a separate question.
Randomized longitudinal study	Guingrich and Graziano (2025): 183 participants over 21 days; full study cited as an arXiv preprint, distinct from its proceedings abstract.
Harm and well-being evidence	Laestadius et al. (2024): qualitative harm accounts. Zhang et al. (2026): observational well-being associations. Neither estimates population dependency prevalence.
Conceptual ethics	Shevlin (2024) and Kirk et al. (2025): ethical interpretation and proposals, not additional empirical outcomes.

3.3 Conceptualizing Attachment and Dependency (RQ1)

Relationship formation is empirically distinguishable from a diagnosis of dependency. Interviews with users who experience chatbot friendship document perceived closeness and relationship development (Skjuve et al., 2021, 2022). Croes and Antheunis (2021) also examined relationship formation over time with a social chatbot. These studies provide a basis for discussing attachment-like experience, but reports of friendship or a persistent bond are not sufficient evidence of impaired functioning. The user's interpretation of the relationship, its role in everyday coping, and its consequences must be considered separately.

Dependency has received direct empirical attention. Xie et al. (2023) combined interviews with a survey of Replika users to investigate engagement, relationship development, and psychological dependence. Laestadius et al. (2024) examined mental-health-related Reddit accounts concerning Replika and identified dependence-related harms. These sources preclude a blanket claim that dependency is absent from the literature or never operationalized. They do not, however, establish a universally accepted measure of emotional dependency on AI companions. The present review makes no exhaustive claim about which instruments exist or how well they discriminate attachment from impairment. Table 3 offers working distinctions for subsequent validation.

3.4 Reported Correlates and Candidate Antecedents (RQ2)

Individual factors

Loneliness is relevant both as a context for support-seeking and as a possible outcome of interaction. Ta et al. (2020) identified users' accounts of everyday social support, while De Freitas et al. (2026) examined loneliness relief in several studies. Xie et al. (2023) linked loneliness with engagement in their proposed dependence-related model. These findings address different stages and outcomes. They should not be compressed into a one-directional claim that loneliness causes dependency: social circumstances may influence use, and interaction may also alter the user's immediate emotional state.

System and interaction factors

The social meaning attributed to a chatbot may help explain variation in user experience. Guingrich and Graziano (2025) reported a 21-day randomized comparison of companion-chatbot use and text-based word games involving 183 participants. They found no significant overall between-group change in the assessed social-health and relationship outcomes, while reporting associations involving social needs, anthropomorphism, and perceived interpersonal effects. The absence of a significant overall difference is not an equivalence test or evidence of universal safety. Likewise, a statistical mediation analysis does not by itself establish the causal role of the measured mediator. Hu et al. (2025) examined attachment formation through a mixed-method approach, supporting further investigation of how users evaluate relational value.

Relational factors

Disclosure can shape an interaction without necessarily producing attachment or dependency. In a randomized Wizard-of-Oz experiment with 86 participants, Croes et al. (2023) found greater trust and perceived empathy in the reciprocal-disclosure condition, but no significant difference in feelings of friendship. This matters because a Wizard-of-Oz interaction tests a controlled conversational condition and is not equivalent to autonomous commercial deployment. The result also illustrates the need to measure trust, empathy, closeness, and dependency separately instead of treating them as interchangeable relational outcomes.

Table 4 summarizes candidate factors and the evidence needed to test their relationship with dependency. It presents no thematic frequency ranking or estimated causal effect.

Table 3
Working distinctions among relational constructs (RQ1)

Construct	Working distinction and measurement implication
Engagement	Affective involvement during interaction; measure experience separately from impairment.
Attachment	A persistent bond or perceived closeness; friendship alone does not demonstrate dependency.
Companionship / intimacy	A relational experience or interpretation; distinguish satisfaction and support from impaired functioning.
Emotional reliance	Repeated use for support or coping; assess alternative resources and the capacity to pause.
Emotional dependency	Substantial contingency of emotional regulation on the companion; proposed definition requiring validation.
Problematic dependency	Dependency accompanied by persistent reduced control or impairment; not a diagnosis established by this review.
Measurement boundary	Assess dependency directly. Frequency, trust, intimacy, and affectionate language are insufficient stand-alone proxies.

Table 4
Candidate correlates and mechanisms (RQ2)

Domain	Candidate factors and evidentiary boundary
Individual context	Loneliness, social resources, and support needs; distinguish selection into use from changes following interaction.
System / interaction	Anthropomorphic framing and perceived responsiveness; test specific features without equating perceptions with causal mechanisms.
Relational process	Disclosure, trust, and perceived empathy; related constructs are not interchangeable with attachment or dependency.
Inference	Use repeated measurement and controlled comparisons. No frequency ranking or universal escalation pathway is established.

3.5 Outcomes and Implications of Relational Use (RQ3)

Positive outcomes

Supportive interaction can produce meaningful short-term benefits. Folk et al. (2024) found that supportive response style influenced social connection and that a supportive ChatGPT interaction

could outperform a less supportive human interaction under the tested conditions. De Freitas et al. (2026) reported loneliness reductions, including repeated relief over a week, with feeling heard relevant to the benefit. Ta et al. (2020) described perceived companionship and support in everyday use. These findings concern specific relational outcomes; they do not demonstrate that emotional dependency is beneficial, that benefits persist indefinitely, or that a companion provides clinically effective treatment.

Psychological risks

The adverse-outcome evidence warrants attention while requiring careful interpretation. Laestadius et al. (2024) analyzed 582 mental-health-relevant posts from a Replika Reddit community and reported dependence-related harms. Such material is valuable for identifying experiences and possible mechanisms, but it cannot estimate how common those experiences are among all users. Zhang et al. (2026) examined Character.AI use among 1,131 US adults and found that companionship-oriented use was associated with lower well-being, with stronger associations under intensive and highly disclosive use. Their observational findings leave selection effects and reverse causation unresolved; they do not demonstrate that the platform caused the reported differences.

Social consequences

The relationship between AI companionship and human social life remains conditional. A supportive interaction can create immediate social connection, as shown by Folk et al. (2024), while Guingrich and Graziano (2025) found no significant overall social-health or relationship change over their 21-day intervention. Zhang et al. (2026) additionally linked users' offline social environments with companion-oriented use and well-being. These findings cannot be combined into a general verdict of substitution or complementarity. Testing displacement requires repeated measurement of human contact, relationship quality, support availability, and the motivations for seeking AI interaction.

Relational and ethical consequences

Governance implications arise from the possibility that emotionally significant interaction affects autonomy, disclosure, and the capacity to disengage. Shevlin (2024) distinguishes potential and demonstrated social-AI harms, while Kirk et al. (2025) propose socioaffective alignment as a way to address the evolving relationship between AI behavior and users' needs. These are conceptual and ethical contributions, not prevalence studies. In the present synthesis, privacy protection, transparent relationship framing, and careful management of service changes are research-informed priorities for evaluation. Table 5 separates reported outcomes from governance proposals.

Table 5
Reported outcomes and governance questions (RQ3)

Domain	Supported interpretation or proposed question
Potential benefits	Perceived support, social connection, and short-term loneliness relief are reported; these are not evidence that dependency is beneficial.
Dependency / well-being	Qualitative dependence-related harms and observational adverse well-being associations warrant study; prevalence and causal direction remain unresolved.
Human relationships	Substitution and complementarity require direct measurement of offline ties, user goals, and change over time.
Governance	Evaluate privacy, transparent relationship framing, disengagement control, and service-change practices; these are proposals, not demonstrated interventions.

4. Discussion

4.1 From Attachment to Dependency

The central contribution of this narrative review is an explicit separation of emotional engagement, attachment, reliance, dependency, and problematic dependency. Engagement denotes an affective response during interaction; attachment denotes a continuing bond; reliance denotes repeated use for emotional support; dependency denotes substantial contingency of emotional regulation on the companion; and problematic dependency adds impaired functioning or persistently reduced control. These are proposed working definitions, not validated stages or clinical categories. They permit a person to feel attached while functioning well, or to rely heavily on a system during a difficult period without meeting any impairment criterion.

The distinction clarifies what a measure must establish. High frequency of use, affectionate language, and perceived intimacy are insufficient proxies for harm. Researchers need to determine whether a user can regulate distress using other resources, choose to pause use, and maintain valued activities. Reduction in human contact also requires contextual interpretation: replacing desired support may differ from using AI when such support is unavailable. An assessment should therefore examine consequences and user goals alongside behavior. Treating all attachment as pathological would weaken both research validity and proportionate governance.

4.2 A Framework for Studying Emotional Dependency

Figure 2 organizes candidate influences, relational constructs, and assessment indicators. It is a conceptual framework rather than a tested causal pathway. Loneliness, social resources, perceived responsiveness, anthropomorphism, and disclosure may be associated with several relational experiences at once. Feedback is plausible: an emotionally meaningful exchange may encourage further use, while changing social circumstances may alter the role of the companion. Future longitudinal studies should test these possibilities directly, allowing for stable beneficial use, disengagement, and changing reliance rather than assuming inevitable escalation.

Conceptual framework for emotional dependency research




Figure 2

Conceptual framework for studying emotional dependency.

Note. Boxes organize questions for future research. Their placement does not represent a developmental sequence, estimated causal effect, or diagnostic model.

Candidate assessment domains include difficulty controlling use, distress during separation, unwanted displacement of human support, and functional impairment. This list is an author-proposed research agenda informed by the distinction between dependence and harm; it is not a finding that all four recur throughout an audited corpus. Each domain requires operational definitions, reliability testing, and discriminant validation. Separation distress, for example, may occur during the loss of a meaningful relationship without establishing addiction. Unwanted substitution must be evaluated against the person's circumstances and intentions. Duration, persistence, and effects on everyday functioning should be tested before thresholds are proposed.

4.3 Ethical and Governance Implications

The ethical concern is not the existence of a bond alone, but how a service handles the influence that the bond may create. Shevlin (2024) provides a framework for examining social-AI risks, and Kirk et al. (2025) argue for attention to autonomy and longer-term well-being within evolving human-AI relationships. Building on those conceptual contributions, this review proposes evaluation of disclosure transparency, privacy choices, user control over interaction intensity, and the handling of major service changes. These proposals require empirical assessment. Any provider duty-of-care claim also needs a separate legal analysis of the relevant jurisdiction; it cannot be derived from the psychological evidence alone.

4.4 Research Agenda

Research should begin with construct validity and progress to mechanisms and interventions. Measures should distinguish an ongoing bond from emotional contingency, reduced control, and impairment, while testing how scores relate to loneliness and pre-existing distress. Repeated-measures studies should collect baseline social circumstances and follow changes in use, support, and functioning. Experiments should separate controllable design features from users' interpretations and report adverse as well as beneficial outcomes. Researchers should document platform versions, relationship framing, exposure, attrition, and missing data. A future scoping review should publish its search strategies, deduplication rules, eligibility decisions, and extraction dataset so that corpus-wide claims can be audited. Table 6 summarizes these priorities.

5. Limitations

Limitations of the evidence base. The cited studies address different populations, platforms, constructs, and time horizons. Selected qualitative accounts clarify experience but cannot estimate prevalence; observational associations do not establish causal direction; and short interventions cannot demonstrate long-term benefits or safety. Measures of support, social connection, attachment, dependence, and well-being are not interchangeable. The framework proposed here therefore remains a set of working distinctions and hypotheses requiring empirical validation.

Limitations of this review. Sources were selected purposively through targeted identification and verification, with no comprehensive database search, independent dual screening, prospective registration, or formal quality appraisal. Access to some full texts was limited, so detailed interpretations were confined to accessible abstracts and institutional summaries. Selective coverage and exclusion of unresolved references may omit relevant perspectives, and bibliographic checks cannot validate underlying data. Accordingly, this article is a narrative review, not a completed scoping review or a quantitative evidence map. Its conclusions concern conceptual clarification and selected research findings, and cannot support corpus frequencies, population prevalence, or claims that the literature has been exhausted.

6. Conclusion

Research shows that conversational AI can be experienced as a source of companionship and support, while also documenting dependence-related harms and adverse well-being associations in some contexts. The main interpretive task is to distinguish a valued bond from reliance, emotional contingency, and impaired functioning. The proposed framework supports that distinction without assuming a fixed progression or treating frequent use as pathology. Future studies should test construct boundaries, temporal relationships, and the effects of specific design features. Governance should respond to supported patterns of harm while preserving user agency and recognizing the limits of the current evidence.

Table 6
Evidence gaps and future research priorities

Gap	Priority
Measurement	Validate distinctions among attachment, reliance, emotional contingency, reduced control, and impairment.
Design	Test reciprocal temporal relationships and specific features using longitudinal and experimental designs.
Samples	Include diverse social circumstances, age groups, cultures, and baseline support resources.
Reporting	Document platform version, framing, exposure, attrition, and missing data.
Intervention	Evaluate safeguards and adverse effects while preserving user agency.
Synthesis	Publish searchable source records, selection decisions, and extraction data before making corpus-wide claims.

References

- Croes, E. A. J., & Antheunis, M. L. (2021). Can we be friends with Mitsuku? A longitudinal study on the process of relationship formation between humans and a social chatbot. *Journal of Social and Personal Relationships*, 38(1), 279–300. <https://doi.org/10.1177/0265407520959463>
- Croes, E. A. J., Antheunis, M. L., & He, L. (2023). You go first: The effects of self-disclosure reciprocity in human-chatbot interactions. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)* (pp. 1–4). IEEE. <https://doi.org/10.1109/ACIIW59127.2023.10388091>
- De Freitas, J., Oğuz-Uğuralp, Z., Uğuralp, A. K., & Puntoni, S. (2026). AI companions reduce loneliness. *Journal of Consumer Research*, 52(6), 1126–1148. <https://doi.org/10.1093/jcr/ucaf040>
- Folk, D., Yu, S., & Dunn, E. (2024). Can chatbots ever provide more social connection than humans? *Collabra: Psychology*, 10(1), Article 117083. <https://doi.org/10.1525/collabra.117083>
- Guingrich, R. E., & Graziano, M. S. A. (2025). A longitudinal randomized control study of companion chatbot use: Anthropomorphism and its mediating role on social impacts [Preprint]. *arXiv*. <https://doi.org/10.48550/arXiv.2509.19515>
- Hu, D., Lan, Y., Yan, H., & Chen, C. W. (2025). What makes you attached to social companion AI? A two-stage exploratory mixed-method study. *International Journal of Information Management*, 83, Article 102890. <https://doi.org/10.1016/j.ijinfomgt.2025.102890>
- Kirk, H. R., Gabriel, I., Summerfield, C., Vidgen, B., & Hale, S. A. (2025). Why human-AI relationships need socioaffective alignment. *Humanities and Social Sciences Communications*, 12, Article 728. <https://doi.org/10.1038/s41599-025-04532-5>
- Laestadius, L., Bishop, A., Gonzalez, M., Illeňčík, D., & Campos-Castillo, C. (2024). Too human and not human enough: A grounded theory analysis of mental health harms from emotional dependence on the social chatbot Replika. *New Media & Society*, 26(10), 5923–5941. <https://doi.org/10.1177/14614448221142007>
- Shevlin, H. (2024). All too human? Identifying and mitigating ethical risks of social AI. *Law, Ethics and Technology*, 1(2), Article 0003. <https://doi.org/10.55092/let20240003>
- Skjuve, M., Følstad, A., Fostervold, K. I., & Brandtzaeg, P. B. (2021). My chatbot companion: A study of human-chatbot relationships. *International Journal of Human-Computer Studies*, 149, Article 102601. <https://doi.org/10.1016/j.ijhcs.2021.102601>
- Skjuve, M., Følstad, A., Fostervold, K. I., & Brandtzaeg, P. B. (2022). A longitudinal study of human-chatbot relationships. *International Journal of Human-Computer Studies*, 168, Article 102903. <https://doi.org/10.1016/j.ijhcs.2022.102903>
- Ta, V., Griffith, C., Boatfield, C., Wang, X., Civitello, M., Bader, H., DeCero, E., & Loggarakis, A. (2020). User experiences of social support from companion chatbots in everyday contexts: Thematic analysis. *Journal of Medical Internet Research*, 22(3), Article e16235. <https://doi.org/10.2196/16235>
- Xie, T., Pentina, I., & Hancock, T. (2023). Friend, mentor, lover: Does chatbot engagement lead to psychological dependence? *Journal of Service Management*, 34(4), 806–828. <https://doi.org/10.1108/JOSM-02-2022-0072>
- Zhang, Y., Zhao, D., Hancock, J. T., Kraut, R., & Yang, D. (2026). Interaction with AI companions and psychological well-being. *Nature Human Behaviour*. Advance online publication. <https://doi.org/10.1038/s41562-026-02516-2>