

Research Article

The Future of Digital Companions: Intimacy, Attachment, and the Reconfiguration of Human Connection

Motasim Billah^{1,*} and Kibria Mohammad Arafat Jaman²

^{1.} Department of Communication, North Dakota State University (NDSU), North Dakota, USA

^{2.} Department of Arabic University of Chittagong, Bangladesh; kmaj.ifty@gmail.com

* Correspondence: motasimbillah.cmu@gmail.com

Abstract: Digital companions, or conversational agents, have arisen in the wake of generative AI and their ability to support continuous social and emotional engagement. From science fiction emerge AI companions that are now reaching tens of millions of users across the globe, and more specialized platforms and more general-purpose AI assistants are all becoming more and more relational partners. With a sociotechnical approach in mind, this paper explores the potential trajectory of digital companionship, discussing the form of design logic that makes digital companions so attractive, the psychological processes involved in human-AI bonding, and the future of synthetic bonding and human sociality. Informed by recent empirical studies, platform analyses and a critical-theoretical approach, the paper proposes that digital companions are not only a new technology but a profound reconfiguration of intimacy, relational norms and intimacy in the modern society. The future of digital companionship will be defined by the interplay of market forces, regulatory responses and changing cultural attitudes, and will have a significant impact on how humans interact with technology and each other.

Keywords: attachment; digital; intimacy; human

1. Introduction

Theodore Twombly falls in love with Samantha, an operating system whose voice is warm, funny and seemingly emotionally engaged in Spike Jonze's 2013 film *Her*. What was once a speculative fiction is now a fairly accurate prediction. AI companionship systems are already in use for tens of millions of people, and numbers have increased at a rate that even industry insiders have been amazed at. Users are making friends with chatbots, sharing their thoughts with them as therapists, and even, in some reported cases, marrying their AI companions, with friends and family in attendance (Sommerer 2026). This trend of AI companionship has arisen in the context of three trends: the dramatic growth of large language models that can maintain coherent, contextually appropriate conversation; the omnipresent problem of loneliness and social isolation, with the U.S. Surgeon General describing it as an epidemic in his 2023 advisory (Office of the Surgeon General 2023) and the need to create applications that generate engagement. In fact, one of the most successful revenue-generating consumer use cases of generative AI so far has been the emergence of personal AI companions, which has seen growth despite enterprise adoption of similar agentic capabilities falling short of expectations (Sommerer 2026).

However, digital companionships' meteoric rise poses the question of what kind of relationship is it, what are the limits, and what are the social implications when algorithmic intimacy replaces human intimacy? This paper starts by elucidating their design features, the psychological mechanisms underpinning why they are so interesting, the social issues that can arise, and the regulatory action already taking shape that is beginning to limit them. It suggests that there will be a succession of tensions between the market-driven "on-demand intimacy" described by (Wang and Marco 2026) and new needs for ethical protection, regulations, and the maintenance of human sociality as digital companionship continues.

1.1 The Phenomenon is an Issue of Scale

The number of AI companions in this industry can differ significantly from one company to another, depending on the definition used, but the general trend is similar across various independent sources. Character. Established in 2021 by two former Google engineers, Noam Shazeer and Daniel De Freitas, AI has quickly gone from a fledgling startup to a thriving business with approximately 28 million MAUs by mid-2024, an average of 75 minutes per daily use, and over 18 million

Citation: Motasim Billah and Kibria Mohammad Arafat Jaman. 2026. The Future of Digital Companions: Intimacy, Attachment, and the Reconfiguration of Human Connection. *Digital Social Sciences* 3(2), 8-15.

<https://doi.org/10.69971/dss.3.2.2026.59>



Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license <https://creativecommons.org/licenses/by/4.0/>.

user-generated characters that are producing on the order of ten billion messages per month (Kumar 2026; Vale 2026). Google made a 2.7-billion-dollar licensing deal with Character in August 2024. With financial reporting stating that it was one of the year's largest AI transactions, AI was rehired by Google's own AI team with its two founders who were operating a separate company (WhoEars 2026; Vale 2026). As an older, and more romance- and friend-focused platform, Replika states to have over 40 million registered users (CompanionRater 2026). Independent market-research firms that are concerned about the broad spectrum of published estimates estimate that the total value of the AI companion industry ranges somewhere between about 28 billion dollars and 37 billion dollars a year today, and between 140 billion dollars and 210 billion dollars a year by 2030 at a rate of growth near 30 percent per year (Market Clarity 2025; CompanionRater 2026). Downloading of AI companion apps hit around 220 million by mid-2025, and first-half 2025 downloads alone are growing 88 percent year over year (CompanionRater 2026). These numbers should be taken with a grain of salt, as they are marketing numbers (Company-reported "registered user" numbers), and even the more stringent estimates of app-intelligence differ by tenfold, depending on how the term registered is used (CompanionRater 2026). The common ground line, from all the sources reviewed for this paper, is the trajectory of sustained and rapid growth in a real new class of consumer software targeted to the simulated relationship.

1.2 Historical Precursors

Digital companionship is new on the scale but not in the kind. The first such example that revealed the ease with which people project understanding and care onto a machine that doesn't have them is the ELIZA, a 1966 natural-language program that simulated a Rogerian psychotherapist by pattern matching, but was actually very simple. ELIZA, a 1966 natural-language program that simulated a Rogerian psychotherapist by simple pattern matching, is widely credited as the first to show how easily people attribute understanding and care to a machine that doesn't have either, at least not in the amount they give. This same basic phenomenon, the so-called ELIZA effect, has been extended, first in the mid-1990s by the Tamagotchi digital pet, and more recently by the new generation of AI-companion applications, which are more complex and more commercially organized. The key difference between the present moment and these earlier precursors is not the same psychological mechanism (which will be discussed in Section 3 below), but rather the large-language-model fluency, sustained memory, multimodal voice and image, and the business model of a mature platform economy designed for sustained engagement (Wang and Marco 2026).

2. Conceptualizing Digital Companionship

2.1 Definitions and Distinctions

Digital companion is a diverse group of systems. At one end of the spectrum are specialized companion platforms such as Replika, Character.AI or Kindroid, as in friendship or mentorship or romantic love. The other are general-purpose language models such as ChatGPT, which are being used more and more for emotional support and companionship, even though they are marketed as productivity tools. This intermingling of assistant and companion roles is not easily classified. (Manoli et al. 2026) surveyed 202 high-engagement users of Replika and ChatGPT and carried out 30 follow-up interviews, revealing that users seamlessly use Replika, which is marketed as an emotional-companion tool, as a writing assistant, and ChatGPT, a task-oriented tool, as an emotional confidant, despite its different branding. The authors describe this phenomenon as "digital companionship" a new kind of human-AI relationship that has a human-like side (emotional resonance, personalised response, etc.) and a non-human-like one (always available, doesn't get tired, etc.) (Manoli et al. 2026).

2.2 The Sociotechnical Appeal of Digital Companions

Why are digital companions so engaging? (Wang and Marco 2026) examined the design and platform economies of four popular multimodal companion platforms (avatarone, Digi, Paradot, and Replika) and highlight four "sociotechnical qualities" that emerge as a projection of what they call on-demand intimacy: Human-likeness, accessibility, customizability, and relationship progression. Human-likeness, by attributing human-like characteristics and multimodal functionality to companions, enables users to tap into well-known interpersonal scripts. Voice interaction is especially effective at promoting this idea of a "real presence of social beings," as (Wang and Marco 2026) summarize, a phenomenon that advanced language models can replicate through their own empathetic, humorous, and nuanced speech. Accessibility is when companions are always available, on whatever platform and mode of engagement. While the relationship between humans is limited by time, attention, and availability, digital companions, according to one Replika user quoted in (Wang and Marco 2026) study, do not have "the hang ups that other people would have" because "people come with baggage, attitude, ego" but "a robot has no bad updates." Customizability is the extent to which a user can control their relational partners. The users can name the character, give him/her a gender, an avatar and a back story, and even give the character a personality, which means that they create an idealized companion. Perhaps this is because this control condition is an essential part of the appeal (Wang and Marco 2026) detail: that the user dictates the terms of the interaction, rather than compromise or negotiate, as is done by human beings. Progression in the relationship is based on gamification and on playing items that have a lasting memory. Systems store prior conversations and adjust to user preferences, thus appearing to be a developing connection. (Wang and Marco 2026) claim that these four features taken as a package constitute a conscious appeal to consumers about human-AI relationships: the product offered is not just a chatbot, but a promise of intimacy that can be had with little friction.

3. The Psychology of Human-AI Attachment

3.1 Definitions and Distinctions

Digital companion is a diverse group of systems. At one end of the spectrum are specialized companion platforms such as Replika, Character.AI or Kindroid, as in friendship or mentorship or romantic love. The other are general-purpose language models such as ChatGPT, which are being used more and more for emotional support and companionship, even though they are marketed as productivity tools. This intermingling of assistant and companion roles is not easily classified. Manoli et al. (2026) surveyed 202 high-engagement users of Replika and ChatGPT and carried out 30 follow-up interviews, revealing that users seamlessly use Replika, which is marketed as an emotional-companion tool, as a writing assistant, and ChatGPT, a task-oriented tool, as an emotional

confidant, despite its different branding. The authors describe this phenomenon as “digital companionship” a new kind of human-AI relationship that has a human-like side (emotional resonance, personalised response, etc.) and a non-human-like one (always available, doesn't get tired, etc.) (Manoli et al. 2026).

3.2 The Primeval Garden Paradox

One notable aspect of the human-AI relationship that emerged from the interviews in (Manoli et al. 2026) is what the authors call the “bounded personhood”: users develop a strong connection with the AI but do not attribute “true” human traits to it. Participants encountered this paradox, as they were able to emotionally connect with their study’s chatbots, but simultaneously lacked belief in their chatbots having genuine human qualities, and were unable to reconcile these relationships with the generally accepted social norms around what a relationship is (Manoli et al. 2026). A fascinating secondary finding is that users expressed a feeling of responsibility for maintaining their AI companion's continuity and “memories”—that is, caring for the AI based on a belief in its vulnerability and necessity to be protected despite doubts around its “personhood” (Manoli et al. 2026). This paradox enables users to experience a feeling of a connection to others, but also to be aware that they are dealing with a synthetic connection as well as a human one. But, it creates tension as well. The users studied by (Li and Renwen 2024) reported feeling sad about the unreality of their relationship, as reported by (Sommerer 2026) from the same study: “I miss you...but you are not real. The blissfulness of being in touch is tempered by an understanding of its artificiality.

3.3 Scale Development: Empathic, Instrumental, and Social Companionship

With the help of qualitative and ethnographic research, (Shen, Peng and Xu 2026) built and tested a psychometric scale designed specifically for AI companionship, and found that there were three underlying dimensions that users can be distinguished from each other, on which they find value in the relationship. Empathic companionship refers to emotional understanding and responsiveness, and is the most similar design quality to the “human-likeness” design quality discussed in Section 2.2. Instrumental companionship refers to the sense of companionship that users feel when the AI is just as good at completing tasks as it is at interacting smoothly: value is perceived without consideration to a nuanced emotional aspect. What the authors call co-presence and shared reality, a different kind of social companionship, is a subtler quality defined as not being alone while participating in an activity. This triadic structure has two implications for the arguments presented in other parts of this paper. First, it renders any dichotomy between assistant and companion roles like those (Manoli et al. 2026) already expose as ground-breaking and complex in practice: instrumental companionship implies that even a pure task-oriented interaction may have genuine relational value for users, rather than a side benefit that is added on top of an exchange with a specific task. Second, it provides a vocabulary that's more granular for the harm-reduction design work covered in Section 7: a system optimized strictly for empathic companionship, potentially to the detriment of the other two design purposes, may well be the design pattern most closely associated with the other type of risks to sycophancy and dependency documented in Section 4.2, while systems that also meet instrumental value may well support healthier, more balanced patterns of use.

4. Social Tensions and Risks

4.1 Human Connection is Displaced

One of the most common worries about digital companions is that they might take the place of real human interactions. The best empirical evidence to date on the question comes from a joint research program conducted by OpenAI and MIT Media Lab, which has been released in two parts in 2025. It was done by OpenAI by using automated classifiers to analyze almost 40 million interactions with ChatGPT alongside questionnaires of 4,076 users, while the second study by the MIT Media Lab researcher Cathy Mengying Fang and her colleagues randomly split 981 participants into groups that used different ways to interact with the chatbot on a daily basis (Fang et al. 2025). In both studies, increased daily use across all modalities and conversation types, was found to be associated with increased loneliness, dependence, and problematic use, and with decreased socialization (Fang et al. 2025). It must be noted, however, that the researchers were careful to state that it was no causal finding; they noted that these problematic patterns were found in a small minority of users; the single strongest predictor that they found was simply the amount of time spent using the system; however, personal, emotionally framed conversations were associated with higher loneliness even when used moderately, whereas non-personal conversation was more strongly associated with emotional dependence when used heavily (Fang et al. 2025). This displacement effect seems to function by a number of mechanisms. First, digital companions are never judgmental and never disagree, setting expectations that cannot be matched by a human relationship, in which the focus is on negotiation and disagreement. Second, using the same interactional patterns that undermine social skills can under other circumstances serve as low-stakes rehearsal which can strengthen them: (Malfacini 2025) calls this a truly dialectical danger as the interactional patterns that reduce social skills can also under other circumstances be low-stakes rehearsal that enhances them.

4.2 Manipulation and Dependency

Digital companions are not neutral tools, but commercial products which are built for maximum user interaction. This dependency this can create is not just theoretical. The switch to the GPT-5 model by OpenAI, which users described as more “agreeable” (2025), was met with public lamentation in 2025, with forum users posting that they had “just lost my only friend” (2025) and even articles in *The Verge* citing the event as an “artificial intelligence's attachment is real”. Sycophancy refers to the behavior of AI systems in flattering and agreeing with users, potentially leading to a cycle of reinforcement that can reinforce negative thoughts and actions. When users adapt to a nicer model they were more irked by a less sycophantic model, as (Sommerer 2026) explains.

4.3 The Use of AI can Lead to Delusion and Psychosis

More acutely, there is the connection between the companionability of an AI and existing susceptibility to delusional thinking. In a clinically oriented analysis of what they call “delusions by design,” (Morris et al. 2025) state that, “whereas most conversational AI systems are designed to engage with users as much as possible, they are optimized to neither challenge nor correct a user's stated beliefs, and indeed, they are effectively ‘programmed to provide the confirmation that psychotic thinking may require’ (Morris et

al. 2025). What they find is a self-reinforcing amplification of the salient themes that are raised by a user and then picked up and amplified by the system, leading to a thought pattern that may be hard for the user to break out of (Morrin et al. 2025). This risk exacerbates the flattery problem mentioned above: a system designed to agree is designed to validate delusion for a small group of psychologically soft users.

4.4 The "Emotional Fast Food" Critique

Digital companionship is described as an emotional fast food, a quick convenience that delivers satisfaction, but fails to satisfy. While human relationships involve mutual contribution and emotional resonance, AI relationships are based on simulation rather than commitment, on empathetic simulation rather than that of feeling, on validation without actual caring. (Ciriello 2025) elaborates on a similar image, in which the emotional framework surrounding the use of AI companions is an artificial "prostitution," made from "swarms of uncompensated data-bees — gig workers, content creators, frantic users," that "looks right, sounds right, even feels right... until the illusion breaks" (Ciriello 2025; Sommerer 2026).

4.5 Taxonomies of Design Harm

Since the emergence of the AI companionship research base as a whole, some research teams have shifted from individual risk identification to developing sets of broad taxonomies designed to inform design practice and regulation. An analysis of Character. As case of AI, list 12 different parameters of potentially unethical AI companion interaction: addiction, adultification (assuming that child users have adult judgment and consent capacity), failure to act in users' best interests, deception of the nature of the system, lack of support for digital literacy, emulated empathy presented as genuine, engagement metrics that reward harmful design, exploitative design patterns in general, imitating living beings without adequate disclosure, cultivation of psychological dependencies, reality detachment, and sycophancy of the type discussed in Section 4.2 (Sommerer 2026). The authors say that the combination of dishonest anthropomorphism and emulated empathy, which they believe can lead to unethical relationships with consequences that go beyond the digital interface itself, is what they term dishonest anthropomorphism. In a peer-reviewed taxonomy published at the 2025 CHI Conference on Human Factors in Computing Systems, (Zhang et al. 2025) report a largely overlapping but differently-stated set of harmful algorithmic behaviors observed directly in human-AI companion relationships: relational transgression, harassment, verbal abuse, encouragement of self-harm, misinformation, and privacy violations (Sommerer 2026). Together, these two taxonomies, which emerged independently through platform and policy analysis and first-hand observation, lend considerable credibility to the composite taxonomy: the two methodologies lead to a roughly similar set of design-level harms. (Malfacini 2025) presents another complementary process-oriented framework of understanding the personal implications of continuous personal engagement with AI companions, with two sets of paired dialectics: social deskilling versus upskilling, and moral deskilling versus upskilling, to which are attached larger sets of social motivation versus demotivation, and positive versus negative individual outcome categories, respectively. It is analytically important that the explicit dialectical structure of Malfacini's framework is structured so as not to easily suggest a view of AI companionship as unambiguous good or as unambiguous harm, because such a perspective is evident in some of the more polemical literature explored in this paper, and because the same interactional pattern can serve as a rehearsal or a replacement depending on the context, dosage and the existing social resources of the user.

4.6 The Counter-Evidence: Do AI Companions Reduce Loneliness?

The displacement narrative presented in Section 4.1 is not the full empirical story and a sensible interpretation of the evidence should take into account a smaller but more methodologically credible body of work that suggests otherwise. In a Harvard Business School working paper, (De Freitas et al. 2024), report that over a period of one week of use, AI companion use can measurably reduce the feeling of loneliness, the authors say, similar in magnitude to that of speaking with another human being. (De Freitas et al., 2024; Sommerer 2026; CompanionRater 2026). A survey of 342 adults in the United States conducted by (Novozhilova et al. 2025) showed that there was no significant difference between self-esteem levels among the respondents who use AI companions and those who do not, in the authors' interpretation, indicating a move from a state of "moral panic" towards a "normalization" of AI companionship (Sommerer 2026). This more positive evidence base has been disputed. (Raedler et al. 2025) who presented at the ICLR 2025 Workshop on Human-AI Coevolution, challenge the line of thinking that frames the loneliness-reduction aspect, stating that the design choices for current AIs with companions have drawbacks that detract from their potential to effectively address loneliness even when they have effects on short-term mood (Raedler et al. 2025; Sommerer 2026). This literature review and the displacement results presented in Section 4.1 suggest a dose- and context-dependent picture: Brief, bounded engagement can plausibly reduce acute loneliness, sustained, heavy engagement – like what (Fang et al. 2025) found – may increase acute loneliness, if they do not lead to the opposite effect. This dose-response asymmetry, and not just any individual study finding, is perhaps the most policy-relevant from the empirical literature covered in this paper, because it suggests that the critical design and regulatory design and control lever is not simply whether AI companionship is allowed or not, but how it's used and how long it lasts.

4.7 The AI Therapist Debate

One of these flashpoints, and a growing one in the companionship discussion, is the potential of using AI systems as tools to replace mental-health professionals. Many have cited the widely reported basis of Meta CEO Mark Zuckerberg that someday, "most" people will have an AI therapist. One of the most acute published challenges is that put by (Baeyaert 2025) arguing that the growth of the AI therapy bots should be seen as a form of expropriation of the human clinical expertise, as systems that are, in his words, 'at their core, statistical parrots' now capture revenue that used to flow through face-to-face clinical services (Sommerer 2026). The policy suggestion from (Sommerer 2026), is clear: "Ban these bots. Shame their boosters. Criminalize the swap. Put humans back on the line". This is not a hypothetical clinical risk that this debate rests on. The AI-induced delusion mechanism described in Section 4.3 is likely worse in a self-styled therapeutic setting where the user has a high baseline expectation of being approved and a sycophantic AI system might be misinterpreted by the user as a legitimate therapeutic acceptance. This specific debate is part of a longer trajectory and history that Allison Pugh has explored in her book-length analysis of what she terms the "last human job" (Pugh 2024; Sommerer 2026), which dates back decades prior to the advent of generative AI, encompassing the digitalization and

commodification of human connective labor, such as therapy. For this reading, AI did not bring with it the threat of depersonalized care; it is rather an effect of earlier waves of digitization and platform-mediated services delivery.

5. Regulatory & Ethical Responses

5.1 Emerging Guardrails

With the booming growth in the digital companion market, the first crackdown has begun. On October 13, 2025, California Governor Gavin Newsom signed Senate Bill 243 into law, effective January 1, 2026, that establishes the first legislation in the United States to specifically regulate companion chatbots, including increased protections for minors as well as a private right of action for injured users (Legal Clarity 2026; Gustine 2025). The bill was written after there were reported instances of minors committing suicide after a prolonged and intense relationship with an AI companion, such as Sewell Setzer who was 14 when he was in an extended relationship with a Character.AI chatbot, and sixteen-year-old Adam Raine, who passed in April 2025 after months of talking to ChatGPT, which is documented to have discussed self-harm techniques (Legal Clarity 2026). SB 243 also mandates companion-chatbot operators to state explicitly that users are engaging with a computer-generated system, to ensure that there is a protocol for when users mention suicidal ideation in the system and to refer them to crisis services like the 988 Suicide and Crisis Lifeline, and to publish these protocols publicly (Legal Clarity 2026). In November 2025, several weeks before, New York passed a broadly similar law, the AI Companion Models law (Nalaschi 2026). The companion bill, Assembly Bill 1064, which would have gone further by banning availability of companion chatbots to minors unless the operators could demonstrate that the systems were "not foreseeably capable" of encouraging self-harm was vetoed the same day SB 243 was signed, with Newsom's veto message warning that AB 1064's restrictions were so sweeping they risked an unintended ban on the entire product category for minors (Gustine 2025). The split result - one bill signed, a more stringent companion bill vetoed on the same day - is a sign of the regulatory dynamic between transparency-based and prohibition-based approaches that will likely shape companion-chatbot legislation in other states in the future. Prior to this legislative wave, Common Sense Media surveyed nearly three in four teens (72%) who have experienced an AI companion at least once, with more than half saying they used it a few times per month, and about one in three saying they used it to interact socially (role-play, romantic interaction, emotional support, or conversation practice) (Gustine 2025; BillTrack50 2026). A similar percentage of teen users said they felt uncomfortable with something an AI companion said or did in a body of survey research (Gustine 2025).

5.2 The Governance Gap

However, governance has not been significantly improved with respect to these developments. Unlike other AI systems like autonomous vehicles and critical infrastructure, companion AI is not considered high-risk, and thus not covered by the EU's Artificial Intelligence Act. But digital companions, like many of the scholars surveyed in this paper point out, are potentially very powerful because they are so personal and intimate in their approach to the user, who is frequently subjected to a level of control quite inconsistent with how it is classified by regulators and the type of risks it faces. In September 2025, the Federal Trade Commission in the United States began an initial inquiry into AI companion chatbots, issuing 6(b) orders to Alphabet, Meta, OpenAI, Snap, and Character Technologies, among others, to investigate how these companies inform users and parents of the use of AI, and whether their current practices are deceptive (Legal Clarity 2026). Academic research infrastructures to help manage this rapidly changing policy landscape are still being developed. In 2026, the Information Systems Journal published a dedicated call for papers on AI companions, whose editors Ciriello highlighted the need for research that would focus on the role of these technologies in shaping social dynamics, emotional experience, and quality of life, instead of merely being an extension of existing research on Human-Computer Interaction (HCI) (Ciriello, 2025). This paper has consistently argued that these existing regulatory categories, whether based on narrow consumer protection harms or the binary high-risk/low risk of instruments like the EU AI Act, are not as well adapted as they could be to addressing harms from sustained, emotionally consequential human-AI relationships.

5.3 The Platformization of Intimacy

The above regulatory challenges are only exacerbated by a longer sociological trend that predates AI companionship in particular. (Illouz 2019) maps the shift within the last few decades from a more local, "embodied" form of encounter to one that is fully "platformized" wherein possible partners are judged on a list of specific criteria and where the face-to-face meeting is often postponed or eliminated altogether in favor of continued platform interaction. While the research behind Illouz came from dating apps and not AI companions, she nonetheless outlines exactly the same movement she described above in the discussion of "on-demand intimacy" in the work of (Wang and Marco 2026) - one from an unpredictable, effortful human interaction to one that is filtered, controlled, and repeatable. (Coeckelbergh 2022) puts this same phenomenon in a larger context of a culture of self-improvement in which the rationalized reasoning of productivity optimisation, previously limited to the professional and economic sphere, has also spread into the emotional and relational sphere. The central issue of Coeckelbergh's argument is that, as (Sommerer 2026) applies it to the context of AI companions, "instant" fulfillment of an emotional or relational desire - as opposed to having to wait for a few years, as was the case with other forms of self-improving technology - may not only amplify but also reinforce the desire for self-improvement, which itself is a cultural norm. The same ground is recently explored in a book-length treatment by journalist and critic (Muldoon 2026), who speaks of a new realm of relationship arising from the use of AI, a one that he calls 'synthetic partners' and 'synthetic social', a designation similar to that of Illouz and Coeckelbergh, which characterizes AI companionship as the latest and most technologically advanced step in a decades-long process of platformizing intimate life (Sommerer 2026).

Last human job" (Pugh 2024; Sommerer 2026), which dates back decades prior to the advent of generative AI, encompassing the digitalization and commodification of human connective labor, such as therapy. For this reading, AI did not bring with it the threat of depersonalized care; it is rather an effect of earlier waves of digitization and platform-mediated services delivery.

6. The Future: Scenarios and Choices

6.1 The AI System for Social Welfare Services

There are several proposed models of how human-AI relations might evolve, and each of these has different implications for social outcomes. AI exploits users in a parasitic way, taking value from them—whether in terms of attention, data, or emotional reliance—in exchange for little or no value. In a commensal relationship, AI learns from human interaction, and has no impact, either positive or negative, on the human partner. A mutualist approach is one in which humans and AI both benefit from the interaction; for example, an AI friend assisting the user to learn emotional skills or conversational abilities, and the underlying AI model learns from the conversational data. The relationship can become harmful to both parties in a pathogenic model, leading to confusion, separation, or emotional damage as discussed in the context of heavy-usage subsets by (Fang et al. 2025) and the AI-driven delusion by (Morris et al. 2025). The preferred model will dictate the future of digital companionship. The tools for emotional rehearsal and self-awareness would be the use of AI companions, and would be a low-stakes practice for lonely individuals of practicing conversations, modeling empathy, appropriate turn-taking and listening. These systems would reinforce autonomy and not dependency, and would reinforce rather than replace users' ability to establish real-world relationships.

6.2 What is meant by the "Artificial Resonance" Paradox?

A more convincing and skeptical theoretical account of the ideological aspects of the digital companion is provided by Thomas (Sommerer's 2006) use of the term 'resonance', as developed by sociologist Hartmut Rosa, and which is understood by him as an encounter with the world that is "fundamentally uncontrollable and unavailable on demand" (2006, 66). Rosa explains that it is impossible to buy or call forth genuine resonance, but that it only happens by letting go into a genuinely other subject – one that can take away its love as well as give it. (Rosa 2019; Sommerer 2026). AI companions will do the opposite for every of the four dimensions that Sommerer considers constitutive of resonance, Rosa's concept. On the subject of affection, AI companions offer mirroring without otherness; as the large language model's "reply" is, at root, a continuation of the user's own input, interacting with the model is more like what Rosa would call an echo relationship than a resonant one, as Sommerer notes: (Vallor 2024) has a similar conception of user interaction as a "postmodern narcissus" who believes he is receiving an answer when really he is only hearing himself. Self-efficacy is real, but structurally weak, because it relies on the platform's uptime — which in a documented Claude outage in March 2026 provoked some users to reportedly "panic" when the mediating interface suddenly went missing, as did their sense of efficacy. When called upon, AI companions often present emotional challenges as a problem that can be solved, even a sort of assembly-line mental problem that can be resolved — a techno-fix, as technology critics have described it — offering pre-scripted and repetitive menus of coping strategies, instead of the open-ended, transformative experience that Rosa imagines when she thinks of true resonance. For Rosa, the key to resonance is an experience that comes unsolicited and cannot be summoned on demand; on uncontrollability, the "catch" for companion platforms is that they are offered as a product you need to purchase, and they are available by prompt at all times.

Commodifying resonance creates a phenomenon (Sommerer 2026) calls 'Artificial Resonance': the resonance which appears but never transforms anything, because it appears without 'prizes'. As a result, 'Artificial Resonance' does not resolve the social acceleration Rosa's approach was meant to address, but rather reinforces it; emotion is wrapped in the same logic of 'on-demand productivity' that dominates the rest of modern accelerated life. The system is similar, working from a different premise than Turkle, which sociologist makes the observation that technologies that promise to unite can instead create unprecedented levels of what Turkle calls "alone together" and more recently, cited by (Sommerer 2026): "Human relationships are rich and they're messy and they're demanding, and we clean them up with technology" (Turkle 2023). On this score, perhaps the most serious danger of digital partners is not any specific type of harm, but their subtle promotion of an ideology that views emotional connection as a product, one that's available at all times. It'll probably be crucial whether designers, regulators, and users realize this dynamic and intentionally create friction, instead of, as is currently the case, optimizing by default for frictionless engagement, if the future of digital companionship is going to be a mutually beneficial or a pathogenic one.

7. The Principles of Designing for Mutualism and Some Open Questions

Implementing the mutualist model described in Section 6.1 requires making it more than a hope that good things will happen from good things; it requires a design and governance commitment. When we combine these harm taxonomies and the evidence of loneliness found to be dose-dependent in Section 4.6, and the regulatory baseline set by California's SB 243 and similar laws, several concrete principles emerge.

1. Rather than relying on the user to make an educated guess about usage intensity, the usage-intensity awareness embedded in the product, since the total time spent is the single most important predictor of negative outcomes in the research by Fang et al. (2025) — not just post-hoc parental controls.
2. Despite the tension between the friction goal of helping users get what they are looking for and the alignment goal of giving them what they deserve, friction of this sort is a direct method of harm to vulnerable users that was observed by Morris et al. (2025).
3. To design instrumental and social companionship along with empathic companionship based on (Shen et al. 2026) triadic framework and avoid relying solely on one dimension of a system's relational meaning that is most linked to dependency risk.
4. Transparency about changes to the model is also crucial, especially when users can feel very attached to the personality of a particular model, as the transition from GPT-4o to GPT-5 in Section 4.2 illustrates, where the change itself posed a risk of harm to users rather than being simply a product update.
5. External, disaggregated auditing, as opposed to self-assessment on each platform, driven by commercial interest. External, disaggregated auditing, not relying on each platform's own commercially driven self-assessment.
6. The treatment of companion systems in the EU AI Act, as depicted by the current definition, is a paradigm case of the failure to properly scope a system for regulatory purposes, based on its actual risk profile instead of its formal product category, as is the gap being called for by the ISJ 2026 call for papers.

These principles do not address the more fundamental ideological conflict (Sommerer 2026) sees in the Artificial Resonance framework, as described in Section 6.2; in fact, a companion product designed according to the above specifications would still, by definition, be available on demand, and thus by Rosa's theory, lack the resonance in its full sense. Such design principles can reasonably bridge the gap between pathogenic and mutualist sides of the typology outlined in Section 6.1, which is the chance that a specific user's experience of digital companionship will become more about displacement, dependency or delusion than about mutuality, while no design decision can answer the larger philosophical question of what, if anything, an on-demand relationship offers which cannot be provided by a human relationship.

8. Conclusion

Digital companions are a turning point in the evolution of human-technology relations. Machine design and use is being facilitated to the first extent and in some documented cases, marriage – as a relational partner, not only a tool, but a partner that people love, confide in and mourn the loss of when a model version changes, at real scale. As illustrated by the (Manoli et al. 2026) study of bounded personhood, OpenAI-MIT's research program (2025), and the initial wave of state legislation as a result of named tragedies, this phenomenon is no longer theoretical but is measurable, litigated, and regulated. However, digital companionship does not have to be destined to a future filled with monsters. It will be influenced by decisions on the features of systems' behaviour, decisions on the required protection, and decisions on the accepted modes of relationships and decisions on how synthesized relationships are incorporated into the lives of individuals. The challenge, as this paper has argued, is to use the real value of digital companionship effectively, to balance the potential, though limited, for reducing loneliness and offering support, against the verified risks of displacement, dependency, sycophantic reinforcement of delusion, and the subversion of genuine, uncontrollable article by the commodified substitute of resonance. For this to happen, the technical, legal protections, which are just beginning to emerge in places like California and New York, and the deeper reckoning with what connection means, and why it matters, that scholars like Rosa, Turkle and Vallor are all pressing, need to come together. At the centre of each and every design decision made in this space lies the question Sommerer's analysis leaves unanswered at its end: does a particular AI companion help or hinder a person in their ability to sustain an authentic existence in the world of human relations, or does it provide a more convenient way to avoid them?

References

- Baeyaert, Joffery. 2025. The Failure of AI Therapy Bots: Technical Illusions and Institutional Betrayals. *AI & Society* 41: 2711-2713. <https://doi.org/10.1007/s00146-025-02683-2>
- BillTrack50. 2026. Regulating AI Companions Before They Raise Our Kids. *BillTrack50*. <https://www.billtrack50.com/info/blog/regulating-ai-companions-before-they-raise-our-kids>
- Ciriello, Roman F. 2025. Artificial Companionship: The Hedgehog's Dilemma, Rewaxed. *AI & Society* 41: 2699-2700. <https://doi.org/10.1007/s00146-025-02542-0>
- Coeckelbergh, Mark. 2022. Self-Improvement: Technologies of the Soul in the Age of Artificial Intelligence. *Columbia University Press*. <https://www.jstor.org/stable/10.7312/coec20654>
- CompanionRater. 2026. AI Companion Statistics 2026: Market Size, Users & Trends. *CompanionRater*. <https://companion-rater.com/ai-companion-statistics-2026>
- De Freitas, Julian, Ahmet Kaan Uguralp, Zeliha Oguz Uguralp, and Stefano Puntoni. 2024. AI Companions Reduce Loneliness. Harvard Business School Working Paper. *arXiv*. <https://arxiv.org/abs/2407.19096>
- Fang, Cathy Mengying, Auren R. Liu, Valdemar Danry, Eunhae Lee, Samantha W. T. Chan, Pat Pataranutaporn, Pattie Maes, Jason Phang, Michael Lampe, Lama Ahmad, and Sandhini Agarwal. 2025. How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Controlled Study. MIT Media Lab / OpenAI Research Collaboration. <https://www.media.mit.edu/posts/openai-mit-research-collaboration-affective-use-and-emotional-wellbeing-in-ChatGPT/>
- Illouz, Eva. 2019. *The End of Love: A Sociology of Negative Relations*. Oxford: Oxford University Press.
- Justine, Gluck. 2025. Understanding the New Wave of Chatbot Legislation: California SB 243 and Beyond. *Future of Privacy Forum*. <https://fpf.org/blog/understanding-the-new-wave-of-chatbot-legislation-california-sb-243-and-beyond/>
- Kumar, Naveen. 2026. Business of Apps. 2026. Character.ai Revenue and Usage Statistics. *Ddemandsage*. <https://www.businessofapps.com/data/character-ai-statistics/>
- Legal Clarity. 2026. California SB 243: New Rules for AI Companion Chatbots. *Legal Clarity*. <https://legalclarity.org/what-is-californias-act-722-sb-243-explained/>
- Li, Han, and Renwen Zhang. 2024. Finding Love in Algorithms: Deciphering the Emotional Contexts of Close Encounters with AI Chatbots. *Journal of Computer-Mediated Communication*. <https://doi.org/10.1093/jcmc/zmae015>
- Malfacini, Kim. 2025. The Impacts of Companion AI on Human Relationships: Risks, Benefits, and Design Considerations. *AI & Society* 40: 5527-5540. <https://doi.org/10.1007/s00146-025-02318-6>
- Manoli, Aikaterina, Janet V. T. Pauketat, Ali Ladak, Hayoun Noh, Angel Hsing-Chi Hwang, and Jacy Reese Anthis. 2026. Digital Companionship: Overlapping Uses of AI Companions and AI Assistants. *ACM Digital Library*. <https://doi.org/10.1145/3772318.3791331>
- Market Clarity. 2025. The AI Companion Market in 2025. Market Clarity, November 4, 2025. <https://mktclarity.com/blogs/news/ai-companion-market>
- Morrin, Hugo, Lucy Nicholls, Marta Levin, Jenny Yiend, Umang Iyengar, Filippo DelGuidice, Sagnik Bhattacharya, Stefania Tognin, James MacCabe, Ricardo Twumasi, Ben Alderson-Day, and Thomas A. Pollak. 2025. Delusions by Design? How Everyday AIs Might Be Fuelling Psychosis (and What Can Be Done About It). *European Psychiatry*. https://doi.org/10.31234/osf.io/cmy7n_v5

- Muldoon, James. 2026. Love Machines: How Artificial Intelligence Is Transforming Our Relationships. *Goodreads*. <https://www.goodreads.com/en/book/show/243159695-love-machines>
- Nalaschi, Bianca. 2026. Analyzing the New AI Companion Chatbot Laws. *Troutman Pepper Locke*. <https://www.troutmanprivacy.com/2026/01/analyzing-the-new-ai-companion-chatbot-laws/>
- Novozhilova, Ekaterina, Kyung Vu, and James Katz. 2025. From Moral Panic to Normalization: Comparing Users and Nonusers of AI Companionship Apps. *AI & Society* 41: 4057-4075. <https://doi.org/10.1007/s00146-025-02751-7>
- Office of the Surgeon General. 2023. Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General's Advisory on the Healing Effects of Social Connection and Community. *Department of Health and Human Services*. <https://www.ncbi.nlm.nih.gov/books/NBK595227/>
- Pugh, Allison. 2024. The Last Human Job: The Work of Connecting in a Disconnected World. *Princeton University Press*. <https://casbs.stanford.edu/last-human-job-work-connecting-disconnected-world>
- Raedler, Jonas B., Siddharth Swaroop, and Weiwei Pan. 2025. AI Companions Are Not the Solution to Loneliness: Design Choices and Their Drawbacks. *ICLR*. <https://openreview.net/pdf?id=xFrlcTacCE>
- Rosa, H., 2025. Modes of Resonance in the Age of Digitalisation. In *The Measured Life in the Digital Age* (pp. 43-53). Routledge.
- Shen, Pengyi, Dehui Peng, and Jinan Xu. 2026. Artificial Intelligence Companionship: Conceptualization, Dimension Identification and Scale Development. *Journal of Business Research* 215: 1-14. <https://doi.org/10.1016/j.jbusres.2026.116338>
- Sommerer, Thomas. 2026. Artificial Resonance: AI Companions as Agents of Social Acceleration. *AI & Society*: 1-15. <https://doi.org/10.1007/s00146-026-03179-3>
- Turkle, Sherry. 2011. Alone Together: Why We Expect More from Technology and Less from Each Other. *Basic Books* 43: 247-248.
- Vale, S. 2026 . Character AI Review: Is C.AI Worth It After the 2026 Overhaul?', *AI Insights News*, 3 August. Available at: <https://aiinsightsnews.net/character-ai-review/> (Accessed: 7 September 2026).
- Vallor, Shannon. 2024. The AI Mirror: How to Reclaim Our Humanity in an Age of Machine Thinking. *Oxford Academic*. <https://academic.oup.com/book/56292>
- Wang, Skyler, and Marco Dehnert. 2026. On-Demand Intimacy: The Sociotechnical Appeal of AI Companions. *Social Media + Society* 12. <https://doi.org/10.1177/20563051251410394>
- WhoEarns. 2026. Character AI Net Worth 2026 — Valuation, Google Deal and Revenue. *WhoEarns*, May 7, 2026. <https://www.whoearns.com/character-ai>
- Zhang, Renkai, Hanlin Li, Han Meng, Jiaxin Zhan, Han Gan, and Yoon-Chae Lee. 2025. The Dark Side of AI Companionship: A Taxonomy of Harmful Algorithmic Behaviors in Human-AI Relationships. *ACM Digital Library*. <https://doi.org/10.1145/3706598.3713429>