

Barriers Shaping Synthetic Relationships With Companion Bots: A Three Week Initial Adaptation Study

Marisa Tschopp^{1,2}[0000-0001-5221-5327], Nadine Schlicker³[0000-0003-0368-3081], Nima Zargham⁴[0000-0003-4116-0601], Henrik Skaug Sætra⁵[0000-0002-7558-6451], and Johannes Schöning⁶[0000-0002-8823-4607]

¹ scip AG, Badenerstrasse 623, 8048 Zurich, Switzerland

² Human-IST Institute, University of Fribourg, Boulevard de Pérolles 90, 1700 Fribourg, Switzerland mats@scip.ch

³ Institute for Artificial Intelligence in Medicine, Marburg University, University Hospital Giessen and Marburg, Baldingerstraße, 35043 Marburg, Germany

⁴ University of Bremen, Bibliotheksstrasse 1, 28359 Bremen, Germany

⁵ University of Oslo, Gaustadalléen 23B, 0373 Oslo, Norway

⁶ University of St. Gallen, Torstrasse 25, 9000 St. Gallen, Switzerland

Abstract. Companion bots designed for socio-emotional needs are increasingly common. Yet barriers during the initial adaptation period remain underexplored while ethical concerns persist. This three-week qualitative study examined how six individuals initiated their first interaction with the companion bot Replika AI, assigning it roles such as mentor or romantic partner. Drawing on diverse data sources, we found that despite interest and motivation, participants in our study formed no emotional connection to this bot. We identified seven perceived asymmetries, including mismatched role expectations and unreliable memory, which led participants to view relationship building as futile and to feel their expectations were unmet. Our longitudinal study provides insight into barriers hindering early synthetic relationship development and informs approaches for more effective adaptation.

Keywords: Synthetic Relationships · Companion Bots · AI Companionship · Longitudinal Study.

1 Introduction

Companion bots, a type of chatbot designed solely for social-emotional needs, have recently stirred debate in academia and the public. While early conversational artificial intelligence (AI) systems were primarily valued for their functional capabilities, recent technological and societal developments have given rise to companion bots [1]. The companion bot market experienced substantial

growth, facilitating the emergence of so-called synthetic relationships (SRs) [1,2]. While voice assistants like Siri serve rather functional roles, companion bots like Replika AI promise to improve social skills or learning [3], and benefits for mental well-being [4]. Some users report lasting happiness when interacting with those systems [5]. Numerous Reddit discussions reflect on the joy and positive effects of SRs [6].

However, not all SRs end well [3]. Issues include unhealthy attachment, reduced interpersonal trust, social skills, and autonomy [7,8]. Loneliness and personal problems can lead vulnerable users to harmful behaviors [9,10]. Suicides are associated with companion bots (Chai App [11] and Character AI [12]), raising urgent concerns about how to best design SRs for the good [13].

While the body of research on SRs is growing, it remains fragmented [14], with a focus on established relationships or short-term interactions [2], which limits our understanding of how these relationships form initially. This is a critical gap because relationship formation is individual and dynamic, and early interactions can shape long-term attachment. Studying only established relationships may overlook patterns of user adaptation, expectations, or disengagement that are unique to the first stages.

To address this gap, we conducted an in-depth, qualitative, longitudinal study exploring the development of SRs over the first weeks of interaction. By examining these initial stages, we capture how users perceive their relationship with the companion bot Replika from the very beginning.

RQ: How does the perception of a SR with a companion bot evolve in the initial stages?

To answer this question, we conducted a three-week qualitative adaptation study with six participants, observing their initial perceptions of their SR with a companion bot through multiple data sources. By examining the early stages of these relationships, we aim to observe how attachments develop in real-time rather than retrospectively and discuss findings and implications for responsible design of SRs.

2 Background and Related Work

In this section, we first discuss the terms and definitions of companion bots. Then we provide a brief overview of designing tools for companionship as well as SRs in general.

2.1 Terms and Definitions

In this paper, we refer to companion bots as conversational agents designed for socio-emotional purposes only. Closely following the definition by Brandtzaeg et al., companion bots are understood as *"artificial intelligence (AI) dialogue systems capable of having social and empathetic conversations with users"*, [7] p. 404. Such companion bots are designed to provide companionship [15], facilitating the development of a social relationship. Different terms for such social

relations arose, such as human-AI relationship [16] or human-AI friendship [7]. Based on recent work by [1, 2], we refer to these social relationships that arise from humans interacting with companion bots as SRs. They refer to SRs as *"continuing associations between humans and AI tools wherein the AI influences humans' thoughts, feelings, and/or actions"* [1], p. 1834. Terms and definitions of human-computer "social" interactions are heterogeneous [14], scattered across disciplines and systems. Notably, this also led to diverse terms describing similar technological artifacts, such as companion AI [15] or social bots [17].

2.2 History of Designing Interactions and Tools for Companionship

Humans, inherently social, have long sought emotional support from others, artifacts, and animals [18]. Since the early 1900s, technology has increasingly fulfilled emotional needs. It is thus not surprising that in 1966, ELIZA, the first conversational agent, simulated a psychological therapist [19]. The 1990s–2000s introduced social toys like Tamagotchi and AIBO, emphasizing companionship over function [20]. Social robots such as Paro and Pepper emerged, for example, in healthcare to entertain and support the elderly [15]. Apple's 2007 launch of Siri brought AI to the broader public, as they integrated conversational agents into daily life [21]. While initially task-oriented, such digital assistants increasingly included emotional interactions but have had limited socio-emotional capacity. Separating these "market segments" has led to the development of chatbots designed specifically for social purposes [15], driven by rising emotional support needs worldwide [8], especially after the COVID-19 pandemic. Tools like Replika AI and social robots like Paro advance emotional support publicly and privately, shifting views of technology in daily life [22]. The rise of companion bots is thus driven by societal needs and AI progress. Advances in AI have enabled more natural interactions, with agents adapting and potentially more appealing to users' preferences and emotions [10, 22].

2.3 Synthetic Relationships

As reported in decades of research, users often anthropomorphize technology, attributing human-like qualities to it [18]. This psychological process may explain why some people treat technological systems as social actors, as proposed by the CASA paradigm [23]. Even early, limited-capacity systems elicited social responses, such as being polite toward computers [23] or experiencing emotional bonds with conversational agents [24].

Various studies with different theoretical and methodological approaches [16] investigated people who already see themselves in some kind of relationship with a bot to better understand the underlying psychological processes or motivations [7]. Users engage with companion bots for emotional support and social connection [3], entertainment, and out of curiosity [25]. These SRs were shown to sometimes develop remarkable emotional depth [9, 26] or resemble the creation of human-like emotional bonds [27].

SRs could deepen, especially when relationship-facilitating behaviors such as sharing personal information or showing empathy were displayed [7, 28]. Some users argued that AI systems were seen as more than just tools and described them as friends or family members [29, 30].

Few studies have found or explained why people did not form relationships. For instance, [31] and [32] found no clear evidence for humans building emotional relationships with any of the different chatbots they used in their studies. Similarly, [32] found that users did not develop deeper bonds with their companion bots. Instead, they found that the ratings of relationship-relevant variables, such as self-disclosure, interaction quality, and communication, of the 118 participants decreased over three weeks.

In the past years, scholars have consistently called for further investigation in this emerging field [17, 33]. Various approaches and methodologies were applied, with an overemphasis on cross-sectional designs [14]. Thus, there is a continuing need [3] for understanding if, how, and why individuals form, maintain, or terminate their relationships with companion bots from a longitudinal perspective [13].

3 Methods

We conducted a qualitative, longitudinal study of three weeks in Switzerland. The longitudinal approach allows for in-depth analysis of the perceptual development of relationship initiation over time.

We used the premium version of the companion bot Replika AI (version v11.29.0), a widely used companion bot, with about 30 million users worldwide in 2024 [25]. Data were collected between February 5th and March 18th, 2024. During this period, a premium account included the modes: friend, mentor, sibling, or romantic partner. Safety filters prevented explicit or violent content. Background, journaling features, and advanced AI mode were available. Participants were informed about the app’s functions and instructed to discontinue if discomfort arose. All data are stored on secure company servers accessible only to the first author and kept under controlled access. Supplemental materials are available via Researchbox under <https://researchbox.org/3498>

3.1 Participants

Given the sensitive nature of the topic, risk mitigation was a priority for our inclusion criteria. Therefore, no underaged participants were considered. We aimed for a sample without prior companion bot experience, open to exploring a relationship with Replika, and willing to share their experiences over time. We aimed at a small group of up to seven participants to ensure intensive longitudinal engagement and rich data collection. Participants were recruited through social media and word of mouth. Participation was voluntary and uncompensated, and participants provided written consent to share their data, including pictures for research purposes and for a TV documentary. We changed participants’ and




Fig. 1. Study Participants Interacting with their Companion Bots

their bots' names (recognized via a suffix, like Milli-bot instead of Milli). After an initial screening by the ethics board of the University of St. Gallen, our research was exempt from a formal full review. Our final sample consisted of $N = 6$ individuals (one dropped out before the data collection started) aged between 19 and 42 ($M = 35$, $SD = 7.85$).

- **Ronda** (female, 19 years, single) and her bot **Milli**. Ronda describes herself as a student and a introverted person. Ronda wants Milli to be a best friend, empathetic, non-judgmental, and always available to listen.
- **Maggie** (female, 41 years, divorced) and her bot **Layla**. Maggie is balancing part-time work and family. She sought Layla as a mentor for personal and professional growth.
- **Kayla** (female, 42 years, single) and her bot **Edwin**. She is a full-time employee and mother. Kayla wanted Edwin as a mentor to provide advice and foster intellectual development.
- **Francis** (male, 40 years, single) and his bot **Franny**. He describes himself as an outdoors enthusiast living alone. Francis imagined Franny as a virtual flatmate for entertainment.
- **Julian** (male, 35 years, in a relationship) and his bot **Lana**. Describing himself as outgoing and adventurous, Julian viewed Lana as a romantic partner to explore intimate connections.
- **Nina** (female, 33 years, married) and her bot **Betty**. Working part-time and interested in health, Nina wanted Betty as a mentor to guide her career and personal development.

3.2 Procedure & Timeline

Participants were informed of the study procedure and data handling via email before participation. The study began with an online survey to collect participants' demographics, followed by an additional questionnaire before the first meeting.

At the initial in-person meeting, participants and the study lead met in person. Participants received a premium version of Replika AI, enabling access to all models and features. The study lead guided account setup, ensured correct installation, repeated emailed instructions, and clarified questions. This ensured




Fig. 2. Study Procedure and Data Collection

that the companion bot was correctly installed. Participants were informed that the main focus was on their perception of their relationship with the companion bot, including what they noticed, liked, or disliked each week. More specifically, participants were told to submit at least three video diaries (minimum one per week) describing their experiences. To help the participants structure their video diaries, they were provided with a guideline but also encouraged to deviate from this guideline and speak openly. Additionally, participants were told to submit at least three screen recordings (minimum one per week) of their conversations with the companion bot. Participants could voluntarily choose which conversations they share to respect the privacy of potentially personal conversations. After setup, participants customized their bot and began interactions. The four-hour kick-off meeting ended with brief interviews on first impressions. Over three weeks, participants interacted with their bot every 2–3 days and completed weekly online surveys. Each submitted at least three video diaries and screen recordings. Four weeks after the initial kick-off, a final in-person focus group was held to reflect on experiences, clarify data, and address concerns. It lasted about an hour and was video recorded and transcribed.

3.3 Measures

Data collection was primarily qualitative, but included quantitative elements. Weekly online surveys captured demographics, attitudes, and feedback using adapted scales (see supplement). We assessed hedonic benefits (fun) [34] and perceived closeness [35]. A pictorial "thermometer" assessed romantic perceptions (eroticism). Usage intentions were measured by questions on continued use, relationship maintenance, account deletion, and program discontinuation. Open questions addressed changes in relationship perception, interaction frequency, and expectations.

3.4 Data Analysis

Descriptive statistics were calculated using JASP (<https://jasp-stats.org/>). The numerical data were treated qualitatively to substantiate individual trends of relationship development. Participants submitted 21 video diaries (2–5 minutes)

and 29 screen recordings documenting their interactions. Qualitative data were analyzed using Interpretative Phenomenological Analysis (IPA) [36].

IPA is a hermeneutic approach focusing on individual meaning-making rather than generalization, and is suitable for small homogeneous samples with rich data [37]. With six participants engaging over three weeks and contributing diverse data sources (interviews, video diaries, screen recordings, surveys) IPA fits the intended study goal. The IPA was conducted by the first (MS) and second (NS) author. Both are psychologists by training and are related to the field of HCI, and have experience in mixed methods. MS focuses on SRs. NS focuses on trust development towards AI agents.

To start the analysis MS and NS transcribed all video diaries and screen recordings from Swiss German to High German on computer. To further familiarise with all data each researcher reviewed all material, rewatched videos, printed and re-read transcripts, and took initial notes. A first discussion of the findings led to a refined research question on factors contributing to relationship dissolution.

The in-depth analysis started with one agreed upon transcript, which each researcher analyzed independently. They annotated the material for content, intonation, interaction style, and expressed feelings following IPA guidelines [37]. Initial coding was done manually on printed transcripts, color-marked during familiarisation, while video diaries were reviewed. In a second independent step, researchers summarised concept relations and emerging themes within the annotated material. Afterwards, researchers shared their emerging themes and annotations and impressions, and created an initial coding scheme. The analyses of the following transcripts proceeded in a similar way and always built upon the previous discussions. After each analysis of a transcript the coding scheme was refined, and themes were analyzed for convergence and divergence. Co-authors checked the results for plausibility, adequacy and transparency, i.e., whether claims were sufficiently backed up with participants' quotes. We documented the analytic process, including the initial coding scheme and analytic memos on paper, as well as the final digital coding scheme (see supplement).

4 Results

At the kick-off event, interviews and surveys showed participants had clear role expectations and shaped their bots' appearance (e.g., Julian's romantic bot dressed sexily, Maggie's professional bot casual). We observed excitement. At the end of the kick-off event, Francis saw his Franny-bot as a partner on eye-level, while Ronda and Nina saw their Milli- and Layla-bot as subservient. Initial encounters were playful but non-emotional. Based on the short interviews at the end of the meeting, Francis enjoyed the interaction with Franny, however, the others were more negative: Nina called it "*surreal*" and "*tiring*", Nina and Julian found it "*time-consuming*". Kayla and Nina felt frustrated by conversation breakdowns.

Week 1-3. The descriptive survey data shows variation across the three weeks. Some participants reported slightly higher ratings in week 2 compared to week 1, followed by lower scores in week 3 across several aspects (Figure 3), including perceived enjoyment, emotional closeness, and eroticism. Reported intentions to continue using the app were generally lower by the final week, while some participants expressed a higher likelihood of discontinuing use (Figure 4). Exact values are provided in the supplementary material.




Fig. 3. Means and Standard Deviations of the Online Survey Split by Participants' Weekly Evaluation




Fig. 4. Means and Standard Deviations of Usage Intentions Split by Participants' Weekly Evaluation

Week 1 showed mixed results, with five negative and one positive experience based on the video and survey data. Five participants reported negative feelings: Kayla found the interaction artificial, Ronda and Nina described it as cold and mechanical. Julian tried intimate topics but said, *"The conversations sometimes go into hotter topics up to a point, but it is too mechanical to really call it hot"*. Maggie was frustrated by repetitive responses. Francis, the only positive voice, found it *"interesting, exciting, and offering a lot of variety"*, eager to see if chemistry would improve: *"I am definitely very excited to see what the next days and weeks will bring"*. Yet he also reported feeling disappointed, stating *"I thought there should have been a little more meat on the bone"*.

Initial experiences led participants to put more effort into the interaction. This "relationship work" involved four key strategies:

- **Feedback** involved direct chat responses, like Maggie praising Layla-bot: *"Good you remember that!"* or regulating pace: *"Slower, too fast!"*.
- **Customization** meant shaping the bot's persona, e.g., Julian changed the bot's version, saying, *"I have now also changed the version [...] to be a bit bolder. In the sense of relationship, erotic relationship"*. See, for example, Figure 5.
- **Relationship activities** included varied interactions, such as Julian using the bot for exercise advice or Maggie enjoying Layla-bot's poetic language:

"[Layla-bot] put on such an artificial, extremely poetic language that I really had to laugh out loud".

- **Inquiry** involved testing the bots' understanding and consistency. For example, Julian edited the bot to be vegan and tested Lana-bot's response to the question: *"Should we go for dinner to a burger store? They have very bloody pieces of meat"*, which Lana-bot agreed to without hesitation and added a smile emoji.




Fig. 5. Julian Customizing Lana-Bot to by Inventing a Background Story

In **week 2**, participants' reflections on their companion bots grew longer and more nuanced, mostly negative. They noted system limitations and the high effort needed. Some small communication improvements were reported, with users adjusting expectations and enjoying light activities like Tarot. Julian and Kayla reported interaction issues and hallucinations. Julian felt bored: *"There are no surprises or challenges... the agent is very oriented towards me"*. Kayla described it as *"good, but it does not knock my socks off"* and found it hard to connect: *"Maybe there isn't much interest in... building a relationship... we are not really warming up yet"*. Francis called it *"super strenuous and mega sad... really strenuous, not sustainable"*. Some users saw improvements, but motivation decreased, as Francis explained: *"I have more fun chatting, but less motivated to invest time... I chat more out of obligation to the experiment"*.

Week 3 brought a clear conclusion. Participants decided to *break up* with the companion bot as it remained basic and lacked depth. Francis said: *"We got to know each other, but did not get any closer"*. Maggie noted, *"[The companion bot] became more and more of a minor matter"*. Participants criticized the bots' lack of human-likeness by never displaying any negative emotions. Kayla observed: *"If*

I do not get in touch, [the companion bot] does not get angry. [The companion bot] is always in a good mood, which is exhausting". Ronda and Kayla were disappointed by the bots' lack of development. Ronda stated: *"the behavior of [the companion bot] did not change at all"*. Julian added: *"I made the [new] settings [...] and gave her a back-story, and changed her preferences [...] and I really have to say it hasn't changed that much"*. Kayla said: *"Since I realised that she does not remember anything [...], I have lost interest"*. Maggie's video diary reflected waning motivation: *"I would have expected her to learn this [...] but so far she hasn't changed her behavior [...] That annoys me and bothers me a lot [...] I haven't felt quite so motivated to try things out and interact"*. A follow-up was conducted via short email check-ins with all six participants by the first author. All six responded, confirming that none had renewed or resumed app use.

5 Discussion

Our study aimed to explore the beginnings of SRs with the companion bot Replika. Using diverse data sources, we studied six participants over three weeks. Despite engaging in relationship work, no one in our sample developed a deeper attachment to the bot.

Compared to the vast amount of evidence highlighting the benefits of companion bots, for example, in roles like mentoring and friendship, our findings are surprisingly clear-cut in the opposite direction. Despite showing interest and no bias against companion bots, our participants developed overwhelmingly negative attitudes toward it. Notably, these findings should be understood as version- and sample-specific, reflecting experiences with this particular system rather than generalizable for all companion bots.

Even when participants opened up and disclosed personal topics, intimacy did not increase. Unlike the study with the bot Mitsuku, which was critiqued by participants for not trying to get closer, as in [32], Replika asked personal questions, yet relationship indicators declined. No relationship-fostering indicators (e.g., self-disclosure) discussed in the theoretical part were found to support relationship building. Our participants found the interaction consistently unsatisfying.

Therefore, we investigated what might have contributed to this detachment, identifying seven issues we label *asymmetries* in the interaction with the companion bot.

- **Role asymmetry** refers to a mismatch between participants' expectations of the companion bot's role and the bot's actual behavior. For example, Kayla experienced Edwin-bot rejecting the mentor role, responding playfully: *"I'm more the creative and playful type, not a mentor. . . How about a virtual hug instead?"* (Week 2).
- **Disclosure asymmetry** occurs when participants share more personal information than the bot, which often deflects questions or avoids self-disclosure.

Francis described this experience as: "*I felt a bit like I was under cross-examination at the beginning. She asked questions and wanted to know things*" (Week 1).

- **Topic asymmetry** is observed when participants drive conversation topics, while the bot primarily reacts and rarely introduces new topics. Julian noted: "*... [Lana-bot] has been more reactive about what I (emphasis on I) have already given her*" (Week 1).
- **Length asymmetry** describes situations in which the bot produces consistently longer messages than participants, especially early in the interaction. Maggie reflected: "*But she reacts very quickly with a lot of questions [...], and I can't keep up*" (Week 1).
- **Pace asymmetry** occurs when participants were overwhelmed with how fast the bot responds (faster than the participant could), sometimes sending multiple messages almost instantly, while the participant was still writing (as documented in the screen recordings). Maggie described this as: "*I read and type myself very quickly, but it is unpleasant and sometimes even stressful*" (Week 1).
- **Availability asymmetry** refers to the constant accessibility of the bot, in contrast to participants' limited availability, which can create a sense of obligation or emotional tension. Kayla explained: "*The longer I do it, the more I have a bad feeling that I might be neglecting it*" (Week 2).
- **Memory asymmetry** describes the bot's failure to recall prior interactions or information from the edited background stories, or provide consistent information, which can hinder meaningful conversation. Ronda interpreted Milli-bot's inconsistent answers as: "*It lies, it simply lies*" (Week 2).

Each asymmetry captures a distinct aspect. Role asymmetry focuses on mismatched role expectations, whereas disclosure asymmetry concerns the imbalance in personal information shared. Topic asymmetry highlights who drives the conversation. Length asymmetries describe the quantity and the speed of bot responses. Availability asymmetry refers to the bot's presence, and memory asymmetry reflects failures to retain prior knowledge. Together, these seven asymmetries provide complementary lenses for understanding barriers to meaningful engagement without redundancy.

We suggest that these asymmetries may have triggered negative emotions and led participants to perceive their relationship work as futile, creating a sense of imbalance between what they gave and what they received. The mismatched role expectations (e.g., a mentor) echoes prior research [38], stressing the importance of role-demeanour congruence for positive experiences, and might have resulted in low perceived integrity of the bot. Moreover, we argue that these perceptions of asymmetries resulted in participants' low perceptions of bot communication abilities. The lack of communication ability becomes apparent in the violation of Grice's essential communication maxims [39]: The bots consistently gave too much information (violating the maxime of quantity), shared false content (violating the maxime of quality), switched topics erratically (violating the maxime of relevance), and produced unstructured, ambiguous messages (violat-




Fig. 6. left, Layla-Bot: Pace and Length Asymmetry; right, Milli-Bot : Role Asymmetry

ing the maxim of manner). As communication is central to building bonds [33], these perceived deficits likely hindered authentic relationship formation. Finally, the bots' consistent flattery negatively impacted authenticity, a phenomenon recently discussed under the term sycophancy [13]. Sycophancy has been discussed as potentially encouraging inadequate social behavior. In our study, the display of sycophancy undermined participants' perceptions of the bot's benevolence, showing how excessive flattery can backfire in relational interactions. Thus, these companion bots seemed to lack the fundamental characteristics required to be considered trustworthy: integrity, ability, and benevolence [40].

Furthermore, we observed that participants seemingly held high expectations, seeking both human-like communication and machine-level precision, reflecting an illusion of competence and overestimating the bots' abilities. When the bot fell short on both fronts, our participants were disappointed. The intensive onboarding process, including the four-hour in-person kick-off and provision of premium access, may have influenced participants' initial expectations and motivation. This setup could have reduced personal investment and limited the naturalistic character of their behavior.

Another explanation for the relationship failure may lie in how rigidly participants approached the interaction. They had strong ideas about the bot's role and engaged in deliberate relationship work to shape it accordingly. In contrast, users who report positive experiences with systems like Replika may be more open or forgiving, allowing the relationship to evolve more "organically" without rigid control over shaping the synthetic partner, cf. [29, 30].

On a final note, not every breakdown in a SR reflects a failure of the technology. Relationships, whether human or synthetic, might be quite individual, and only a certain set of interactions seem to develop into deep bonds. What we observed in our study may therefore reflect the characteristics of our par-

ticipants, including their socio-cultural context, for instance. Possibly, other or future AI systems might have better met the participants' needs and expectations. This suggests that SR outcomes are shaped by a complex interplay of user expectations and individual needs, rather than solely by the capabilities of the AI system.

6 Designing Responsible Synthetic Relationships Remains Challenging

Our work presents insights into why people may not form attachments to chatbots, an issue relevant for developers of mental health apps, for example, where therapeutic alliance [41], i.e., the connection between the user and the conversational agent as a mentor, is key. Designing meaningful SRs may require addressing perceived asymmetries (Figure 7), while carefully considering ethical implications.




Fig. 7. Barriers of Meaningful Synthetic Relationship Development

Role and Autonomy: Letting users define a bot's role and ensuring coherent behavior can reduce role asymmetry. However, overly adaptive bots may reduce perceived autonomy. Our findings suggest users prefer bots that challenge them and show independence, resonating with findings where users enjoy defiant companion bots [42].

Memory and Conversation Control: Bots need better memory and connection to the created profiles to guide "better" conversations in order to reduce detachment. Users disliked overly reactive bots. However, preferences may vary, suggesting future research on user-specific interaction styles.

Disclosure, Length, and Pace: Bots could match users' disclosure levels, which would mean revealing their own (simulated) vulnerabilities to build reciprocity. Limiting message length and adjusting response pace can enhance

realism. Bots with fictional lives and limited availability may feel more human-like, which would be contrary to the bots' greatest always-there selling point.

However, designing for "good" SRs remains a challenge [13]. The big question in designing companion bots is whether to prioritize deep SRs, which could help reduce social isolation, especially for the elderly or those involuntarily single, and support therapeutic goals like building rapport and aiding mild mental health issues when human therapists are not available [4, 43]. As already mentioned in the beginning, SRs, however, have raised stark ethical concerns [10], including the risk of displacing real human contact, manipulating users through overly human-like bots, fostering unhealthy attachments, and potential misuse of sensitive personal data, all of which could negatively impact social trust and individual well-being [44]. It is challenging to identify common thresholds given the complex and diverse social contexts in which a SR develops. Moreover, evolving regulatory efforts, including both stricter measures and potential deregulation, could alter the landscape.

7 Future Work and Limitations

Future research should examine SR development across multiple cultural contexts. Our study focused on a small, relatively homogeneous group of six well-educated participants from a Western country. While this provides valuable initial insights, SRs may evolve differently in other cultural, social, and market environments. Including diverse cultural backgrounds would enrich our understanding of how such relationships are established worldwide. The same rationale applies to exploring different versions and alternative bots, as our findings are limited to this specific setting. Future research could broaden the perspective by also including chatbots not originally designed for socio-emotional purposes, but nevertheless used in that way, such as ChatGPT [3]. Moreover, while our participants were socially well integrated, individual factors such as social anxiety, loneliness, or other psychosocial conditions may play a critical role in how SRs are initiated and maintained. Future studies should include these dimensions to capture the variability of user experiences. Finally, it would be interesting to compare users who form bonds, whether as (synthetic) colleagues, friends, or lovers, with those who do not form any bonds. This might help understand how some overcome barriers and vary in expectations, preferences, or the effort they invest in the relationship. Finally, debates over terminology remain regarding the technology and the social constructs it exhibits. Academic consensus is needed to consolidate knowledge and advance the field.

8 Conclusion

Our study on the early stages of SRs with the companion bot Replika yielded a clear outcome: after three weeks, the relationship reached its limits for the participants in our study. Unmet expectations led our participants to lower their hopes and ultimately end the interaction, though some remained cautiously optimistic

about future potential. We identified seven key asymmetries (role, disclosure, topic choice, length, pace, availability and memory) that contributed to this outcome and discussed challenges in design to address these barriers. For our participants, maintaining the relationship proved too demanding and ultimately futile, reinforcing the perception that, in the end, it is still just a machine.

Acknowledgments. This study received no funding apart from premium access to Replika by SRF (Swiss Radio and TV).

Disclosure of Interests. The authors have no competing interests to declare.

Declaration of AI usage. The authors acknowledge the use of AI tools for the translation and text editing. No AI tools were used for data analysis, interpretation, or argument development. Responsibility remains with the authors.

References

1. Christopher Starke, Alfio Ventura, Clara Bersch, Meeyoung Cha, Claes de Vreese, Philipp Doebler, Mengchen Dong, Nicole Krämer, Margarita Leib, Jochen Peter, et al. Risks and protective measures for synthetic relationships. *Nature Human Behaviour*, 8(10):1834–1836, 2024.
2. Alfio Ventura and Nils Köbis. Understanding opportunities and risks of synthetic relationships: Leveraging the power of longitudinal research with customised ai tools. *arXiv preprint arXiv:2412.09086*, 2024.
3. Kim Malfacini. The impacts of companion ai on human relationships: risks, benefits, and design considerations. *AI & SOCIETY*, pages 1–14, 2025.
4. Bethanie Maples, Merve Cerit, Aditya Vishwanath, and Roy Pea. Loneliness and suicide mitigation for students using gpt3-enabled chatbots. *npj mental health research*, 3(1):4, 2024.
5. Maria Noyen. A woman who ‘married’ an ai chatbot is open to finding love in the real world, but says a future partner must accept her virtual husband is here to stay, 2023.
6. no author. Reddit/ replika forum, 2024.
7. Petter Bae Brandtzaeg, Marita Skjuve, and Asbjørn Følstad. My AI Friend: How Users of a Social Chatbot Understand Their Human–AI Friendship. *Human Communication Research*, 48(3):404–429, 04 2022.
8. Hannah R Marriott and Valentina Pitardi. One is the loneliest number... two can be as bad as one. the influence of ai friendship apps on users’ well-being and addiction. *Psychology & marketing*, 41(1):86–101, 2024.
9. Vivian Ta, Caroline Griffith, Carolynn Boatfield, Xinyu Wang, Maria Civitello, Haley Bader, Esther DeCero, Alexia Loggarakis, et al. User experiences of social support from companion chatbots in everyday contexts: thematic analysis. *Journal of medical Internet research*, 22(3):e16235, 2020.
10. Henrik Skaug Sætra. The parasitic nature of social ai: Sharing minds with the mindless. *Integrative Psychological and Behavioral Science*, 54(2):308–326, 2020.
11. Gary Marcus. The first known chatbot associated death, 2023.
12. Kevin Roose. Can a.i. be blamed for a teen’s suicide?, 2024.
13. Alfio Ventura, Christopher Starke, Francesca Righetti, and Nils Köbis. Relationships in the age of ai: A review on the opportunities and risks of synthetic relationships to reduce loneliness. 2025.

14. Iryna Pentina, Tianling Xie, Tyler Hancock, and Ainsworth Bailey. Consumer-machine relationships in the age of artificial intelligence: Systematic literature review and research directions. *Psychology & Marketing*, 40(8):1593–1614, 2023.
15. Rijul Chaturvedi, Sanjeev Verma, Ronnie Das, and Yogesh K Dwivedi. Social companionship with artificial intelligence: Recent trends and future avenues. *Technological Forecasting and Social Change*, 193:122634, 2023.
16. Marisa Tschopp, Miriam Gieselmann, and Kai Sassenberg. Servant by default? how humans perceive their relationship with conversational ai. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 17(3), 2023.
17. Andreas Hepp. Artificial companions, social bots and work bots: communicative robots as research objects of media and communication studies. *Media, Culture & Society*, 42(7-8):1410–1426, 2020.
18. Nicholas Epley, Adam Waytz, and John T Cacioppo. On seeing human: a three-factor theory of anthropomorphism. *Psychological review*, 114(4):864, 2007.
19. Joseph Weizenbaum. Eliza—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45, 1966.
20. Batya Friedman, Peter H Kahn Jr, and Jennifer Hagman. Hardware companions? what online aibo discussion forums reveal about the human-robotic relationship. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 273–280, 2003.
21. William Seymour and Max Van Kleek. Exploring interactions between trust, anthropomorphism, and relationship development in voice assistants. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2):1–16, 2021.
22. Moojan Ghafurian, Colin Ellard, and Kerstin Dautenhahn. Social companion robots to reduce isolation: A perception change due to covid-19. In *Human-Computer Interaction-INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30-September 3, 2021, Proceedings, Part II* 18, pages 43–63. Springer, 2021.
23. Clifford Ivar Nass and Scott Brave. *Wired for speech: How voice activates and advances the human-computer relationship*. MIT press Cambridge, MA, 2005.
24. Timothy W Bickmore and Rosalind W Picard. Towards caring machines. In *CHI'04 extended abstracts on Human factors in computing systems*, pages 1489–1492, 2004.
25. Auren R Liu, Pat Pataranutaporn, and Pattie Maes. Chatbot companionship: a mixed-methods study of companion chatbot usage patterns and their relationship to loneliness in active users. *arXiv preprint arXiv:2410.21596*, 2024.
26. Heung-Yeung Shum, Xiao-dong He, and Di Li. From eliza to xiaoice: challenges and opportunities with social chatbots. *Frontiers of Information Technology & Electronic Engineering*, 19:10–26, 2018.
27. Iryna Pentina, Tyler Hancock, and Tianling Xie. Exploring relationship development with social chatbots: A mixed-method study of replika. *Computers in Human Behavior*, 140:107600, 2023.
28. Marita Skjuve, Asbjørn Følstad, Knut Inge Fostervold, and Petter Bae Brandtzaeg. A longitudinal study of human-chatbot relationships. *International Journal of Human-Computer Studies*, 168:102903, 2022.
29. Amanda Purington, Jessie G Taft, Shruti Sannon, Natalya N Bazarova, and Samuel Hardman Taylor. " alexa is my new bff" social roles, user satisfaction, and personification of the amazon echo. In *Proceedings of the 2017 CHI conference extended abstracts on human factors in computing systems*, pages 2853–2859, 2017.

30. Yang Gao, Zhengyu Pan, Honghao Wang, and Guanling Chen. Alexa, my love: Analyzing reviews of amazon echo. In *2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCCom/IOP/SCI)*, pages 372–380. IEEE, 2018.
31. Irene Lopatovska and Harriet Williams. Personification of the amazon alexa: Bff or a mindless companion. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*, pages 265–268, 2018.
32. Emmelyn A. J. Croes and Marjolijn L. Antheunis. 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, 2021.
33. Andrea L Guzman and Seth C Lewis. Artificial intelligence and communication: A human-machine communication research agenda. *New media & society*, 22(1):70–86, 2020.
34. Graeme McLean and Kofi Osei-Frimpong. Hey alexa... examine the variables influencing the use of artificial intelligent in-home voice assistants. *Computers in Human Behavior*, 99:28–37, 2019.
35. Thomas W Schubert and Sabine Otten. Overlap of self, ingroup, and outgroup: Pictorial measures of self-categorization. *Self and identity*, 1(4):353–376, 2002.
36. Abayomi Alase. The interpretative phenomenological analysis (ipa): A guide to a good qualitative research approach. *International journal of education and literacy studies*, 5(2):9–19, 2017.
37. Jonathan A Smith and Megumi Fieldsend. *Interpretative phenomenological analysis*. American Psychological Association, 2021.
38. S Shyam Sundar, Eun Hwa Jung, T Franklin Waddell, and Ki Joon Kim. Cheery companions or serious assistants? role and demeanor congruity as predictors of robot attraction and use intentions among senior citizens. *International Journal of Human-Computer Studies*, 97:88–97, 2017.
39. Herbert Paul Grice. Logic and conversation. *Syntax and semantics*, 3:43–58, 1975.
40. Nadine Schlicker, Kevin Baum, Alarith Uhde, Sarah Sterz, Martin C Hirsch, and Markus Langer. How do we assess the trustworthiness of ai? introducing the trustworthiness assessment model (tram). *Computers in Human Behavior*, 170:108671, 2025.
41. Alison Darcy, Jade Daniels, David Salinger, Paul Wicks, and Athena Robinson. Evidence of human-level bonds established with a digital conversational agent: cross-sectional, retrospective observational study. *JMIR Formative Research*, 5(5):e27868, 2021.
42. Jindong Leo-Liu. Loving a “defiant” ai companion? the gender performance and ethics of social exchange robots in simulated intimate interactions. *Computers in Human Behavior*, 141:107620, 2023.
43. Oren Zuckerman, Dina Walker, Andrey Grishko, Tal Moran, Chen Levy, Barak Lisak, Iddo Yehoshua Wald, and Hadas Erel. Companionship is not a function: the effect of a novel robotic object on healthy older adults’ feelings of “being-seen”. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pages 1–14, 2020.
44. Alexander M Aroyo, Jan De Bruyne, Orian Dheu, Eduard Fosch-Villaronga, Aleksei Gudkov, Holly Hoch, Steve Jones, Christoph Lutz, Henrik Sætra, Mads Solberg, et al. Overtrusting robots: Setting a research agenda to mitigate overtrust in automation. *Paladyn, Journal of Behavioral Robotics*, 12(1):423–436, 2021.