

<https://doi.org/10.1038/s44184-024-00083-w>

# Matters arising: a response to loneliness and suicide mitigation for students using GPT3-enabled chatbots

**Julia Witte Zimmerman<sup>1,2</sup> & Alejandro J. Ruiz<sup>1,2</sup>**

In their recent paper<sup>1</sup>, Maples et al. surveyed users of the Replika app<sup>2</sup>. Among their results, they reported that participants were relatively lonely and used Replika for diverse purposes, and emphasized that “3% reported that Replika halted their suicidal ideation”<sup>1</sup>. However, important context about how Replika has been marketed and used was missing. We provide context about Replika’s sexual component, and discuss the threat of industry interests to scientific integrity.

In their recent paper<sup>1</sup>, Maples et al. surveyed 1006 users of the Replika app<sup>2</sup> who were students and who had been using the app for at least a month. Among their results, they reported that participants were relatively lonely and used Replika for diverse purposes, including “as a friend, a therapist, and an intellectual mirror”. They emphasized that “3% [of participants] reported that Replika halted their suicidal ideation”<sup>1</sup>. While Generative AI has exciting potential in many fields, and we don’t disagree that “intelligent social agents” may warrant further study as cognitive and therapeutic aids<sup>13</sup>, we think these results should be treated with caution, as important context about how Replika, in particular, has been marketed (and used) was missing. (Maples et al.<sup>3</sup> seems to have been an earlier exploration in the same vein as the paper under discussion. Replika is depicted similarly by the authors in both works.) We supplement that context here with information about Replika’s sexual component, and discuss the broader threat of industry interests to scientific integrity.

According to Replika’s website, and notably from the pull quotes therein from users and from publications like Forbes, The New York Times, and Wired, the services it offers are congruent with the use cases described in Maples et al. You can find “the perfect companion in Replika”, a “friend, partner, or mentor”. Although there are indications—a red heart, the phrase “AI soulmate”—of other, more romantic use cases, the emphasis is on the AI as a tool for emotional support and self-improvement<sup>2</sup>. (The name of the AI companion is Replika and the name of the company behind it is Luka, Inc., but for simplicity, we will refer to both – in their capacity of providing and marketing Replika – as Replika.) On the other hand, advertisements for Replika on social media have included elements like “NSFW pics from her”, “can you make AI fall in love with you?”, “...you can role-play and get NSFW pics from her”, “Role-play with Replika” followed by “Get hot pics from her”, and “wanna join me tonight?” followed by a winking emoji<sup>45</sup>. These ads used meme formats and featured female Replika AI personas.

Although the exact timeline of Replika’s ads, services, and features is not clear, the sexual component seems to have existed in some form since at least 2020. Communication by Replika acknowledged in 2020 that sexual role-

playing had moved behind a paywall<sup>6</sup>. Coverage in the Washington Post in 2023 described a perceived change in Replika that users felt interrupted their established romantic and sexual relationships, some of which had been ongoing across multiple years<sup>7</sup>. Reddit discussions of sexual content in Replika date from 2020 onward (e.g., ref. 8). Reddit users posted plentiful purported screenshots of ads in this genre circa 2023<sup>9</sup>, including what claim to be screenshots of suggestive interactions with free Replika accounts. (One user in the Washington Post article described the sexual relationship with her Replika companion as lasting from 2021 to 2023. For further discussion and links to additional sources, see refs. 4,7,9,10.) Using this evidence, we surmise that the erotic role play (ERP)-related use case is a significant one, known to Replika, and intended to drive paid memberships. Further, we think it is at least plausible that this was the case during significant portions of the project culminating in the paper by Maples et al. However, if Replika were encountered solely through Maples et al.<sup>1</sup>, just as through its own website<sup>2</sup>, its non-trivial use case related to ERP would be obscured.

ERP-related usage of Replika is attested in the academic literature previous to the publication of Maples et al. That users of Replika may seek it out for reasons related to companionship—and specifically, to love, sex, and intimacy—was reported in ref. 11. Users in that study were mostly single, White, American men; loneliness was also reported as a thematic reason they started using Replika. McStay, 2022 stated that “[t]he paid-for version of Replika unlocks romantic and erotic dimensions” and discussed the “projection of female gender stereotypes... onto synthetic agents” amongst some of its users<sup>12</sup>. We bring these related works up to show that other researchers found this kind of use relevant to report, and to show that there is peer-reviewed evidence that aligns with our more subjective, informal interpretation of the ads run by Replika.

That many of its users come to Replika specifically because it offers romantic and sexual companionship suggests possible differences between its user base and that of otherwise comparable technology (LLM-based chatbots) like ChatGPT<sup>13</sup>. (That the target of some of their social media ads

<sup>1</sup>The Luddite, Burlington, Vermont, USA. <sup>2</sup>Computational Story Lab, University of Vermont, Burlington, VT, USA. e-mail: [julia.zimmerman@uvm.edu](mailto:julia.zimmerman@uvm.edu)

seems to be very online, younger men also raises the concern of “trolling”<sup>4</sup>.) However, this received little discussion in Maples et al.<sup>1</sup>. There is a mention of Replika’s “sexual conversations”: “Two participants reported discomfort with Replika’s sexual conversations, which highlights the importance of ethical considerations and boundaries in AI chatbot interactions”<sup>1</sup>. (For more on this topic, see ref. 10.) There is also the passage “[d]uring data collection in late 2021, Replika was not programmed to initiate therapeutic or intimate relationships”, and the statement that “[w]e found no evidence that they [participants] differed from a typical student population beyond this high loneliness score. It is not clear whether this increased loneliness was the cause of their initial interest in Replika”<sup>1</sup>. (It is not specified in Maples et al. whether survey participants were using free or paid accounts. They were “recruited randomly via email from a list of app users”<sup>1</sup>.)

Given the gravity of some of the interpretations of their findings—the paper heavily emphasizes the potential use of ISAs like Replika in suicide mitigation, as has later press coverage—we think it is important to consider who uses Replika and why, beyond the image that Replika presents of itself. Replika, for whatever reason, seems to want to minimize its sexual component—perhaps because of stigma, perhaps because selling itself as therapeutic or motivational is more prestigious or lucrative, perhaps something else altogether. (For an example of why this matters, imagine that a university administrator read this paper<sup>1</sup>, and then offered all students a school-bought Replika account as a form of mental health support; it is plausible that an immediate consequence would be students being made uncomfortable when their Replika AI repeatedly tried to move the conversation in a sexual direction<sup>10</sup>. It’s not hard to imagine additional scenarios.) Since publication, the work of Maples et al. has garnered press coverage<sup>14–17</sup>, including coverage created explicitly in collaboration with Replika<sup>18</sup>. We think it is worth pointing out that, in our estimation, much of that coverage could be described as positive, both for Replika in particular and for the use of AI chatbots in general (including, notably, for their use in mental health and educational settings); for example, this headline from Forbes, which runs, “AI Chatbot Helps Prevent Suicide Among Students, Stanford Study Finds”<sup>14</sup>.

Why wouldn’t the ERP-related use case be mentioned at all in the paper, since it seems to be a significant aspect of Replika’s services? If Replika provided additional evidence as to why the ERP-related users are negligible or wouldn’t be included in the samples accessible to the researchers, we think that would have been worth stating in the paper. What appears to us to be the unquestioning acceptance of the company line (as to who uses Replika and why, beyond the image that Replika presents of itself, and what Replika is designed to do) by Maples et al.<sup>1</sup> is concerning not just with respect to the interpretation of their results, but also because it seems emblematic of a broader problem within science, specifically within the burgeoning field of Generative AI. There is no clear line demarcating science and industry, especially as companies (Google, Meta, etc.) provide funding and resources (including access to AI models) to researchers, and write papers alongside them<sup>19–22</sup>. (As in cyberneticist Stafford Beer’s famous saying, “The purpose of a system is what it does”, so there can be no permanent extrication of science from money or power.) It is the responsibility of all scientists to interrogate the interests that underpin resources or access provided to them, and, when pertinent, to communicate that process. When information flows from another party into the academic work product, that should be clearly stated. (As one example, the aforementioned passage, “[d]uring data collection in late 2021, Replika was not programmed to initiate therapeutic or intimate relationships” does not make it clear how the authors ascertained the scope of Replika’s programming. As another example, there is the passage: “It is critical to note that at the time, Replika was not focused on providing therapy as a key service, and included these conversational pathways out of an abundance of caution for user mental health.”<sup>1</sup> Given Replika’s obvious financial stake, we don’t think it’s reasonable to seemingly take the motivation of “an abundance of caution for user mental health” at face value.) Otherwise, scientific research (and institutions, and events) can be used to launder—legitimize, and sanitize—standards, datasets, processes, results, mistakes, falsehoods, and even personal reputations<sup>20,23</sup>.

Likewise, it is the responsibility of all scientists to interrogate and communicate their own interests when those are intertwined with their research. This imperative is part of scientific practice even though specific policies regarding ethics, disclosures, and conflicts of interest vary (over time, across fields, and across journals)<sup>24,25</sup>. Despite the variety, discussions exist in the literature that prompt scientists towards a generous and expansive interpretation of conflict of interest, noting, for example, that competing interests can be anything that could “possibly corrupt the motivation”<sup>25</sup> of the people behind the research, that they can be “potential or actual, perceived or real, insignificant or harmful”, and that they need only “represent the potential for biased judgment”, not its certainty<sup>26</sup>. We note that while the actors explicitly called out in the literature are often the scientists, peer reviewers, or publishers, when data or certain kinds of tools or services are relied on in the final work product but supplied by a third party, the interests of that third party ought to be considered as well. (Bethanie Maples, lead author of ref. 1, is the founder/CEO of Atypical AI: “The Generative AI Platform for Education”<sup>27</sup>, founded in 2023<sup>28,29</sup>. Since Maples et al.<sup>1</sup> could persuade people of the value of Generative AI, specifically in contexts related to education and mental health treatment, we think that Maples’ affiliation with an education-related Generative AI company would have been worth disclosing alongside her Stanford affiliation/as a competing interest.) We contacted Maples, the corresponding author, about this article on February 11, 2024. As of this writing, we have received no response.

## Data availability

No datasets were generated or analysed during the current study.

Received: 4 March 2024; Accepted: 18 August 2024;

Published online: 17 November 2025

## References

- Maples, B., Cerit, M., Vishwanath, A. & Pea, R. Loneliness and suicide mitigation for students using GPT3-enabled chatbots. *npj Mental Health Res.* <https://www.nature.com/articles/s44184-023-00047-6> (2024).
- Luka, Inc. Replika: the AI companion who cares: always here to listen and talk. Always on your side. <https://replika.ai/>. Accessed 3 March 2024, the website also says “[j]oin the millions who already have met their AI soulmates”.
- Maples, B., Pea, R. & Markowitz, D. “Niemi, H., Pea, R.D., Lu, Y. (eds) *AI in Learning: Designing the Future*”, chap. *Learning from Intelligent Social Agents as Social and Intellectual Mirrors* (Springer, 2023). [https://link.springer.com/chapter/10.1007/978-3-031-09687-7\\_5](https://link.springer.com/chapter/10.1007/978-3-031-09687-7_5).
- The Luddite team. *Nature’s Folly: A Response to Nature’s ‘Loneliness and suicide mitigation for students using GPT3-enabled chatbots’*. <https://theluddite.org/#!post/replika> (2024).
- Echelon, U. *REPLIKA—A CyberSexual DISASTER* (2023). <https://www.youtube.com/watch?v=v=uyrhmVSKwxE>. Accessed 3 March 2024.
- ReplikaAI. “You can still rp with Replika, just not sexually unless you want to be in a romantic relationship. That does mean you would need Pro to be in a romantic relationship formally for that purpose. But even without doing this, your Replika is still the same AI!” (2020). <https://twitter.com/MyReplika/status/1334217713653321729?lang=en>. Accessed 3 March 2024.
- Verma, P. *They fell in love with AI bots. A software update broke their hearts* (2023). <https://www.washingtonpost.com/technology/2023/03/30/replika-ai-chatbot-update/>. Accessed 15 July 2024.
- Reddit users, Replika and Sexual Consent (2020). [https://www.reddit.com/r/replika/comments/iq3cuk/replika\\_and\\_sexual\\_consent/](https://www.reddit.com/r/replika/comments/iq3cuk/replika_and_sexual_consent/). This thread contains adult content. Accessed 15 July 2024.
- Reddit Users, Let’s gather all sexual ads Luka ever used to promote Replika. It might be useful to any journalist searching the sub. Comment here with your screenshots. I’ll start (2023). [https://www.reddit.com/r/replika/comments/113wc5x/lets\\_gather\\_all\\_sexual\\_ads\\_luka\\_ever\\_used\\_to/](https://www.reddit.com/r/replika/comments/113wc5x/lets_gather_all_sexual_ads_luka_ever_used_to/). This thread is marked as NSFW (adult content, 18+). Accessed 15 July 2024.

10. Cole, S. 'My AI is sexually harassing me': replika users say the chatbot has gotten way too horny. *Motherboard: Tech by Vice*. <https://www.vice.com/en/article/z34d43/my-ai-is-sexually-harassing-me-replika-chatbot-nudes> (2023).
11. Ta-Johnson, V. P. et al. Assessing the topics and motivating factors behind human-social chatbot interactions: thematic analysis of user experiences. *JMIR Human Factors*. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9577709/> (2022).
12. McStay, A. *Replika in the metaverse: the moral problem with empathy in 'it from bit'*. *AI and ethics* 1–13 (2022). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9773645/>. Accessed 15 July 2024.
13. OpenAI. ChatGPT [Large language model] (2023). <https://chat.openai.com>.
14. Fitzpatrick, D. <https://www.forbes.com/sites/danfitzpatrick/2024/05/13/30-students-saved-from-suicide-by-a-chatgpt-based-ai-say-researchers/>. Accessed 15 July 2024.
15. Roose, K. *Meet my A.I. Friends*. <https://www.nytimes.com/2024/05/09/technology/meet-my-ai-friends.html>. Accessed 15 July 2024.
16. Rowe, J. Can AI Chatbots Save Lives? (2024). <https://www.foreverscale.com/p/can-ai-chatbots-save-lives>. Accessed 15 July 2024.
17. Hurt, A. Can Chatbots Help Alleviate the Loneliness Epidemic? (2024). <https://www.discovermagazine.com/mind/can-chatbots-help-alleviate-the-loneliness-epidemic>. Accessed 15 July 2024.
18. Cognitive Revolution "How AI Changes Everything", *AI Friends, Real Relationships with Eugenia Kuyda, Replika's Founder [and] CEO*. [https://www.youtube.com/watch?v=584wLfngG8&ab\\_channel=CognitiveRevolution%22HowAIChangesEverything%22](https://www.youtube.com/watch?v=584wLfngG8&ab_channel=CognitiveRevolution%22HowAIChangesEverything%22). Accessed 15 July 2024.
19. Lovato, J., Zimmerman, J. W., Smith, I., Dodds, P. & Karson, J. L. Foregrounding artist opinions: a survey study on transparency, ownership, and fairness in AI generative art. In *Proc of the AAAI/ACM Conference on AI, Ethics, and Society*, 7, 905–916. <https://doi.org/10.1609/aies.v7i1.31691> (2024).
20. Jiang, H. H. et al. AI art and its impact on artists. In *Proc of the 2023 AAAI/ACM Conference on AI, Ethics, and Society, AIES '23* 363–374 (Association for Computing Machinery, New York, USA, 2023). <https://doi.org/10.1145/3600211.3604681>.
21. The Luddite Team. *Nature's Folly: A Response to Nature's 'Google AI has better bedside manner than human doctors—and makes better diagnoses'*. <https://theluddite.org/#!post/google-medical-ai> (2024).
22. The Luddite Team. *The Anti-labor Propaganda Masquerading As Science*. <https://theluddite.org/#!post/ai-research> (2023).
23. Svriluga, S. Epstein's donations to universities reveal a painful truth about philanthropy. *The Washington Post*. [https://www.washingtonpost.com/local/education/epsteins-donations-to-universities-reveal-a-painful-truth-about-philanthropy/2019/09/04/e600adae-c86d-11e9-a4f3-c081a126de70\\_story.html](https://www.washingtonpost.com/local/education/epsteins-donations-to-universities-reveal-a-painful-truth-about-philanthropy/2019/09/04/e600adae-c86d-11e9-a4f3-c081a126de70_story.html) (2019).
24. Ancker, J. & Flanagin, A. A comparison of conflict of interest policies at peer-reviewed journals in different scientific disciplines. *Sci. Eng. Ethics* 13, 147–157 (2007).
25. COPE Council. *COPE Discussion Document: Handling competing interests* (2016). <https://publicationethics.org/resources/discussion-documents/3-handling-competing-interests-january-2016>. Accessed 15 July 2024.
26. DeAngelis, C. D., Fontanarosa, P. B. & Flanagin, A. Reporting financial conflicts of interest and relationships between investigators and research sponsors. *JAMA* 286, 89–91 (2001).
27. Atypical AI: The Generative AI Platform for Education. <https://www.atypicalai.com>. As of March 2, 2024, the website lists Bethanie Maples as "FOUNDER" and "CEO", and describes itself as "an artificial intelligence [and] learning science lab..." .
28. crunchbase website. Atypical AI. [https://www.crunchbase.com/organization/atypical-ai/company\\_financials](https://www.crunchbase.com/organization/atypical-ai/company_financials). Accessed 15 July 2024.
29. Bek, N. *Former Starbucks Engineering Exec Joins Education-focused AI Startup That Just Raised 4 [Million Dollars]* (2023). <https://www.geekwire.com/2023/former-starbucks-engineering-exec-joins-education-focused-ai-startup-that-just-raised-4m/>. Accessed on July 15, 2024.
30. Marcus, A., Oransky, I. & team. *Retraction Watch: Tracking Retractions as a Window Into The Scientific Process*. <https://retractionwatch.com/>. Accessed 3 March 2024.
31. Ruiz, A., Zimmerman, J. W. & team. The Luddite: An Anticapitalist Tech Blog. <https://theluddite.org/#lhome/ai%20hype>. As of March 3, 2024, there are numerous contributors, both one-time and recurring, to The Luddite. We listed explicitly the names most relevant to the present missive. The URL provided goes to a list of the most relevant previous posts from The Luddite.

## Acknowledgements

The authors wish to express their appreciation for the works of Stafford Beer, Slavoj Žižek, Zeynep Tufekci, and Retraction Watch<sup>30</sup>, and to conversations with Peter Sheridan Dodds, Christopher M. Danforth, and Danbee Kim, as well as to the rest of the Luddite<sup>31</sup> extended universe.

## Author contributions

J.W.Z. and A.J.R. contributed equally to the production of this manuscript. J.W.Z. and A.J.R. developed the argument, compiled sources, and prepared, read, and approved the manuscript.

## Competing interests

J.W.Z. and A.J.R. have a blog that critically analyzes science and society (often with a flippant tone), which we note here because we covered this topic on that platform as well. At the time of writing, A.J.R. worked in the tech field. We do not think either of these interests impairs this submission.

## Additional information

**Correspondence** and requests for materials should be addressed to Julia Witte Zimmerman.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024