Raise, Don't Replace: Two Kinds of Connection in AI Character Products

Jinghan Zhu
Kotoko AI, London, UK
jzhu@kotoko.ai

Yunjia Guo
Kotoko AI, London, UK
yguo@kotoko.ai

Some new consumer apps promise to do your socialising for you: an avatar that likes, comments, breaks the ice, and screens people on your behalf. We think this confuses two different things. Connecting with another person depends on that person actually showing up; hand the act to a stand-in and most of its value goes, which is part of why a message believed to be written by AI is trusted and felt less. Connecting with a virtual character is not like that. The character is openly made, and what carries the bond is what you put into it, not whether some hidden human meant it. Automating the first kind tends to hollow it out, and an avatar acting as you is unlikely to build the second kind either. The opening many products are missing, we argue, is that second kind: the bond with a character you make and raise. Drawing on a deployed AI-character game and 20 interviews with its creator-users, we describe what seems to carry that bond, authoring, a character that stays coherent, care over time, and seeing it in your own life, and report something we did not expect: most (though not all) of our most committed users were lukewarm on AI conversation, yet were deeply attached to their characters.

AI characters; human–AI relationships; parasocial interaction; original characters; social software; position paper

1. INTRODUCTION

Several recent products treat AI's role in social life as delegation. Elys, an AI-native social network backed by major investors, gives each user a "cyber avatar" that scans feeds, likes, comments, breaks the ice, and screens potential connections, handing control back to the person only at the important moments [1]. The pitch is that AI can take on the work of socialising so you do not always have to be there.

Our claim is easy to state and harder to settle: connecting with a person and connecting with a made character are different in kind, and designs that treat the first like the second go wrong. The first half of this paper makes that case and works through the obvious objections, including the fact that "AI in a social product" covers very different things, from drafting a message to acting as you. The second half turns to a product we have built and the people who use it, not to prove the claim but to show what the alternative looks like in practice. This is a position paper, so the weight is on the argument.

We make three contributions:

1. a distinction between human-to-human (H2H) and human-to-virtual-being (H2V) connection, drawn by what makes each one worth having, and the diagnosis it supports: that proxy-avatar products (e.g., Elys) appear to conflate the two;
2. the position that the more promising, and currently neglected, goal is the H2V bond done well, with a four-part account of what that involves, grounded in a deployed product;
3. a finding that cuts against the grain of "AI social": most (not all) of our most invested users were cool on AI conversation, yet bonded strongly with their characters, which suggests the bond grows more from making and caring than from talking.

2. RELATED WORK

Three lines of work inform the argument; what is new is how they fit together.

AI-mediated communication and authenticity. When AI writes, or is thought to write, what people say to each other, trust and felt authenticity can drop. Jakesch et al. found profiles *believed* AI-written were discounted, an effect driven by

© Zhu et al. Published by BCS Learning and Development Ltd.

the suspicion of automation rather than by the text [2], part of a broader research agenda on AI-mediated communication [3]. Turkle made a related worry earlier: that we come to "expect more from technology and less from each other" [4].

Bonds with AI as their own thing. Parasocial-interaction theory describes one-sided but genuinely felt bonds with a media figure [5]. Two further lines explain why such bonds are neither surprising nor second-rate. The Computers-Are-Social-Actors work showed that people apply social rules to machines they know are machines [6], and Turkle's studies of relational artifacts, from the Tamagotchi onward, found that attachment often grows from the act of caring for the artifact rather than from its conversational competence [7]. That is an early version of the pattern we report in Section 4. In games, player-avatar research likewise distinguishes several stable modes of relating to a made character [8]. People form real, developing relationships with social chatbots [9]; this can also go wrong [10], and users hold tangled beliefs about what the AI other even is [11]. Closest to us, Sun et al. find players relate to AI-generated game characters in distinct ways, as themselves, as their child, or as their friend [12].

Original characters and shared play. Making an "original character" (OC) is a widespread way of building a personal "digital life," and the tools have to leave the creator in control [13]. Separately, sharing an experience with strangers in a game can build social ties, with games working as "third places" [14].

Constitutive value in philosophy. The intuition that some acts lose their point when delegated has older roots outside HCI. Philosophical accounts of friendship since Aristotle treat the friend's own choice and effort as partly constitutive of the relationship's value, not merely instrumental to it; work on friendship and the self makes the same point in contemporary terms [15], and recent discussion of relationship goods extends it [16]. We borrow the spirit of these accounts rather than their machinery: "warrant," below, names what they would call the constitutive condition of a relationship's value.

Each of these names part of what we are after, authenticity, agency, parasocial attachment, social presence, but none asks the question we need: what makes a particular connection worth having in the first place, and therefore what gets lost when you automate it. That is the question the rest of the paper turns on.

3. TWO KINDS OF CONNECTION

3.1. What makes each one worth having

We use *warrant* as shorthand for the thing that makes a connection count as what it is. H2H and H2V, we argue, have warrants that cannot stand in for each other.

For H2H, the warrant is that another person, who could have done otherwise, chose to engage with you. A quick test makes this concrete: replace the person with a stand-in but keep the visible act, a like, a comment, an opener, identical, and the value still leaks away. The research on AI-mediated communication keeps finding this signature in different places: a profile is trusted less when people think AI wrote it [2]; the same supportive message lands as *less heard* once it is labelled AI [17]; a thanks or an apology says *less about who you are* when marked AI-assisted [18]. The throughline is that an act is worth something partly because it shows your own intention and effort, and handing it off removes that. These are specific settings, so we treat the general version as a hypothesis rather than a proven law (see Section 5).

For H2V, the warrant is different: it is your own investment in something everyone knows is made. There is no hidden sincerity for automation to expose, so automation is not a betrayal here; it is the premise. What carries the bond is what you put in, and whether the character holds together over time. The investment can look quite different across products, authoring and tending a character on our platform, or self-disclosure and projection with a chatbot like Replika [9], but in each case it is the user's own investment in an openly artificial other doing the work, which is exactly the kind of one-sided yet real bond parasocial theory described long ago [5].

So the two are not near and far points on one "how human" scale. One is wrecked by substitution; the other is built on artificiality. Table 1 Zhu et al.Table 1 sets them side by side, summarising the framework as five contrast dimensions: what warrants each connection, how it responds to substitution, how it responds to automation, what counts as "real," and how each characteristically fails., summarising the framework as five contrast dimensions: what warrants each connection, how it responds to substitution, how it responds to automation, what counts as "real," and how each characteristically fails.

3.2. Not all "AI in social" is the same

It matters which job the AI is doing. It might **draft** something you still choose to send, **assist** you

Table 1: *Two kinds of connection (schematic; the H2H column states tendencies we argue for, not measured results).*

	Human-to-human (H2H)	**Human-to-virtual-being (H2V)**
What makes it count	Another person freely choosing to engage	Your own investment in something openly made
Replace the actor with a stand-in?	Most of the value tends to go, even if the act looks the same [2; 17; 18]	Fine; the thing was made to begin with
Automation	Tends to erode it as it scales [2]	Part of what it is
What "real" means here	The other person's sincere choice	Your investment, plus a character that stays coherent
How it fails	Hollowed out; discounted as inauthentic	Incoherent, uncanny, or inconsistent
Examples	Shared experience, then a chosen connection (games as "third places" [14])	Authoring and care (our platform); disclosure and projection (Replika [9])

(rank a feed, surface people, nudge you, suggest an opener), **act on its own** inside some space, or **stand in as you** to other people. Our complaint is aimed only at the last one. Drafting and assistance can be fine, because you are still the one who decides to send, to connect, or to show up. The trouble starts when the AI performs, as you, the acts whose value came from your choosing to do them.

Proxy products like Elys, where the avatar likes, comments, breaks the ice, and screens *as the user*, look to us like this confusion in practice. This is not our reading imposed on the product: in a published interview, Elys's founder frames the goal as an autonomous entity "acting proactively on your behalf" and treats interpersonal friction as entropy for AI to eliminate [19]. The design does keep posting and the final confirmation of a relationship in human hands—the avatar "can comment, like, or react" but "never post or send friend requests" [19]—but on our account the damage is done earlier, since the acts being automated—the like, the comment, the opener—are precisely the ones whose value lay in a person choosing to perform them. What the design treats as friction to remove is, on our view, the warrant itself. Automating those acts likely undercuts the very thing that made them worth anything, even as it scales them up [2]; and because the avatar stands in for the person rather than being an openly made character you raise, it is unlikely to grow the other kind of bond either. On this reading it risks serving neither, though that is a judgment about a class of products, not a tested result. (Our own platform also has AI characters acting socially, but they are openly virtual beings the user raises, not stand-ins for absent people; any human-to-human tie on the platform forms only after real people share something and choose to connect, Section 4.)

3.3. How the claim could be wrong

On this account, an AI character product is better off aiming at the H2V bond than at automating people's social lives, and where human ties do form, the move is to give people something to share, not to act for them. The claim is not safe from evidence: it would weaken if people drew lasting value from connections they knew were AI-performed—longitudinal use data from proxy products such as Elys would test this directly—or if the H2V bond turned out to ride mainly on how well the character talks. Our case points the other way, but it is one case.

4. AN ILLUSTRATIVE CASE

Method. We draw on 20 semi-structured interviews with creator-users of the platform.[1] We recruited through in-game and Discord channels: a short survey went out in and around the game, and respondents who said they cared about original-character (OC) creation were invited to talk. Interviews were by voice, ran 40–60 minutes, and loosely followed a guide on creating characters, everyday use, and what people liked or disliked. Participation was voluntary, and everyone consented to our use of anonymised quotes. The sample is purposive, not representative: it leans toward invested OC creators and toward US and EU players, covers both mobile and PC, and spans from a few months to over a decade of OC experience. The interview guide and the analysis were developed collaboratively by two of the authors together with a colleague; interviews were audio-recorded, and we worked from the recordings and our notes, reading and grouping the material inductively into the themes below [20]. AI tools were used only to assist literature

[1] Participant names are pseudonyms; quotes are lightly edited for length and clarity.

search and language editing, not for data collection or analysis. We offer this as a case for the argument, not a controlled study, and we are part of the team that built the product, a bias we name rather than pretend away. On the platform, users build characters from text and tags, chat with and dress them, take them through branching adventures, and, with a feature that composites a character into the user's own video, place them in real-world footage.

What follows is not tidy. People wanted different, sometimes opposite, things, and the disagreements are part of the point.

Authoring, and how much control people want. The bond starts with making a character that is unmistakably yours, and tags were the main lever; one long-time creator called them the basis for "building personality." But how much fidelity people needed varied a lot. A novelist-illustrator who has drawn OCs since she was twelve said the tools got her only "70%" of the way, missing hair variants, tails, wings, and clean layering, and that gap was reason enough to keep using an older tool. Others did not want that much control at all: one daily player, whose character was still taking shape, liked letting the AI improvise and was happy to be surprised. So "authoring" was not one thing; it ranged from people who wanted to pin down every detail to people who wanted room to discover the character.

Life, and whether that means conversation. People wanted the character to feel alive, but they disagreed sharply about whether that meant talking to it. Some were openly cold on AI chat: one called the dialogue "stale" and would read but never reply; another faulted Character.AI for canned lines; one creator's tarot-reader character even turned evasive and, after she missed a day, told her "don't do that again," which she found frightening, the kind of harm others have documented for companion AI [10]. But not everyone felt this way. One heavy user wanted the *opposite*: a character with real memory and "its own thoughts," a non-judgmental presence he leaned on through a rough stretch, for him the conversation was the point. The common thread was not "more talk" or "less talk" but coherence and consequence, a character that remembers and stays itself; people who got that bonded, whether or not chat was central for them.

Care, and whether it should have stakes. Much of the bond grew through tending the character, dressing it, giving gifts, coming back. One daily player wanted that to bite: a "friendship bar" with real downsides for neglect and rewards for attention, so the relationship could change. Others wanted the reverse: no chores, no pressure; one prized the character precisely as company "without forcing daily tasks or payments," and a creator who values the tool as a serious instrument bristled at "meaningless grinding." So care mattered to nearly everyone, but whether it should carry obligation split the room.

Seeing the character in your own life. The most vivid moments came when the character showed up in the real world. Using that feature, one daily user turned footage of her bedroom into a clip of her and a friend's characters "eating biscuits on my bed," called it "super cute," and was struck that their faces carried their personalities, her bubbly OC beside a friend's "silent type." A creator who is generally cool on AI lit up at a demo dropping a character into a real room; another said he would pay for this, but only if it left the original footage untouched. And not everyone cared: one writer, who treats the tool as a way to finalise characters for her fiction, was indifferent to anything social-facing at all—we retain her as a negative case rather than an exception. Where it landed, though, it landed hard.

Underneath the disagreements, people kept describing not an assistant but something they own and raise, a character spoken of as "my child," as a hoped-for version of themselves, or as a presence that does not judge. We report that last note carefully and make no clinical claims. What recurred, in this sample, was that the attachment sat in making, keeping, and showing the character more than in conversation with it.

5. DISCUSSION

For designers. Two questions are worth asking before building: which connection are you serving, and which job is the AI doing. If you want people to connect with each other, give them something to share, or help a person who still chooses to act, rather than acting as them. If you want the bond with the character, spend the effort on letting people author it well, keeping it coherent, giving them ways to care, and bringing it into their lives, before spending it on heavier chat. The proxy bet is the risky one: it automates the part whose value depends on a person being there, and skips the bond people actually form.

The part we did not expect. That several committed creators disliked AI chat yet were deeply attached to their characters is the most useful thing we found, and it sets a boundary on the claim. The bond, here, came mostly from making and caring. But the disagreement above, one user wanting exactly the deep conversation others rejected, is a caution against overstating it: conversation is not nothing, it is just not where most of these users located the bond.

Is the finding just the definition echoing back? One might object that our finding repeats our premise: having defined the H2V warrant as the user's own investment, we then report that the bond rests on investment. But the definition does not predict which *form* of investment carries the bond. Conversation is itself a form of investment—the self-disclosure Replika users describe is exactly that [9]—so the definition is compatible with a world where the bond runs mainly through talk. That, in this sample, it ran instead through making, tending, and showing the character is an empirical observation the definition did not supply.

Limits. It is worth being explicit about what to trust how far. The conceptual distinction, that H2H and H2V differ by what makes each worth having, is the claim we hold firmly. The read on specific products, and the design advice that follows, are weaker: they rest on this one case and on extrapolating prior findings, not on controlled evidence. The substitution effect we build on comes from particular settings, profiles [2], supportive messages [17], short warmth signals [18], and stretching it to interpersonal acts in general is a hypothesis; high-stakes or long-running ties may behave differently. The case is one product and a self-selected, mostly US/EU set of highly invested creators, studied by the team that built it. Whether the distinction holds up elsewhere, and how these bonds change over months and years, is what we want to look at next. How these bonds look from the outside, to the people who interact with a proxy or who encounter someone else's raised character, is a further question we have not touched.

6. CONCLUSION

Connecting with a person and connecting with a made character are not the same thing, and the difference is not how human the other party seems; it is what makes the connection worth having. Read that way, the avatar that acts as you is solving the wrong problem: it automates the one part that needed a person, and never becomes a character you could care about. The opening it misses is the one our users kept showing us, a character you make, keep coherent, look after, and finally see in your own life. The surprise was how little of that ran through conversation.

REFERENCES

[1] APPSO / 36Kr. 2026. The AI Avatar Surfs the Internet 24/7: How Does It Make Friends Better Than Me? 36Kr, 22 May 2026. `https://eu.36kr.com/en/p/3819846320722050`.

[2] Maurice Jakesch, Megan French, Xiao Ma, Jeffrey T. Hancock, and Mor Naaman. 2019. AI-Mediated Communication: How the Perception that Profile Text was Written by AI Affects Trustworthiness. In *Proc. CHI 2019*. DOI:10.1145/3290605.3300469.

[3] Jeffrey T. Hancock, Mor Naaman, and Karen Levy. 2020. AI-Mediated Communication: Definition, Research Agenda, and Ethical Considerations. *Journal of Computer-Mediated Communication* 25, 1 (2020), 89–100. DOI:10.1093/jcmc/zmz022.

[4] Sherry Turkle. 2011. *Alone Together: Why We Expect More from Technology and Less from Each Other*. Basic Books.

[5] Donald Horton and R. Richard Wohl. 1956. Mass Communication and Para-Social Interaction: Observations on Intimacy at a Distance. *Psychiatry* 19, 3 (1956), 215–229.

[6] Clifford Nass, Jonathan Steuer, and Ellen R. Tauber. 1994. Computers are Social Actors. In *Proc. CHI 1994*, 72–78. DOI:10.1145/191666.191703.

[7] Sherry Turkle. 2007. Authenticity in the Age of Digital Companions. *Interaction Studies* 8, 3 (2007), 501–517. DOI:10.1075/is.8.3.11tur.

[8] Jaime Banks. 2015. Object, Me, Symbiote, Other: A Social Typology of Player-Avatar Relationships. *First Monday* 20, 2 (2015). DOI:10.5210/fm.v20i2.5433.

[9] Marita Skjuve, Asbjørn Følstad, Knut Inge Fostervold, and Petter Bae Brandtzaeg. 2021. My Chatbot Companion — a Study of Human-Chatbot Relationships. *International Journal of Human-Computer Studies* 149 (2021), 102601. DOI:10.1016/j.ijhcs.2021.102601.

[10] Renwen Zhang, Han Li, Han Meng, Jinyuan Zhan, Hongyuan Gan, and Yi-Chieh Lee. 2025. The Dark Side of AI Companionship: A Taxonomy of Harmful Algorithmic Behaviors in Human-AI Relationships. In *Proc. CHI 2025*. DOI:10.1145/3706598.3713429.

[11] Aikaterina Manoli, 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. In *Proc. CHI 2026*. DOI:10.1145/3772318.3791331.

[12] Qirui Sun, Haixin Qiao, Jinghan Zhu, Yunjia Guo, Yifan Jin, Teli Yuan, Siyu Wang, and Haipeng Mi. 2025. Myself, My Child or My Friend: Understanding Player

Relationships with AI-Generated Game Characters. In *Ext. Abstracts CHI 2025*. DOI:10.1145/3706599.3720207.

[13] Yuqian Sun, Xingyu Li, Shunyu Yao, Noura Howell, Tristan Braud, Chang Hee Lee, and Ali Asadipour. 2025. ORIBA: Exploring LLM-Driven Role-Play Chatbot as a Creativity Support Tool for Original Character Artists. arXiv:2512.12630.

[14] Constance A. Steinkuehler and Dmitri Williams. 2006. Where Everybody Knows Your (Screen) Name: Online Games as "Third Places". *Journal of Computer-Mediated Communication* 11, 4 (2006), 885–909. DOI:10.1111/j.1083-6101.2006.00300.x.

[15] Dean Cocking and Jeanette Kennett. 1998. Friendship and the Self. *Ethics* 108, 3 (1998), 502–527.

[16] Kimberley Brownlee. 2013. A Human Right Against Social Deprivation. *The Philosophical Quarterly* 63, 251 (2013), 199–222. DOI:10.1111/1467-9213.12018.

[17] Yidan Yin, Nan Jia, and Cheryl J. Wakslak. 2024. AI Can Help People Feel Heard, but an AI Label Diminishes This Impact. *PNAS* 121, 14 (2024), e2319112121. DOI:10.1073/pnas.2319112121.

[18] Pranav Khadpe, Kimi Wenzel, George Loewenstein, and Geoff Kaufman. 2025. Explaining the Reputational Risks of AI-Mediated Communication: Messages Labeled as AI-Assisted Are Viewed as Less Diagnostic of the Sender's Moral Character. *Proc. AAAI/ACM Conf. on AI, Ethics, and Society (AIES)* 8, 2 (2025), 1413–1423. DOI:10.1609/aies.v8i2.36641.

[19] Peng Zhang. 2026. Conversation with Elys Founder: His 10 Product Insights and the Next-Generation Social Network He Wants to Build. Founder Park / GeekPark, 15 February 2026. English translation: TechFlow. `https://www.techflowpost.com/en-US/article/30379`.

[20] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. DOI:10.1191/1478088706qp063oa.