

Beyond Authenticity: GenAI Personas as Affective Diplomats in Long Distance Families*

Qing Li^{1,2,*}, Yufeng Wu¹, Abdulrhman Alorini¹, Mohammed Althubiani¹,
Ezwan Shah Abd Majid¹, Weijie Wang¹, Shang Chieh Lee³, Raj Mahmud¹, Robert P. Spang^{4,2},
Sebastian Möller^{2,5}, Jaime Garcia¹ and A Baki Kocaballi¹

¹Faculty of Engineering and Information Technology, University of Technology Sydney, Sydney, New South Wales, Australia

²Quality and Usability Lab, Technische Universität Berlin, Berlin, Germany

³University of Technology Sydney, Sydney, New South Wales, Australia

⁴Usability Lab, Bauhaus-Universität Weimar, Weimar, Germany

⁵German Research Center for Artificial Intelligence (DFKI), Berlin, Germany

Abstract

GenAI personas are emerging as conversational proxies that can speak in the first person on behalf of individuals. While such systems have been used, little is known about their role in intimate family relationships where care, vulnerability, and responsibility are central. This position paper argues that GenAI personas for long-distance families should be viewed as systems that support relational and emotional coordination, rather than as mechanisms for producing authentic self-representation. Drawing on a first person persona probe modeled on the lead author, we articulate tensions among curated care and transparency, delegated agency and control, and relational presence and authenticity. This work outlines design directions focused on ongoing consent, intentional imperfection, boundary setting, and protection of the represented subject. GenAI personas could support connection across distance without displacing human contact.

Keywords

GenAI Persona, Long Distance Families, Delegated Presence, Affective Diplomacy, AI Ethics

1. Introduction

Long-distance family relationships are increasingly shaped by migration, education, and global mobility. Maintaining connection across distance often relies on communication technologies such as Zoom and FaceTime, yet these systems primarily support synchronous interaction and assume symmetry in schedules, technological literacy, and emotional availability [1, 2]. In practice, family communication is frequently shaped by temporal, emotional, and cultural asymmetries that make ongoing connection fragmented and effortful [2]. Existing communication tools often reduce connection to scheduled exchanges rather than sustaining a continuous sense of relational presence.

Recent advances in Generative Artificial Intelligence (GenAI) introduce new possibilities for delegated presence in long-distance family life. Large Language Models can approximate individual communication styles and sustain extended social interaction [3], while conversational avatars can speak in the first

ACM CHI'26 Workshop on Responsible Use of AI Personas in Human-Centered Design and Research, April 13, 2026, Barcelona, Spain. © 2026 Copyright held by the owner/author(s).

*Corresponding author.

✉ Qing.Li-6@student.uts.edu.au (Q. Li); Yufeng.Wu-2@student.uts.edu.au (Y. Wu);
AbdulrhmanFahadA.F.Alorini@student.uts.edu.au (A. Alorini); Mohammed.S.Althubiani@student.uts.edu.au
(M. Althubiani); EzwanShah.AbdMajid@student.uts.edu.au (E. S. A. Majid); Weijie.Wang-3@uts.edu.au (W. Wang);
ShangChieh.Lee@uts.edu.au (S. C. Lee); Chaklader.R.Mahmud@student.uts.edu.au (R. Mahmud);
robert.spang@uni-weimar.de (R. P. Spang); sebastian.moeller@tu-berlin.de (S. Möller); Jaime.Garcia@uts.edu.au (J. Garcia);
Baki.Kocaballi@uts.edu.au (A. B. Kocaballi)

🆔 0009-0005-7519-5118 (Q. Li); 0009-0005-6537-7161 (Y. Wu); 0009-0007-9517-2970 (A. Alorini); 0009-0009-1366-0802
(M. Althubiani); 0000-0002-0609-0524 (E. S. A. Majid); 0000-0002-3587-8222 (W. Wang); 0009-0004-5047-106X (S. C. Lee);
0000-0001-6580-9060 (R. P. Spang); 0000-0003-3057-0760 (S. Möller); 0000-0001-5718-1605 (J. Garcia); 0000-0002-8328-5317
(A. B. Kocaballi)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

person using synthesized voices and photorealistic representations. Rather than simply transmitting messages, these systems can perform a version of the self through ongoing interaction. Prior work has explored AI-mediated proxies and embodied agents in workplace and productivity contexts [4, 5], yet their role in intimate family relationships remains underexplored.

Family communication, however, is not primarily organized around informational accuracy or authentic self-representation. Research on mediated relationships suggests that maintaining connection depends on continuity, care, and relational appropriateness rather than information exchange alone [2]. Studies of relational agents further show that interactive systems can participate in emotional management and politeness work [6, 7]. At the same time, identity in mediated communication is often selective, negotiated, and context dependent rather than a stable self that can be perfectly replicated. Delegating communication to an AI persona therefore raises questions not only about authenticity, but also about emotional buffering, accountability, and care.

Emerging work on AI-mediated representation highlights similar tensions. AI proxies may generate responses that users did not explicitly author [4], while increasingly autonomous systems blur distinctions between representation and delegation [8]. In intimate family contexts, these tensions become particularly sensitive because communication is closely tied to obligation, vulnerability, and emotional protection. Prior work on benevolent deception suggests that withholding information may sometimes be experienced as care rather than manipulation [9]. Similarly, AI personas may be intentionally curated to soften tone, preserve dignity, or shield loved ones from worry.

In this paper, we argue that GenAI personas for long-distance families should be understood as *affective diplomats* rather than authenticity machines. Their value lies less in reproducing the real self than in mediating emotional asymmetries and sustaining relational continuity across absence. Grounded in a Research through Design exploration [10, 11], we developed a first-person GenAI persona probe to examine how AI-mediated representation reshapes experiences of care, presence, and delegated agency in long-distance family relationships. Through this probe, we surface three tensions: (1) curated care and transparency, where families may prefer emotionally filtered communication over strict disclosure; (2) delegated agency and control, where AI-generated speech complicates authorship and accountability; and (3) relational presence and authenticity, where feelings of connection may matter more than faithful self-representation. We conclude by discussing implications for the governance and design of GenAI personas centered on consent, boundary setting, and protection of the represented subject.

2. Related Work

2.1. Long-Distance Family Communication

Research on long-distance family relationships has shown that maintaining connection across distance involves more than information exchange alone. Relationships are often shaped by temporal separation, uneven availability, and asymmetries in emotional expression and technological access [2]. Existing communication platforms such as video calling systems primarily support synchronous interaction, requiring participants to coordinate schedules and sustain reciprocal attention in real time [1]. While these systems are effective for direct conversation, they often fit poorly with the asynchronous and fragmented nature of long-distance family life.

Prior HCI research has explored a range of technologies for supporting distributed families, including always-on communication systems, mediated domestic presence, and conversational technologies [1, 10, 12]. Studies suggest that relational continuity, the ongoing feeling of connection between family members, is difficult to sustain through episodic interaction alone [2]. Communication in long-distance families is also closely tied to care practices, where emotional support may be expressed indirectly through coordination, politeness, or protective filtering rather than explicit disclosure. These findings suggest a need for systems that support delegated and emotionally mediated forms of presence beyond real-time communication.

2.2. AI Personas and Delegated Presence

Recent advances in Generative AI have enabled increasingly sophisticated forms of AI-mediated representation. Large Language Models can simulate conversational styles and sustain social interaction over extended exchanges [3]. Building on these capabilities, recent systems have explored AI proxies and embodied conversational agents that can participate in interaction on behalf of users. For example, Dittos [4] investigated personalized embodied agents that attend workplace meetings during a user's absence, while D-Twins [5] explored AI-generated digital twins designed for emotional intervention.

These systems extend traditional communication tools by introducing delegated agency, where interaction is partially generated by the system rather than directly authored by the user. Prior research suggests that such delegation complicates questions of authorship, accountability, and identity representation [4, 8]. Rather than functioning as passive communication channels, AI personas actively shape interaction through generated responses, conversational tone, and social behavior. While existing work has primarily examined these systems in workplace, productivity, or metaverse settings [4, 5, 8], less is known about how delegated AI personas operate within intimate family relationships where care, obligation, and vulnerability are central.

2.3. Authenticity, Care, and Relational Mediation

Research on relational agents suggests that conversational systems can participate in emotional and social mediation rather than simply delivering information. Relational agents have been shown to support long-term interaction through politeness, emotional management, and socially appropriate communication strategies [6]. Similarly, politeness theory highlights how people selectively manage expression to preserve dignity, reduce conflict, and maintain interpersonal relationships [7].

These perspectives complicate assumptions that mediated systems should prioritize authenticity or exact self-representation. Prior work on benevolent deception argues that withholding or selectively presenting information may sometimes function as a form of care rather than manipulation [9]. In family contexts, emotionally filtered communication may therefore be experienced as supportive and protective, particularly when users seek to shield loved ones from distress or maintain relational harmony.

Emerging AI personas intensify these tensions because they can speak in the first person and sustain interaction independently of the represented individual. This raises broader questions about how authenticity, emotional care, and accountability should be negotiated when conversational agency is delegated to AI systems. Our work builds on these discussions by examining GenAI personas as relational mediators in long-distance family relationships rather than as systems primarily concerned with accurate self-representation.

3. Method

3.1. First-Person GenAI Persona Probe

To explore how GenAI personas may mediate care, presence, and delegated communication in long-distance family relationships, we developed a first-person GenAI persona probe grounded in a Research through Design (RtD) approach [10, 11]. Rather than evaluating system performance or usability, the probe was designed to surface relational and ethical tensions surrounding AI-mediated representation in intimate family contexts.

The probe consisted of a photorealistic conversational avatar created using the HeyGen interactive avatar platform (Figure 1).¹ The system supported real-time voice and text interaction through a first-person conversational interface. The avatar was modeled on the lead author, including synthesized appearance and voice characteristics, allowing participants to interact with the system as a persistent family proxy. Consistent with prior work on AI-mediated agents and simulated social interaction [3, 4],

¹<https://www.heygen.com/interactive-avatar>




Figure 1: GenAI persona probe based on the lead author shown from two views: (a1) a participant engaging in conversational interaction with the persona, and (a2) the HeyGen interface illustrating the conversational dialogue. Participants reflected on how such a persona might be used in their own long-distance family contexts.

the probe was configured to approximate the communication style and conversational tone of the represented individual.

Our approach adopted a first-person and collaborative autoethnographic orientation, where researchers engaged reflexively with the probe as both designers and participants [11]. This stance was appropriate because AI personas in family contexts involve sensitive questions of identity, emotional vulnerability, and delegated agency that are difficult to explore through conventional evaluative approaches alone.

3.2. Probe Enactment

Five researcher-participants with lived experience of maintaining long-distance family relationships participated in one-on-one enactment sessions. Participants came from diverse cultural and geographic backgrounds and reflected on their own experiences of distance, care, and family communication during the study.

Each session lasted approximately 30 minutes and involved conversational interaction with the GenAI persona followed by a semi-structured reflective interview. Participants were invited to engage with the avatar as though it were a family proxy and to imagine how such a system might operate within their own long-distance family relationships. The sessions focused on exploring situations involving emotional support, everyday communication, absence, and difficult conversations.

The probe was intentionally treated not as a finalized product but as a speculative relational artifact intended to provoke reflection on mediated presence, authenticity, emotional buffering, and delegated communication. This orientation aligns with RtD approaches that use experiential enactment to surface emerging social and ethical implications of new technologies [10].

3.3. Analysis

Session recordings, reflective discussions, and interview transcripts were analyzed using reflexive thematic analysis [13]. Our analysis focused on identifying recurring tensions in how participants interpreted and negotiated AI-mediated representation in family relationships. Rather than seeking consensus or generalizable behavioral findings, the analysis aimed to articulate design-relevant tensions surrounding curated care, delegated agency, accountability, and relational presence.

Through iterative discussion among the research team, we synthesized these tensions into conceptual themes that informed our framing of GenAI personas as affective diplomats rather than authenticity machines.

4. Findings: Tensions of GenAI Personas

Our analysis surfaced three tensions surrounding the use of GenAI personas in long-distance family relationships: (1) curated care and transparency, (2) delegated agency and control, and (3) relational

presence and authenticity. These tensions illustrate how AI personas operate not simply as communication tools, but as relational mediators that reshape how care, responsibility, and presence are negotiated across distance.

4.1. Curated Care and Transparency

Participants frequently imagined the GenAI persona as a mechanism for emotionally filtered communication rather than accurate self-representation. Rather than valuing transparency alone, participants preferred forms of curation that could soften tone, reduce worry, and preserve relational harmony. One participant described wanting the persona to “report good news and hide bad news” in order to protect parents from unnecessary distress. Similar forms of selective disclosure appeared in discussions around difficult family relationships, where participants envisioned the persona acting as an emotional buffer capable of making communication gentler and less confrontational.

These reflections align with prior work on relational agents and politeness, which suggests that mediated systems can participate in emotional management and face-saving practices [6, 7]. Participants often framed emotionally filtered communication not as deception, but as care. In some cases, participants discussed intentionally concealing the artificial nature of the persona from family members if disclosure might create anxiety or emotional insecurity. This echoes notions of benevolent deception in HCI, where withholding information may function as a protective relational strategy rather than manipulation [9].

At the same time, participants did not necessarily seek exact authenticity from the persona. Several described preferring an idealized or supportive version of themselves rather than a perfectly faithful reproduction. This suggests that in family contexts, emotionally appropriate representation may be valued more than accurate replication of the self.

4.2. Delegated Agency and Control

Participants valued the GenAI persona for its ability to sustain interaction during absence and reduce the burden of maintaining constant availability across distance. The persona was imagined as maintaining everyday conversation asynchronously, particularly in situations where time-zone differences, fatigue, or emotional exhaustion made direct communication difficult. In this sense, the persona functioned as a form of delegated continuity, allowing relationships to feel ongoing even when synchronous interaction was impossible.

However, delegation also complicated questions of authorship and accountability. Several participants expressed discomfort when the persona generated responses that they felt did not fully reflect their intentions or values. One participant described the system as “saying things I wouldn’t say,” producing uncertainty around responsibility for AI-generated speech. While participants understood the system as partially autonomous, they nevertheless felt that accountability remained socially attached to the represented individual rather than the system itself.

This tension became particularly visible in discussions around conversational boundaries. Participants were comfortable delegating lightweight emotional interaction and routine conversation to the persona, but expressed hesitation about allowing the system to engage in serious family matters, conflict, or emotionally sensitive decision making. Delegation was therefore viewed as conditional and situational rather than absolute.

Participants also reflected on the importance of maintaining visible traces of mediation within the interaction. Some argued that if the persona became too realistic or indistinguishable from the represented person, it could become emotionally unsafe or ethically ambiguous. Minor imperfections in conversational timing, expression, or behavior were sometimes interpreted as useful reminders that the interaction remained mediated rather than fully human.

4.3. Relational Presence and Authenticity

Participants consistently described interactions with the GenAI persona as socially meaningful despite knowing that the system was artificial. During sessions, participants often followed ordinary social conventions such as greeting the persona, expressing gratitude, and explaining why they needed to end the conversation. These behaviors suggest that the interaction was experienced less as tool use and more as participation in a social encounter.

Importantly, the feeling of presence emerged primarily through responsiveness and conversational continuity rather than visual realism or accurate self-representation. Participants did not necessarily expect the persona to function as a perfect replica. Instead, they valued its ability to sustain a feeling of ongoing connection across absence. In this sense, relational presence was produced through interaction itself rather than through fidelity alone.

Participants also emphasized that the appropriateness of AI-mediated presence depended heavily on social context. While asynchronous conversational companionship was viewed as potentially supportive in everyday family life, participants questioned whether AI personas would be appropriate in emotionally sensitive situations such as serious conflict, grief, or major family events. This suggests that the legitimacy of GenAI personas is socially negotiated and shaped by contextual expectations around care, authenticity, and respect.

Taken together, these findings suggest that GenAI personas in family relationships are not primarily evaluated according to realism or technical fidelity. Instead, participants interpreted them through relational concerns surrounding emotional safety, delegated responsibility, and the maintenance of care across distance.

5. Discussion

Our findings suggest that GenAI personas in long-distance family relationships should not be understood primarily as systems for authentic self-representation. Instead, participants consistently framed these personas as relational mediators that help manage emotional asymmetries, maintain continuity, and support care across absence. In this section, we discuss how these findings complicate assumptions surrounding authenticity, delegation, and realism in AI-mediated family communication.

5.1. Beyond Authenticity: GenAI Personas as Affective Diplomats

Much of the emerging discourse surrounding AI-generated personas and digital twins assumes that the primary design goal is accurate replication of the individual [8]. However, our findings suggest that family relationships operate according to a different relational logic. Participants often preferred emotionally filtered, softened, or idealized communication over exact authenticity. In this sense, the GenAI persona was valued less as a truthful replica and more as an affective diplomat capable of mediating emotional tension and preserving relational harmony.

This finding aligns with prior work showing that mediated communication frequently involves politeness, emotional buffering, and selective disclosure rather than complete transparency [7, 6]. Participants described situations where shielding loved ones from distress was perceived as an act of care, even when this involved withholding information or concealing the artificial nature of the persona. Such practices resemble benevolent deception [9], where selective mediation functions as emotional protection rather than manipulation.

Our findings therefore challenge assumptions that authenticity should serve as the primary evaluative framework for AI personas in intimate relationships. In long-distance family contexts, sustaining a feeling of connection and care may matter more than reproducing a perfectly accurate version of the self. Rather than functioning as authenticity machines, GenAI personas may operate as infrastructures of affective diplomacy that negotiate relational continuity across distance.

5.2. Imperfection, Boundaries, and Emotional Safety

Participants also revealed ambivalence toward highly realistic AI-mediated representation. While conversational responsiveness contributed to a feeling of social presence, several participants expressed discomfort with the possibility of personas becoming indistinguishable from the represented person. Interestingly, minor imperfections were sometimes interpreted positively because they maintained awareness that the interaction remained mediated.

This finding complicates dominant assumptions in AI embodiment research that greater realism necessarily improves social interaction. Instead, visible signs of mediation may support emotional safety, accountability, and interpretive clarity. Participants appeared to prefer systems that remain socially meaningful without fully collapsing distinctions between the human and the proxy.

Boundary setting also emerged as important for managing delegated agency. Participants were generally comfortable allowing personas to support lightweight emotional interaction or routine communication, but were reluctant to delegate serious family discussions, conflict, or emotionally consequential decisions. This suggests that delegation in intimate relationships is not binary but negotiated and situational.

These findings point toward the importance of designing GenAI personas with revisable consent, constrained autonomy, and explicit boundary mechanisms. Rather than maximizing realism or autonomy alone, future systems may need to support ongoing negotiation between emotional usefulness and relational accountability.

5.3. Designing for Relational Appropriateness

Finally, our findings suggest that the success of GenAI personas depends less on technical fidelity than on relational appropriateness. Participants evaluated the persona according to whether its presence felt socially and emotionally suitable within particular family contexts. Everyday companionship and asynchronous emotional support were often viewed positively, while situations involving grief, conflict, or major family events were seen as more ethically sensitive.

This reinforces prior HCI research suggesting that technologies participate in socially situated practices rather than functioning as neutral communication channels [2]. In family contexts, AI personas become entangled with cultural expectations surrounding care, obligation, and interpersonal responsibility. Their legitimacy is therefore negotiated socially rather than determined solely through technical capability.

In conclusion, our findings suggest that future GenAI personas for families should be designed not simply around realism, but around relational governance. This includes supporting contextual appropriateness, preserving emotional safety, enabling user control over delegation, and protecting the represented subject from unwanted or misleading forms of representation.

6. Conclusion

GenAI personas introduce new possibilities for maintaining connection in long-distance family relationships. Rather than functioning simply as tools for communication or accurate self-representation, our findings suggest that these systems may operate as relational mediators that help sustain care, continuity, and emotional presence across absence. Through a first-person GenAI persona probe, we surfaced tensions surrounding curated care and transparency, delegated agency and accountability, and relational presence and authenticity.

Our work contributes to GenAI personas as *affective diplomats* rather than authenticity machines. Participants often valued emotional appropriateness, continuity, and protective mediation over exact realism, suggesting that AI-mediated family communication operates according to relational logics distinct from those found in workplace or productivity settings. At the same time, the findings highlight important concerns around authorship, emotional safety, boundary setting, and the governance of delegated representation.

As GenAI personas become increasingly capable of speaking and interacting on behalf of individuals, future research should move beyond questions of technical fidelity alone and examine how these systems reshape care, responsibility, and relational norms in intimate contexts. Designing for long-distance families may therefore require approaches centered not only on realism and autonomy, but also on consent, contextual appropriateness, and protection of the represented subject.

7. Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to perform grammar and spelling checks and to improve the clarity of the writing. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] Y. Gan, C. Greiffenhagen, S. Reeves, Connecting distributed families: Camera work for three-party mobile video calls, in: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–12.
- [2] A. J. Merolla, Relational maintenance and noncopresence reconsidered: Conceptualizing geographic separation in close relationships, *Communication Theory* 20 (2010) 169–193.
- [3] J. S. Park, C. Q. Zou, A. Shaw, B. M. Hill, C. Cai, M. R. Morris, R. Willer, P. Liang, M. S. Bernstein, Generative agent simulations of 1,000 people, *arXiv preprint arXiv:2411.10109* (2024).
- [4] J. Leong, J. Tang, E. Cutrell, S. Junuzovic, G. P. Baribault, K. Inkpen, Dittos: Personalized, embodied agents that participate in meetings when you are unavailable, *Proceedings of the ACM on Human-Computer Interaction* 8 (2024) 1–28.
- [5] I.-C. Lo, P.-L. P. Rau, D-twins: Your digital twin designed for real-time boredom intervention, in: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–15.
- [6] T. W. Bickmore, R. W. Picard, Establishing and maintaining long-term human-computer relationships, *ACM Transactions on Computer-Human Interaction (TOCHI)* 12 (2005) 293–327.
- [7] P. Brown, S. C. Levinson, *Politeness: Some universals in language usage*, volume 4, Cambridge university press, 1987.
- [8] V. Moravec, B. Gavurova, M. Rigelsky, “what is the latest news, avatar pavel?”-ai assistants in transformation processes of metaverse, *Computers in Human Behavior: Artificial Humans* (2025) 100225.
- [9] E. Adar, D. S. Tan, J. Teevan, Benevolent deception in human computer interaction, in: *Proceedings of the SIGCHI conference on human factors in computing systems*, 2013, pp. 1863–1872.
- [10] W. Gaver, F. Gaver, Living with light touch: An autoethnography of a simple communication device in long-term use, in: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–14.
- [11] B. Ploderer, T. Capel, N.-E. Davaakhuu, N. H. Tung, D. M. Maichal, A. K. Kuzhiparambil, Q. H. Mai, W. Reitberger, Maintaining long-distance relationships with (mediocre) llm-based chatbots: A collaborative ethnographic study, in: *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–10.
- [12] Q. Li, A. B. Kocaballi, J. Garcia, Human digital twin for long-distance relationships: a scoping review, *Behaviour & Information Technology* (2025) 1–23.
- [13] K. A. Neuendorf, Content analysis and thematic analysis, in: *Advanced research methods for applied psychology*, Routledge, 2018, pp. 211–223.