

Dittos: Mimetic, Reciprocal Agents in AI Mediated Communication

EDWARD CUTRELL, Microsoft Research

JOHN TANG, Microsoft Research

SASA JUNUZOVIC, Microsoft Research

MARTEZ MOTT, Microsoft Research

DENAE FORD, Microsoft Research

KORI INKPEN, Microsoft Research

Recent advances in generative AI have enabled agents that can represent specific people in social interactions when they are unavailable. These agents—often described as digital twins—can look and sound like an individual and participate in conversations on their behalf. In this position paper, we focus on a particular class of such systems that we refer to as Dittos: mimetic, reciprocal AI agents that not only interact with others as a proxy for a person but also report back what occurred so that human relationships can continue to develop over time. Drawing on our experiences designing, deploying, and studying Dittos, we argue that this class of systems surfaces a distinct set of challenges to core assumptions in AI mediated communication research. These challenges concern how presence is evoked through representation, how trust is established when AI speaks in someone’s voice, how people remain meaningfully informed about Ditto mediated interactions, and how ethical risks can be anticipated and mitigated. We position these challenges as central to understanding a new phase of AI mediated communication, where systems increasingly act *on behalf of* people rather than merely between them.

CCS CONCEPTS • Human-centered computing → Human computer interaction (HCI) Interaction paradigms

Additional Keywords and Phrases: digital twins, mimetic agent, trust, human-AI interaction, embodied agents, mediated communication

1 INTRODUCTION

AI mediated communication research within CHI and CSCW has historically focused on systems that support or structure interaction between people, such as messaging, meetings, and collaborative work tools. With the emergence of generative AI, systems are increasingly able to participate in communication as autonomous actors rather than remaining infrastructural; AI becomes an intermediary in communication between people.

Recent public demonstrations of AI “digital twins,” including conversational and embodied representations of real individuals [2, 9], have made this shift visible while also raising concerns about authenticity, trust, and responsibility. Our experience in designing, deploying and studying Dittos—AI agents that explicitly represent a specific person in social interaction—introduce qualitatively new dynamics to mediated communication.

We take the position that Dittos represent a distinct form of AI mediated communication between people. Dittos are not intended to replace people, but to extend their availability across time and context. As such, they foreground research challenges related to identity, agency, social context, and relational continuity.

2 DITTOS AS A DISTINCT DESIGN SPACE

We characterize Dittos by three defining properties.

First, Dittos are mimetic: they are designed to evoke a specific person (we call this person the Source) through voice, language use, visual appearance, interaction style, and knowledge. Prior work on mimetic models highlights the ethical and social implications of AI systems that act “like you,” particularly when grounded in personal data and behavioral traces [8].

Second, Dittos are socially situated. They interact with collaborators who typically already have an established relationship with the Source. Work on personalized embodied agents in meetings shows that even limited participation by such agents can alter conversational dynamics and perceptions of presence [7].

Third, Dittos are reciprocal. Interactions with a Ditto are not an endpoint; instead, the Source is kept informed about what occurred, enabling follow up, repair, and continuation of the relationship, as well as informing modifications to the Ditto’s behavior. This connects Dittos to long standing CSCW research on meeting capture and sensemaking [11, 12, 13], but with the added complexity that the interaction was conducted in one’s name.

Together, these properties distinguish Dittos from generic conversational agents and create rich new possibilities and complex challenges for collaboration and communication.

3 KEY RESEARCH CHALLENGES

Over the past three years, we’ve built, deployed and studied Dittos in a variety of contexts, including meetings [7], informal 1:1 interactions [6], and family relationships [10]. Most recently, we explored a deployment in which a large group of student interns interacted with the Dittos of two senior mentors. Together, these experiences uncovered several important considerations and research challenges.

3.1 Evoking Presence Through Ditto Representation

A central challenge concerns how Dittos should represent the people they stand in for. We found that vocal characteristics and linguistic style can evoke a sense of presence more strongly than visual fidelity alone, even when visual representations are deliberately stylized [7]. Subtle features such as vocabulary choice, cadence, and dysfluencies play an important role in how collaborators experience a Ditto.

This creates a genuine design dilemma: realistic representations are overwhelmingly preferred by both Sources and collaborators, yet the more convincing a Ditto becomes, the more ethically fraught its use. We argue that the field should resist the instinct to resolve this tension through disclosure mechanisms alone. If a Ditto is effective precisely because people forget it is one, then mandatory disclosure undermines the very social function it is designed to serve. This is not a call to deceive—it is a call to acknowledge that Dittos occupy an uncomfortable middle ground that existing frameworks for transparency do not adequately address.

3.2 Trust When a Ditto Speaks for Someone

Trust is foundational for Dittos but operates at multiple levels. Sources must trust that the Ditto will act consistently with their intentions and values, while collaborators must trust that what a Ditto says is meaningful and will be conveyed accurately back to the Source.

Prior work on AI clones and digital twins raises concerns about identity fragmentation and misplaced trust when AI systems speak in a person’s voice [5]. Notably, collaborators may place greater trust in a Ditto than in a generic assistant, even when both are generated by the same underlying AI system [7]. This suggests that trust in Ditto mediated communication is deeply entangled with identity and accountability. Dittos do not just inherit trust; they manufacture it through identity and may carry more persuasive weight than the underlying AI has earned.

3.3 Making Ditto Mediated Interaction Legible

Reciprocity requires that people remain meaningfully informed about interactions conducted by their Ditto. However, reviewing full recordings of Ditto mediated conversations is often impractical, and generic AI summaries may omit socially salient details such as facial expressions, gestures and vocal expressivity. Reciprocity is not binary; it spans a spectrum from grief bots, which can never close the loop, to telepresence robots, which approximate full co-presence. Dittos occupy a middle ground where the loop is partially closed: Sources can review interactions and update their Ditto, but current recaps may omit details that would be apparent in direct interaction.

This challenge builds on prior work in meeting understanding and summarization [11, 12, 13], but extends it in an important way: the goal is not simply awareness, but the ability to re-enter a relationship and continue it seamlessly after a Ditto interaction. For example, in our intern deployment, mentors would often use topics discussed with their Dittos as starting points for direct conversations with interns.

Further, these interaction summaries or recaps are used by the Source to update the context or knowledge for future interactions. For example, one mentor observed similar questions from different interns and updated his Ditto with an extensive FAQ that was then available to everyone.

3.4 Ethical and Responsible Design of Dittos

Finally, Dittos raise significant ethical concerns around consent, disclosure, security, and misuse. Prior work highlights risks such as doppelganger anxiety, misrepresentation, and unintended delegation [5, 8]. Because many risks may only emerge through long term social use, speculative and anticipatory approaches—such as structured “dark Ditto” or Black Mirror inspired exercises—offer promising methods for proactively identifying harms and informing responsible design [3, 1]. And while such exercises can be helpful in revealing potential risks, three years of deployment have taught us that such risks with Dittos are not hypothetical, they are operational. We have observed Dittos confidently stating things the Source would never say, fabricate memories, act as a confidant, and derail meetings all in the relative safety of our lab.

4 IMPLICATIONS FOR CHI AND CSCW

Dittos illustrate a broader shift in AI mediated communication: systems are increasingly acting on behalf of people rather than merely between them. This shift challenges assumptions about agency, responsibility, and interaction boundaries that underlie much CHI and CSCW research. Major platforms are already exploring digital twin capabilities [9], and the underlying technology is advancing faster than our frameworks for reasoning about it. The question is not whether people will delegate social interaction to AI—it is whether we will have developed adequate norms for accountability, consent, and relational repair when they do. We think this is an important consideration in navigating a responsible design process for developing digital twins. We offer this position paper as a call for the CHI and CSCW communities to engage with Dittos as a concrete lens on the future of AI mediated communication, helping to shape design principles, research methods, and ethical norms before such systems become commonplace.

REFERENCES

- [1] Jed R. Brubaker, Casey Fiesler, Michael Madaio, John Tang, Richmond Y. Wong. 2024. Generative AI Going Awry: Enabling Designers to Proactively Avoid It in CSCW Applications, *CSCW Companion '24: 2024 Conference on Computer-Supported Cooperative Work and Social Computing*, 125 – 127. <https://dl.acm.org/doi/abs/10.1145/3678884.3689133>
- [2] Reid Hoffman. 2024. Reid Hoffman meets his AI twin - Full. YouTube. April 24, 2024. <https://youtu.be/rgD2gmwCS10?si=DqNn8BFerDHPHSuO>
- [3] Shamika Klassen & Casey Fiesler. 2022. “Run Wild a Little With Your Imagination”: Ethical Speculation in Computing Education with Black Mirror, *SIGCSE 2022: Proceedings of the 53rd ACM Technical Symposium on Computer Science Education*, February 2022, pp. 836 – 842. <https://doi.org/10.1145/3478431.3499308>
- [4] Patrick Yung Kang Lee, Ning F. Ma, Ig-Jae Kim, and Dongwook Yoon. 2023. Speculating on Risks of AI Clones to Selfhood and Relationships: Doppelgänger-phobia, Identity Fragmentation, and Living Memories. *Proc. of the ACM Human-Computer Interaction* 7, CSCW1 (2023), 1–28. <https://doi.org/10.1145/3579524>
- [5] Patrick Yung Kang Lee, Ning F. Ma, Ig-Jae Kim, and Dongwook Yoon. 2023. Speculating on Risks of AI Clones to Selfhood and Relationships: Doppelgänger-phobia, Identity Fragmentation, and Living Memories. *Proc. of the ACM Human-Computer Interaction* 7, CSCW1 (2023), 1–28. <https://doi.org/10.1145/3579524>
- [6] Seonghee Lee, Denae Ford, John Tang, Sasa Junuzovic, Asta Roseway, Ed Cutrell, Kori Inkpen. 2025. IRL Dittos: Embodied Multimodal AI Agent Interactions in Open Spaces. <https://arxiv.org/abs/2504.21347>
- [7] Joanne Leong, John Tang, Edward Cutrell, Sasa Junuzovic, Gregory Paul Baribault, Kori Inkpen. 2024. Dittos: Personalized, Embodied Agents That Participate in Meetings When You Are Unavailable. *Proc. ACM Human-Computer Interaction* 8, CSCW2, Article 494 (November 2024), 28 pages. <https://doi.org/10.1145/3687033>
- [8] Reid Mellroy-Young, Jon Kleinberg, Siddhartha Sen, Solon Barocas, and Ashton Anderson. 2022. Mimetic Models: Ethical Implications of AI that Acts Like You. In *5th AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, 479–490. <https://doi.org/10.1145/3514094.3534177>
- [9] Nilay Patel. 2024. The CEO of Zoom wants AI clones in meetings, *The Verge*, June 3, 2024, <https://www.theverge.com/2024/6/3/24168733/zoom-ceo-ai-clones-digital-twins-videoconferencing-decoder-interview>
- [10] Teerapaun Tanprasert, Jiamin Dai, John Tang, Kori Inkpen, Edward Cutrell, and Joanna McGrenere. 2025. FamilyDittos: Reimagining Intergenerational Interaction through Mimetic Agents. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW419 (November 2025), 31 pages. <https://doi.org/10.1145/3757600>
- [11] Simon Tucker and Steve Whittaker. 2005. Accessing Multimodal Meeting Data: Systems, Problems and Possibilities. In *Machine Learning for Multimodal Interaction (Lecture Notes in Computer Science)*, Samy Bengio and Hervé Bourlard (Eds.). Springer, Berlin, Heidelberg, 1–11. https://doi.org/10.1007/978-3-540-30568-2_1
- [12] Ruotong Wang, Lin Qiu, Justin Cranshaw, Amy Zhang. 2024. Meeting Bridges: Designing Information Artifacts that Bridge from Synchronous Meetings to Asynchronous Collaboration, *Proc. ACM Human-Computer Interaction* 8, CSCW2, Article 35 (November 2024), 29 pages <https://dl.acm.org/doi/10.1145/3637312>
- [13] Steve Whittaker, Rachel Laban, and Simon Tucker. 2006. Analysing Meeting Records: An Ethnographic Study and Technological Implications. In *Machine Learning for Multimodal Interaction (Lecture Notes in Computer Science)*, Steve Renals and Samy Bengio (Eds.). Springer, Berlin, Heidelberg, 101–113. https://doi.org/10.1007/11677482_9
- [14] Sicheng Xu, Guojun Chen, Yu-Xiao Guo, Jiaolong Yang, Chong Li, Zhenyu Zang, Yizhong Zhang, Xin Tong, and Baining Guo. 2024. VASA-1: Lifelike Audio-Driven Talking Faces Generated in Real Time. *NeurIPS Oral Poster*, <https://arxiv.org/abs/2404.10667>