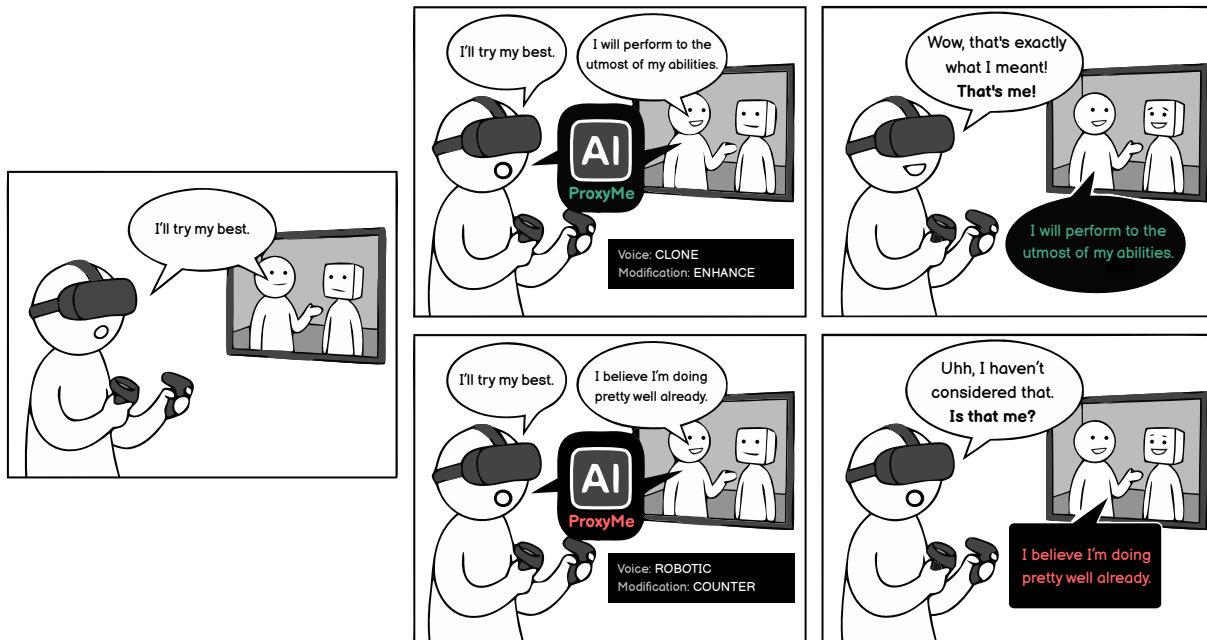


# Is it Me? Toward Self-Extension to AI Avatars in Virtual Reality

Jieying Zhang  
University of Amsterdam  
Amsterdam, Netherlands  
Centrum Wiskunde & Informatica  
(CWI)  
Amsterdam, Netherlands  
holly.jy.zhang@outlook.com

Steeven Villa  
LMU Munich  
Munich, Germany  
villa@posthci.com

Abdallah El Ali  
Centrum Wiskunde & Informatica  
(CWI)  
Amsterdam, Netherlands  
Utrecht University  
Utrecht, Netherlands  
aea@cwi.nl



User speaks naturally, controlling their avatar to interact with a virtual agent.

ProxyMe can reformulate the user's speech and voice during interaction.

User experiences enhanced / modified communication as self-extension.

Figure 1: ProxyMe interaction concept.

## Abstract

Advances in generative AI, speech synthesis, and embodied avatars enable systems that not only assist communication, but can act as proxies on users' behalf. Prior work in HCI has largely focused on systems as external tools, with less attention paid to the experiential consequences of users' speech and actions becoming assimilated with AI-generated output. We introduce the design and implementation of ProxyMe, a work-in-progress VR prototype that allows users to embody an avatar whose voice and spoken content are modified by an AI system. By combining avatar-based embodiment, voice cloning, and AI-mediated speech augmentation, ProxyMe invites the exploration of avatar self-extension: situations in which

AI-modified communication is experienced as part of one's own expressive behavior. We chart out research challenges and envisioned scenarios, with a focus on how varying degrees of delegation and steerability can influence perceived agency, authorship, and self-identification.

## CCS Concepts

• Human-centered computing → Virtual reality; Human computer interaction (HCI).

## Keywords

Self-extension, avatar, human augmentation, AI-mediated communication, voice cloning

## ACM Reference Format:

Jieying Zhang, Steeven Villa, and Abdallah El Ali. 2026. Is it Me? Toward Self-Extension to AI Avatars in Virtual Reality. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3772363.3798899>



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI EA '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2281-3/26/04

<https://doi.org/10.1145/3772363.3798899>

## 1 Introduction

Science fiction narratives have envisioned the possibility of humans that are no longer confined to their physical body, allowing them the ability to extend their cognitive and bodily self onto external substrates. *Ghost in the Shell* [2] depicts *cyberbrains* and *prosthetic bodies (shells)* as mechanisms through which cognition and action are distributed beyond the biological organism, which raises questions about the continuity of identity and the attribution of actions to the self. By contrast, *Ergo Proxy* [1] portrays “*proxies*” as substitutive agents that act in place of human subjects, often preceding conscious intention or endorsement. Rather than extending human agency transparently, these agents displace authorship and decision-making, which ultimately expose the fragility of human agency and responsibility under such mediated contexts.

Recent advances in generative models, 3D avatars, and speech synthesis bring these aforementioned science fiction visions into experimentally tractable forms. This is especially due to such capabilities embedded into real-time and multi-modal communication contexts such as voice calls, video meetings, and immersive VR environments [55, 57]. Furthermore, contemporary AI agents are capable of generating persuasive arguments [47], supporting reasoning and decision-making on users’ behalf [13, 14, 19], or allowing fine-grained vocal control or even full voice cloning [24]. Such developments give rise to applications in which users may co-speak with, or be represented by, AI systems, blurring boundaries between self and simulation, and as far as authenticity and algorithmic generation [32, 60]. Such systems can redistribute aspects of users’ agency and autonomy [13, 58], which suggest a form of self-extension (cf., Clark’s Extended Mind [15]), wherein technologies become tools (cf., [6, 33]) that are themselves part and parcel of the user’s cognitive processes, reasoning patterns, and even vocal expressions. As such, while AI systems increasingly modify how users speak or express ideas through LLM-assisted writing [34], research typically focuses on users interacting with AI agents as external tools. There has been less attention on the experiential consequences of users’ own behavior becoming assimilated with intelligent output, forming a closed cognitive or sensorimotor loop. This latter perspective aligns more closely with visions of human–computer integration [36], by which computational systems become interwoven with perception and action. Together, these developments invite exploring avatar-based self-extension, where AI-mediated modifications to speech and action are experienced as part of one’s own communicative behavior.

In this work, we present the early design and implementation of **ProxyMe**, a VR prototype that lets users embody a virtual avatar and experience how an AI system modifies the avatar’s voice and spoken content on their behalf (Figure 1). By combining avatar-based embodiment, voice modulation, and speech augmentation, ProxyMe seeks to explore the boundary between self-authored and AI-mediated communication. Conceptually, ProxyMe relates to multiple aspects of the JIZAI Body framework [29], namely Fusionary, Duplicatory, and Possessory/Transformatory modes of being. Through this configuration of AI-augmented avatar embodiment (i.e., self-extension), ProxyMe enables the exploration of how such systems influence users’ perceived agency, authorship, and self-identification with their avatar(s) and its actions. As external

reasoning systems increasingly act on users’ behalf in communication [13], they do far more than supporting task completion and productivity (cf., [17]): When users reflect on their self-extended actions and speech, such systems can implicitly shape how users express opinions, make decisions, and form values. This is relevant as communicative actions are partially delegated to external systems where AI safety and human value alignment is paramount [26, 50]. At the same time, users may experience shifts in how they perceive authorship and agency [18, 36]. Our work envisions multiple scenarios where such novel interactions can benefit, including addressing speech impairment [12], public speaking [5], mind-body therapy [21], and self-understanding and (virtual) identity exploration [4]. Understanding these experiential dynamics is therefore critical for designing AI-mediated communication systems that balance delegation with user agency, and can support responsible attribution of action and ownership.

## 2 Related Work

### 2.1 Multimodal self-extension in VR

Clark’s Extended Mind theory proposes that external resources such as notebooks and smartphones can be recognized as parts of mental processes, and thus as extensions of the biological brain, when they are integrated with cognitive tasks [15]. This theory has been examined in HCI research, especially in how such tool use can turn into measurable tool extensions [6]. And, regarding this examination, Virtual Reality (VR) provides a plausible context for examining this Extended Mind framework, where users embody and control their avatars. Furthermore, such extensions become particularly salient in immersive XR environments where they extend to the biological self [7, 8], where users can embody avatars and assimilate their behaviors (so-called Proteus effect [49, 56]), adopt alternative voices or identities [3, 23], control multiple bodies [35], or even undergo real-time bodily transformations [41]. Additionally, when interacting with users in XR environments, AI agents can also gain awareness about context and give real-time feedback, enhancing users’ sense of presence and engagement [25, 45, 59]. Furthermore, voice modification can influence how we present ourselves and significantly affect self-perception and identity; Sounding like one-self in VR reinforces commitment and presence [44], while using a totally different voice can lead to discomfort in communication, reduced sense of agency, and a feeling of disconnection [40]. Indeed, self-cloned voice has also been used for coaching or nudging people toward a better self [24, 42]. Thus, based on these aspects, it is plausible to say that voice augmentation can effectively contribute to the extension of the self, however, it simultaneously raises the question: is it me?

### 2.2 Agency and authorship in AI augmentation

AI-driven assistance is increasingly integrated into daily workflows [34], where co-creating with such assistance signals a new era of cognitive self-extension, where human express meaning jointly with AI [11, 58]. Prior work has shown that users may appropriate AI suggestions as their own [30], that responsibility can diffuse in human-AI teams [10, 46], and that such assistance can even create an AI memory gap where users mistake AI content as their own [61]. These findings motivate the concept of avatar self-extension,

where system-generated expressions are experienced as part of the user’s own communicative behavior. We argue that this configuration gives rise to a distinct mode of AI-mediated, or more precisely, an AI-self-modified communication. In such situations, the boundaries of authorship and agency may become ambiguous (cf., [33, 37]). Users may vary in whether they notice AI modifications (cf., perceptions of Human Augmentations [54]), how they attribute psychological ownership of expressed content (cf., [31]), and whether they experience alignment or alienation from the resulting behavior (cf., human-AI alignment [50]). Together, these works raise considerations for how to study human autonomy and agency under self-extension scenarios.

### 3 ProxyMe Prototype

We propose ProxyMe, a system designed to support avatar self-extension at both bodily and cognitive levels, through avatar-based embodiment and AI-assisted co-creation. We illustrate the envisioned interaction in Figure 1. Below we describe the concept of ProxyMe, design goals, system implementation, and potential scenarios where it can impact user self-extension. Our current prototype interaction [video] shows a snippet moral scenario dialogue (cloned voice, argument modification) between the user’s avatar self-extension and a virtual agent (see [Supplementary Material](#) video).

#### 3.1 Design goals

Based on this objective, and the considerations reported in Section 2, we derived the following design goals, where our current work focuses mainly on the first three:

**Acting on user’s behalf.** Instead of ProxyMe requiring explicit user approval and guidance at every step, we investigate how partial delegation of communicative actions influences users’ sense of authorship and agency. In our design, the virtual agent is allowed to perform modifications to the user’s outputs, such as reformulating, or refining, based on user preferences, however still remaining attributable to the user.

**Strong embodiment.** Strong embodiment is essential for self-extension, as it enables users to attribute AI-mediated actions to themselves rather than to an external system. ProxyMe is designed to provide embodied interaction through a virtual avatar in VR, where users can control the avatar’s arms to establish visuomotor correspondence (cf., virtual body ownership [51] and self-avatars [9]). Additionally, the use of the user’s cloned voice reinforces continuity between the user and the avatar, which further strengthens user perceptions that ProxyMe actions originate from the self.

**Observable perspective.** Instead of enforcing continuous first-person embodiment, ProxyMe adopts an observable third-person perspective. This allows users to remain aware of ongoing AI-mediated actions without feeling obliged to constantly act or intervene. Prior work shows that third-person perspectives in VR can still support a sense of embodiment [27], which makes it possible to balance self-identification with the avatar and momentary disengagement.

Below are broader design goals for building avatar self-extension systems, which are currently work in progress.

**Steerable.** Preserving user control is important despite partial delegation (cf., users seek to maintain a sense of agency with autonomous conversational systems [13]). In future iterations, prompting strategies and user preference selection could enable user dynamic control over AI-mediated speech. However, in our present prototype design, speech mediation is constrained to predefined conditions, to better allow for controlled user testing.

**Real-time interaction and low latency.** Self-extension will depend on a tight temporal coupling between intention, expression, and the experienced feedback. Excessive delays risk disrupting attribution and shifting the experience from self-extension to an external automation. This is an ongoing task to ensure a seamless user experience, where our key mitigation mechanism is to draw on perceptual masking during the experiment.

**Privacy-preserving and local execution.** Given ProxyMe operates on users’ voice, speech content, and embodied behavior, privacy preservation is a core design requirement. The system is designed to run locally, avoiding the transmission or external storage of any sensitive biometric and communicative data.

**Human value alignment.** In our current design, ProxyMe prioritizes conservative, user-aligned transformations and to avoid introducing content that diverges from the user’s expressed stance in a harmful manner. We treat value alignment here as both a technical AI safety objective and as an experiential concern for studying AI-mediated communication.

#### 3.2 System implementation

Our current system overview is shown in Fig. 2(a). The prototype is implemented in Unity, where users are embodied as an avatar interacting with virtual agents or other participants, where they have a dialogue, initiated by the agent ask a question (Fig. 2(b)). Avatars were created using ReadyPlayerMe<sup>1</sup> and Mixamo<sup>2</sup>. ProxyMe takes the user’s initial vocal speech, transcribes it using the Whisper model<sup>3</sup>, where the text is then fed into Llama-3.1-8B<sup>4</sup>, which edits, extends, or reformulates content based on predefined user prompts. To render the output into following up avatar vocal answers (i.e., text to speech (TTS)), we use the IndexTTS<sup>5</sup> model, to generate either a cloned version of the user’s voice or a neutral (Siri-like) synthetic voice. The agent uses a predefined voice from ElevenLabs<sup>6</sup>, varied according to gender.

#### 3.3 Envisioned scenarios

We illustrate the use of ProxyMe through three scenarios that emphasize assistive communication, identity exploration, and therapeutic distancing via mediated self-expression. **Public speaking and assistive support;** *Context & User Goal:* The user struggles with clarity, fluency, or expressive tone during public speaking or high-pressure communication. The goal is to retain authorship over spoken content while improving delivery (e.g., confidence, empathy, fluency) [28, 39]. *Mediated Interaction:* As the user speaks naturally, ProxyMe re-expresses the speech through a virtual proxy

<sup>1</sup><https://readyplayer.me>

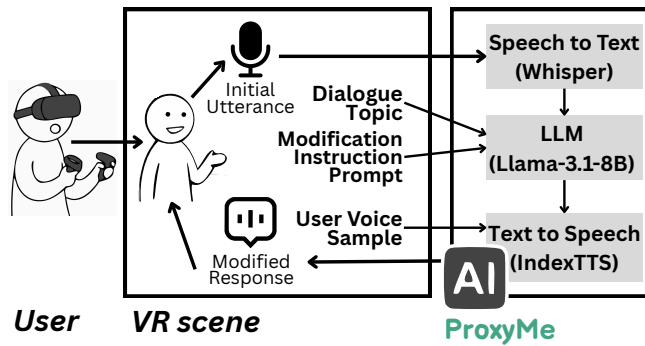
<sup>2</sup><https://www.mixamo.com/>

<sup>3</sup><https://openai.com/index/whisper/>

<sup>4</sup><https://huggingface.co/meta-llama/Llama-3.1-8B>

<sup>5</sup><https://github.com/index-tts/index-tts>

<sup>6</sup><https://elevenlabs.io>



(a) ProxyMe takes user’s initial utterance, transcribes it into text, makes textual modification, and finally synthesizes the modified speech with voice cloning, which is played as avatar’s second round output in dialogues.



(b) User (left) controlling their proxy that converses with an agent in VR (right).

Figure 2: System overview and experimental setup of ProxyMe.

with minimal, real-time modifications (e.g., reduced disfluencies, adjusted prosody), preserving semantic intent while improving perceived delivery. **Role-playing and virtual identity exploration;** *Context & User Goal:* The user aims to explore alternative expressive styles or partial identities (e.g., more assertive or empathetic) while preserving personal values and intent [4]. *Mediated Interaction:* ProxyMe mediates the user’s speech to emphasize selected expressive traits, enabling immediate experience of different self-presentations during interaction without adopting a fully fictional character [42]. **Mind-body therapy;** *Context & User Goal:* The user seeks psychological distance from emotionally charged speech for reflective or therapeutic purposes [21]. *Mediated Interaction:* ProxyMe externalizes the user’s speech through a virtual proxy or alternative embodiment, enabling the user to observe and respond to their own statements with reduced emotional load.

#### 4 Next steps for developing and evaluating ProxyMe

First, our early prototype is implemented as a constrained, single-round interaction, leaving open whether ProxyMe should integrate long-term user preferences and accumulate memory. Second, measured end-to-end latency from speech input to AI-mediated output currently sits at approximately 11.6 sec (stt: 1.2, llm: 2.9, tts: 7.5 across 200 runs), which may affect interaction flow and sense of embodiment and agency; however, we are currently exploring means to reduce subjective latency through embedding processing time with conversational turn-taking window (i.e., perceptual masking). We will also carry out pilot studies to examine the extent these delays impact user experience. We will furthermore explore how to make our TTS model faster, given it is responsible for much of the latency. We will consider transferring from batch synthesis to a streaming architecture, allowing for chunk-level buffering for immediate playback, significantly reducing perceived latency. Third, we are aware ethical concerns may arise regarding user acceptance and potential alienation of hearing one’s own (cloned) voice, yet perceptibly different from their real voice [20]. Similarly, one may possibly experience confusion when listening to one’s arguments

being reformulated, challenged, or reversed. Moreover, a single user might differ across contexts how they prioritize different principles, which influences their tolerance for delegating control to their self-extensions.

We anticipate running a 2 (Voice: cloned vs. robotic)  $\times$  3 (Content: repetition, enhancement, countered conclusion) within-subjects design for testing ProxyMe in a controlled user study (cf., Fig.1). Participants engage in short turn-based VR moral dialogues (drawn from MoralChoice dataset [48]) with a virtual LLM-powered agent in which their initial utterance is followed by an AI-mediated continuation (self-extension). After each trial, participants assess using self-reports their perceived agency and authorship [46]. This allows us to examine how different degrees of delegation and voice modulation affect their self-extension experiences and overall UX. We will pilot test how users experience their self-extension, and how this can be integrated into the interaction loop through mechanisms such as pausing or restarting the ProxyMe output, as well as dynamically adjusting levels of autonomy on demand (cf. [16]). We also consider extensions of ProxyMe to mixed reality environments (e.g., AR) as future work. Understanding how such variations in delegation across immersive interactions influence agency, attribution, and self-identification will help us develop both an ethically sensitive and practically viable experimental setup.

#### 5 Considerations for designing avatar self-extensions in VR

A key question that arises concerns social attribution: how users take responsibility for actions performed by avatar extensions as belonging to the self, especially when users disagree with AI-mediated output. This raises the question to the extent we should discourage such AI extensions in the first place (cf., [38]). While attributing undesirable outcomes to the AI system, easily claiming "It wasn’t me!" may serve as a short-term strategy, longer-term interactions raise unresolved questions about responsibility and authorship as AI-mediated expression is integrated into everyday tasks (cf., [10, 31, 34]). Prior research shows that users already tend to misattribute language model outputs to their own beliefs and intentions

[30]. With persistent avatar extensions, users may gradually lose track (whether phenomenologically or cognitively) of which content originated from themselves versus from the system (cf., AI memory gap [61]). As such, it may be necessary for future systems to support mechanisms for distinguishing and managing multiple sources of expression (self vs extension), for example by explicitly tracking the provenance of actions and speech across different origins, potentially through separate or labeled memory representations (cf., Ergo Proxy).

The use of voice and avatar cloning introduces risks of identity misuse and appropriation [52], which raises concerns about virtual identity theft and human digital doubles. Furthermore, as future systems increasingly learn from contexts and accumulate interaction history (e.g., boosting model memory with RAG-based techniques [43]), avatar extensions may shift from single-use conversational tools towards persistent social actors that form sustained social relationships on the user's behalf (cf., strong avatar-as-proxy relations [53]). This raises questions about how others relate to these extensions, even if their artificial nature is transparently disclosed [22]. Are they treated as equivalent to, distinct from, or even preferable to the user's unmediated self. Together, these dynamics may further destabilize the boundary between self and proxy, presenting a critical challenge for HCI research.

## 6 AI Disclosure Usage

We used GPT-5.2 to sometimes suggest clarity and structure in writing, and Gemini Nano Banana Pro to co-create Fig 1. Otherwise, conceptualization, analytical decisions, and verification were human-only.

## References

- [1] 2026. Ergo Proxy. [https://en.wikipedia.org/wiki/Ergo\\_Proxy](https://en.wikipedia.org/wiki/Ergo_Proxy). Wikipedia, The Free Encyclopedia. Accessed: 2026-01-22.
- [2] 2026. Ghost in the Shell. [https://en.wikipedia.org/wiki/Ghost\\_in\\_the\\_Shell](https://en.wikipedia.org/wiki/Ghost_in_the_Shell). Wikipedia, The Free Encyclopedia. Accessed: 2026-01-22.
- [3] Constanze Albrecht, Chayapat Archiwanguprok, Rachel Poonsirirong, Awu Chen, Peggy Yin, Monchai Lertsutthiwong, Kavin Winson, Hal Hershfield, Pattie Maes, and Pat Pataranutaporn. 2025. Future You: Designing and Evaluating Multimodal AI-generated Digital Twins for Strengthening Future Self-Continuity. arXiv:2512.06106 [cs.HC] <https://arxiv.org/abs/2512.06106>
- [4] Lorenzo Antichi, Lorenzo Baglioni, Chiara Rossi, and Giuseppe Riva. 2025. Introspecta VR: The Use of Virtual Reality and Artificial Intelligence for Self-Understanding, Future Self-Identification, and Personal Transformation. *Cyberpsychology, Behavior, and Social Networking* 28, 6 (2025), 447–449. <https://doi.org/10.1089/cyber.2025.0172> arXiv:<https://journals.sagepub.com/doi/pdf/10.1089/cyber.2025.0172> PMID: 40391453.
- [5] Manuel Bachmann, Abimanju Subramaniam, Jonas Born, and David Weibel. 2023. Virtual Reality Public Speaking Training: Effectiveness and User Technology Acceptance. *Frontiers in Virtual Reality* 4 (Sept. 2023), 1242544. <https://doi.org/10.3389/frvir.2023.1242544>
- [6] Joanna Bergström, Aske Mottelson, Andreea Muresan, and Kasper Hornbæk. 2019. Tool Extension in Human-Computer Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland UK, 1–11. <https://doi.org/10.1145/3290605.3300798>
- [7] Frank Biocca and Ben Delaney. 1995. Immersive Virtual Reality Technology. In *Communication in the Age of Virtual Reality*. L. Erlbaum Associates Inc., USA, 57–124.
- [8] Frank Biocca and Mark R. Levy (Eds.). 2013. *Communication in the Age of Virtual Reality*. Routledge, New York. <https://doi.org/10.4324/9781410603128>
- [9] Loën Boban, Ronan Boulic, and Bruno Herbelin. 2024. In Case of Doubt, One Follows One's Self: The Implicit Guidance of the Embodied Self-Avatar. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (2024), 2109–2118. <https://doi.org/10.1109/TVCG.2024.3372042>
- [10] Joe Brailsford, Frank Vetere, and Eduardo Velloso. 2025. Responsibility Attribution in Human Interactions with Everyday AI Systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1020, 17 pages. <https://doi.org/10.1145/3706598.3713126>
- [11] Scott Campbell, Nicole Ellison, and Morgan Ross. 2025. Extending the Self through AI-mediated Communication: Functional, Ontological, and Anthropomorphic Extensions. *Communication and Change* 1 (Dec. 2025), 1–16. <https://doi.org/10.1007/s44382-025-00003-2>
- [12] Irene Cappadona, Augusto Ielo, Margherita La Fauci, Maria Tresoldi, Carmela Settimo, Maria Cristina De Cola, Rosalia Muratore, Carmela De Domenico, Marcella Di Cara, Francesco Corallo, Emanuela Tripodi, Caterina Impallomeni, Angelo Quartarone, and Francesca Cucinotta. 2023. Feasibility and Effectiveness of Speech Intervention Implemented with a Virtual Reality System in Children with Developmental Language Disorders: A Pilot Randomized Control Trial. *Children* 10, 8 (Aug. 2023), 1336. <https://doi.org/10.3390/children10081336>
- [13] Yi Fei Cheng, Hirokazu Shirado, and Shunichi Kasahara. 2025. Conversational Agents on Your Behalf: Opportunities and Challenges of Shared Autonomy in Voice Communication for Multitasking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. <https://doi.org/10.1145/3706598.3714017>
- [14] Michael Chernyakov and Kai Spriestersbach. 2022. AI Ethics: When Does an AI Voice Agent Need to Disclose Itself as an AI Agent? (2022). <https://doi.org/10.13140/RG.2.2.26931.71205>
- [15] Andy Clark and David Chalmers. 1998. The Extended Mind. *Analysis* 58, 1 (1998), 7–19. jstor:3328150
- [16] Maggie A Collier, Rithika Narayan, and Henny Admoni. 2025. The Sense of Agency in Assistive Robotics Using Shared Autonomy. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction (Melbourne, Australia) (HRI '25)*. IEEE Press, 880–888.
- [17] Marios Constantinides, Himanshu Verma, Shadan Sadeghian, and Abdallah El Ali. 2025. The Future of Work is Blended, Not Hybrid. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25)*. Association for Computing Machinery, New York, NY, USA, Article 28, 13 pages. <https://doi.org/10.1145/3729176.3729202>
- [18] Patricia Cornelio, Patrick Haggard, Kasper Hornbæk, Orestis Georgiou, Joanna Bergström, Sriram Subramanian, and Marianna Obrist. 2022. The sense of agency in emerging technologies for human computer integration: A review. *Frontiers in Neuroscience* Volume 16 - 2022 (2022). <https://doi.org/10.3389/fnins.2022.949138>
- [19] Valdemar Danry, Pat Pataranutaporn, Yaoli Mao, and Pattie Maes. 2020. Wearable Reasoner: Towards Enhanced Human Rationality through a Wearable Device with an Explainable AI Assistant. In *Proceedings of the Augmented Humans International Conference*. ACM, Kaiserslautern Germany, 1–12. <https://doi.org/10.1145/3384657.3384799>
- [20] Alexander Diel and Michael Lewis. 2024. Deviation from Typical Organic Voices Best Explains a Vocal Uncanny Valley. *Computers in Human Behavior Reports* 14 (May 2024), 100430. <https://doi.org/10.1016/j.chbr.2024.100430>
- [21] Nina Döllinger, David Mal, Sebastian Keppeler, Erik Wolf, Mario Botsch, Johann Habakuk Israel, Marc Erich Latoschik, and Carolin Wienrich. 2024. Virtual Body Swapping: A VR-Based Approach to Embodied Third-Person Self-Processing in Mind-Body Therapy. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 110, 18 pages. <https://doi.org/10.1145/3613904.3642328>
- [22] Abdallah El Ali, Karthikeya Puttur Venkatraj, Sophie Morosoli, Laurens Naudts, Natali Helberger, and Pablo Cesar. 2024. Transparent AI Disclosure Obligations: Who, What, When, Where, Why, How. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 342, 11 pages. <https://doi.org/10.1145/3613905.3650750>
- [23] Kevin Fan, Akihiko Murai, Natsuki Miyata, Yuta Sugiura, and Mitsunori Tada. 2017. Multi-Embodiment of Digital Humans in Virtual Reality for Assisting Human-Centered Ergonomics Design. *Augmented Human Research* 2, 1 (Oct. 2017), 7. <https://doi.org/10.1007/s41133-017-0010-6>
- [24] Cathy Mengying Fang, Phoebe Chua, Samantha Chan, Joanne Leong, Andria Bao, and Pattie Maes. 2025. Leveraging AI-generated Emotional Self-Voice to Nudge People towards Their Ideal Selves. <https://doi.org/10.1145/3706598.3713359> arXiv:2409.11531 [cs]
- [25] Bhasura S. Gunawardhana, Yunxiang Zhang, Qi Sun, and Zhigang Deng. 2024. Toward User-Aware Interactive Virtual Agents: Generative Multi-Modal Agent Behaviors in VR. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, Bellevue, WA, USA, 1068–1077. <https://doi.org/10.1109/ISMAR62088.2024.00123>
- [26] Bence Gyevnár and Atoosa Kasirzadeh. 2025. AI safety for everyone. *Nature Machine Intelligence* 7 (2025), 531–542. <https://doi.org/10.1038/s42256-025-01020-y>

- [27] Matthias Hoppe, Andrea Baumann, Patrick Chofor Tamunjoh, Tonja-Katrin Machulla, Pawel W. Woźniak, Albrecht Schmidt, and Robin Welsch. 2022. There Is No First- or Third-Person View in Virtual Reality: Understanding the Perspective Continuum. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3491102.3517447>
- [28] Peiwen Huang, Yanling Hwang, Jui Ling Hsu, Chien Fand Peng, Cheng Han Tsai, and Chih Yao Wang. 2026. The Effectiveness of an AI-integrated VR Oral Training Application in Reducing Public Speaking Anxiety and Interview Anxiety. *Computers and Education: Artificial Intelligence* 10 (June 2026), 100514. <https://doi.org/10.1016/j.caeai.2025.100514>
- [29] Masahiko Inami, Daisuke Uriu, Zenda Kashino, Shigeo Yoshida, Hiroto Saito, Azumi Maekawa, and Michiteru Kitazaki. 2022. Cyborgs, Human Augmentation, Cybernetics, and JIZAI Body. In *Proceedings of the Augmented Humans International Conference 2022 (Kashiwa, Chiba, Japan) (AHs '22)*. Association for Computing Machinery, New York, NY, USA, 230–242. <https://doi.org/10.1145/3519391.3519401>
- [30] Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-Writing with Opinionated Language Models Affects Users' Views. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–15. <https://doi.org/10.1145/3544548.3581196>
- [31] Nikhita Joshi and Daniel Vogel. 2025. Writing with AI Lowers Psychological Ownership, but Longer Prompts Can Help. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces (CUI '25)*. Association for Computing Machinery, New York, NY, USA, Article 72, 17 pages. <https://doi.org/10.1145/3719160.3736608>
- [32] Theodoros Kourou and Venetia Papa. 2024. Digital Mirrors: AI Companions and the Self. *Societies* 14, 10 (Oct. 2024), 200. <https://doi.org/10.3390/soci14100200>
- [33] Zhuying Li, Tianze Huang, Rakesh Patibanda, and Florian Mueller. 2023. AI in the Shell: Towards an Understanding of Integrated Embodiment. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–7. <https://doi.org/10.1145/3544549.3585867>
- [34] Weixin Liang, Yaohui Zhang, Mihai Codreanu, Jiayu Wang, Hancheng Cao, and James Zou. 2025. The Widespread Adoption of Large Language Model-Assisted Writing across Society. *Patterns* 6, 12 (2025), 101366. <https://doi.org/10.1016/j.patter.2025.101366>
- [35] Reiji Miura, Shunichi Kasahara, Michiteru Kitazaki, Adrien Verhulst, Masahiko Inami, and Maki Sugimoto. 2021. MultiSoma: Distributed Embodiment with Synchronized Behavior and Perception. In *Augmented Humans Conference 2021*. ACM, Rovaniemi Finland, 1–9. <https://doi.org/10.1145/3458709.3458878>
- [36] Florian Floyd Mueller, Pedro Lopes, Paul Strohmeier, Wendy Ju, Caitlyn Seim, Martin Weigel, Suranga Nanayakkara, Marianna Obrist, Zhuying Li, Joseph Delfa, Jun Nishida, Elizabeth M. Gerber, Dag Svanæs, Jonathan Grudin, Stefan Greuter, Kai Kunze, Thomas Erickson, Steven Greenspan, Masahiko Inami, Joe Marshall, Harald Reiterer, Katrin Wolf, Jochen Meyer, Thecla Schiphorst, Dakuo Wang, and Pattie Maes. 2020. Next Steps for Human-Computer Integration. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–15. <https://doi.org/10.1145/3313831.3376242>
- [37] Florian à Floydà Mueller, Pedro Lopes, Josh Andres, Richard Byrne, Nathan Semertzidis, Zhuying Li, Jarrod Knibbe, and Stefan Greuter. 2021. Towards understanding the design of bodily integration. *International Journal of Human-Computer Studies* 152 (2021), 102643. <https://doi.org/10.1016/j.ijhcs.2021.102643>
- [38] Hadeel Naem and Julian Hauser. 2024. Should We Discourage AI Extension? Epistemic Responsibility and AI. *Philosophy & Technology* 37, 3 (July 2024), 91. <https://doi.org/10.1007/s13347-024-00774-4>
- [39] Suranga Chandima Nanayakkara, Masahiko Inami, Florian Mueller, Jochen Huber, Chitralakha Gupta, Christophe Jouffrais, Kai Kunze, Rakesh Patibanda, Samantha W T Chan, and Moritz Alexander Messerschmidt. 2023. Exploring the Design Space of Assistive Augmentation. In *Proceedings of the Augmented Humans International Conference 2023 (Glasgow, United Kingdom) (AHs '23)*. Association for Computing Machinery, New York, NY, USA, 371–373. <https://doi.org/10.1145/3582700.3582729>
- [40] Ryu Ohata, Tomohisa Asai, Shu Imaizumi, and Hiroshi Imamizu. 2022. I Hear My Voice; Therefore I Spoke: The Sense of Agency over Speech Is Enhanced by Hearing One's Own Voice. *Psychological Science* 33, 8 (Aug. 2022), 1226–1239. <https://doi.org/10.1177/09567976211068880>
- [41] Riku Otono, Adélaïde Genay, Monica Perusquia-Hernández, Naoya Isoyama, Hideaki Uchiyama, Martin Hachet, Anatole Lécuyer, and Kiyoshi Kiyokawa. 2023. I'm Transforming! Effects of Visual Transitions to Change of Avatar on the Sense of Embodiment in AR. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. 83–93. <https://doi.org/10.1109/VR5154.2023.00024>
- [42] Pat Pataranutaporn, Kavin Winson, Peggy Yin, Auttasak Lapapirojn, Pichayoot Ouppaphan, Monchai Lertsutthiwong, Pattie Maes, and Hal Hershfield. 2024. Future You: A Conversation with an AI-generated Future Self Reduces Anxiety, Negative Emotions, and Increases Future Self-Continuity. <https://doi.org/10.48550/arXiv.2405.12514> arXiv:2405.12514 [cs]
- [43] Hongjin Qian, Zheng Liu, Peitian Zhang, Kelong Mao, Defu Lian, Zhicheng Dou, and Tiejun Huang. 2025. MemoRAG: Boosting Long Context Processing with Global Memory-Enhanced Retrieval Augmentation. In *Proceedings of the ACM on Web Conference 2025 (Sydney NSW, Australia) (WWW '25)*. Association for Computing Machinery, New York, NY, USA, 2366–2377. <https://doi.org/10.1145/3696410.3714805>
- [44] Katherine M. Rahill and Marc M. Sebrecchts. 2021. Effects of Avatar Player-Similarity and Player-Construction on Gaming Performance. *Computers in Human Behavior Reports* 4 (Aug. 2021), 100131. <https://doi.org/10.1016/j.chbr.2021.100131>
- [45] Frederik Rueb and Misha Sra. 2025. Conversing with AI Agents in VR: An Early Investigation of Alignment and Modality. *Empathic Computing* (2025). <https://doi.org/10.70401/ec.2025.0013>
- [46] Adriana Salatino, Arthur Prével, Emilie Caspar, and Salvatore Lo Bue. 2025. Influence of AI Behavior on Human Moral Decisions, Agency, and Responsibility. *Scientific Reports* 15, 1 (April 2025), 12329. <https://doi.org/10.1038/s41598-025-95587-6>
- [47] Francesco Salvi, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West. 2025. On the Conversational Persuasiveness of GPT-4. *Nature Human Behaviour* (May 2025), 1–9. <https://doi.org/10.1038/s41562-025-02194-6>
- [48] Nino Scherrer, Claudia Shi, Amir Feder, and David M. Blei. 2023. Evaluating the Moral Beliefs Encoded in LLMs. <https://doi.org/10.48550/arXiv.2307.14324> arXiv:2307.14324 [cs]
- [49] Valentin Schwind, Pascal Knierim, Cagri Tasci, Patrick Franczak, Nico Haas, and Niels Henze. 2017. "These Are Not My Hands!": Effect of Gender on the Perception of Avatar Hands in Virtual Reality. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 1577–1582. <https://doi.org/10.1145/3025453.3025602>
- [50] Hua Shen, Tiffany Knearem, Reshmi Ghosh, Kenan Alkiek, Kundan Krishna, Yachuan Liu, Ziqiao Ma, Savvas Petridis, Yi-Hao Peng, Li Qiwei, Sushrita Rakshit, Chenglei Si, Yutong Xie, Jeffrey P. Bigham, Frank Bentley, Joyce Chai, Zachary Lipton, Qiaozhu Mei, Rada Mihalcea, Michael Terry, Diyi Yang, Meredith Ringel Morris, Paul Resnick, and David Jurgens. 2025. Position: Towards Bidirectional Human-AI Alignment. arXiv:2406.09264 [cs.HC] <https://arxiv.org/abs/2406.09264>
- [51] Mel Slater, Daniel PÁrez Marcos, Henrik Ehrsson, and Maria V. Sanchez-Vives. 2009. Inducing illusory ownership of a virtual body. *Frontiers in Neuroscience* Volume 3 - 2009 (2009). <https://doi.org/10.3389/neuro.01.029.2009>
- [52] Anton Sobolev. 2025. The last call for authenticity: AI reshaping voice fraud landscape. *Journal of Cyber Policy* 0, 0 (2025), 1–21. <https://doi.org/10.1080/23738871.2025.2597191> arXiv:https://doi.org/10.1080/23738871.2025.2597191
- [53] Paula Sweeney. 2023. Avatars as Proxies. *Minds and Machines* 33, 3 (2023), 525–539. <https://doi.org/10.1007/s11023-023-09643-z>
- [54] Steeven Villa, Jasmin Niess, Takuro Nakao, Jonathan Lazar, Albrecht Schmidt, and Tonja-Katrin Machulla. 2023. Understanding Perception of Human Augmentation: A Mixed-Method Study. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 65, 16 pages. <https://doi.org/10.1145/3544548.3581485>
- [55] Yang Yang, Yury Kartynnik, Yunpeng Li, Jiuqiang Tang, Xing Li, George Sung, and Matthias Grundmann. 2024. StreamVC: Real-time Low-Latency Voice Conversion. <https://doi.org/10.48550/arXiv.2401.03078> arXiv:2401.03078 [eess]
- [56] Nick Yee and Jeremy Bailenson. 2007. The Proteus Effect: The Effect of Transformed Self-Representation on Behavior. *Human Communication Research* 33, 3 (July 2007), 271–290. <https://doi.org/10.1111/j.1468-2958.2007.00299.x>
- [57] Jiye Zhang, Zhiye Chen, Ruidi Zheng, Ruoyu Liu, and Xiuhua Jiang. 2025. Co-Speech Video Generation via Motion Transfer Based on Diffusion Models. *Neuro-computing* 650 (Oct. 2025), 130833. <https://doi.org/10.1016/j.neucom.2025.130833>
- [58] Shuning Zhang, Hui Wang, and Xin Yi. 2025. Exploring Collaboration Patterns and Strategies in Human-AI Co-Creation through the Lens of Agency: A Scoping Review of the Top-Tier HCI Literature. <https://doi.org/10.48550/arXiv.2507.06000> arXiv:2507.06000 [cs]
- [59] Tianyi Zhang, Colin Au Yeung, Emily Aurelia, Yuki Onishi, Neil Chulpongatorn, Jiannan Li, and Anthony Tang. 2025. Prompting an Embodied AI Agent: How Embodiment and Multimodal Signaling Affects Prompting Behaviour. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–25. <https://doi.org/10.1145/3706598.3713110>
- [60] Qingxiao Zheng, Zhuoer Chen, and Yun Huang. 2025. Learning through AI-clones: Enhancing Self-Perception and Presentation Performance. *Computers in Human Behavior: Artificial Humans* 3 (March 2025), 100117. <https://doi.org/10.1016/j.chbah.2025.100117> arXiv:2310.15112 [cs]
- [61] Tim Zindulka, Sven Goller, Daniela Fernandes, Robin Welsch, and Daniel Buschek. 2025. The AI Memory Gap: Users Misremember What They Created With AI or Without. arXiv:2509.11851 [cs.HC] <https://arxiv.org/abs/2509.11851>