



CUI@CHI 2024: Building Trust in CUIs – From Design to Deployment

Smit Desai
smitad2@illinois.edu
University of Illinois at
Urbana-Champaign
Urbana-Champaign, USA

Mateusz Dubiel
mateusz.dubiel@uni.lu
University of Luxembourg
Esch-sur-Alzette, Luxembourg

Martin Porcheron
martin@boldinsight.uk
Bold Insight
London, UK

Heloisa Candello
heloisacandello@br.ibm.com
IBM Research
São Paulo, Brazil

Christina Wei
christina.wei@mail.utoronto.ca
University of Toronto
Toronto, Canada

Nima Zargham
zargham@uni-bremen.de
Digital Media Lab, University of
Bremen
Bremen, Germany

Anastasia Kuzminkyh
anastasia.kuzminkyh@utoronto.ca
University of Toronto
Toronto, Canada

Joel Fischer
Joel.Fischer@nottingham.ac.uk
University of Nottingham
Nottingham, UK

Benjamin R Cowan
benjamin.cowan@ucd.ie
University College Dublin
Dublin, Ireland

Jaisie Sin
jaisie.sin@ubc.ca
University of British Columbia
Toronto, Canada

Shashank Ahire
shashank.ahire@hci.uni-hannover.de
Leibniz University
Hannover, Germany

Minha Lee
m.lee@tue.nl
Eindhoven University of Technology
Eindhoven, The Netherlands

Cosmin Munteanu
cosmin@taglab.ca
University of Waterloo
Waterloo, Canada

ABSTRACT

Conversational user interfaces (CUIs) have become an everyday technology for people the world over, as well as a booming area of research. Advances in voice synthesis and the emergence of chatbots powered by large language models (LLMs), notably ChatGPT, have pushed CUIs to the forefront of human-computer interaction (HCI) research and practice. Now that these technologies enable an elemental level of usability and user experience (UX), we must turn our attention to higher-order human factors: trust and reliance. In this workshop, we aim to bring together a multidisciplinary group of researchers and practitioners invested in the next phase of CUI design. Through keynotes, presentations, and breakout sessions, we will share our knowledge, identify cutting-edge resources, and fortify an international network of CUI scholars. In particular, we will engage with the complexity of trust and reliance as attitudes and

behaviours that emerge when people interact with conversational agents.

CCS CONCEPTS

• **Human-centered computing** → **HCI theory, concepts and models; HCI design and evaluation methods; Natural language interfaces.**

KEYWORDS

conversational user interfaces, chatbots, conversational agents, voice assistants, conversational AI, trust, reliance

ACM Reference Format:

Smit Desai, Christina Wei, Jaisie Sin, Mateusz Dubiel, Nima Zargham, Shashank Ahire, Martin Porcheron, Anastasia Kuzminkyh, Minha Lee, Heloisa Candello, Joel Fischer, Cosmin Munteanu, and Benjamin R Cowan. 2024. CUI@CHI 2024: Building Trust in CUIs – From Design to Deployment. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3613905.3636287>

1 MOTIVATION

Conversational User Interfaces (CUIs), such as virtual assistants like Apple Siri and Amazon Alexa, alongside Large Language Models

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CHI EA '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0331-7/24/05
<https://doi.org/10.1145/3613905.3636287>

(LLMs) like OpenAI’s ChatGPT and Google Bard, are steadily becoming integral to our daily routines. Users engage in a wide range of activities with these technologies, from setting reminders and shopping online to seeking information and having casual conversations [37]. Regardless of the type of interaction—transactional or social—it necessitates the establishment of trust [36]. This trust is paramount because it not only ensures the accuracy and reliability of information and services but also influences how comfortable users feel when sharing personal data or forging deep connections with these AI-powered systems [19]. In this workshop, our main goal is to unite dedicated researchers from the CHI community for an in-depth exploration of CUIs from a trust perspective. We will delve into the unique nuances of trust within CUIs, addressing open questions and charting the course for future developments in this evolving landscape.

Despite trust being crucial in the adoption and continued use of automated systems [31, 35, 36, 38], it is understudied in CUI, as exemplified by calls for research in CUI’19 [21] and again in CUI’22 [19]. This lack of research is perhaps most evident in the opacity associated with the fundamental nature of trust, and specifically its conceptualization [24]. While the term ‘trust’ finds frequent usage within HCI [4] and the more specialized literature on CUIs [19, 21, 27, 43], a clear and generally accepted definition within this context remains elusive. In CUIs, trust can manifest in diverse ways: (1) it may encompass both cognitive elements rooted in thinking and judgment and affective aspects rooted in feelings [42], (2) it often comprises attitudinal [25] and behavioral components [41], (3) it can be user-centric or agent-centric (including the owner of the agent) [42], and (4) its nature is inherently context-dependent [4, 44]. This complexity is particularly pronounced due to the natural modality employed by users when interacting with these systems, which introduces a unique set of trust dynamics. Given the ‘black-box’ nature of conversational interfaces, users frequently rely on anthropic tendencies [28]. Anthropomorphism only serves as a means of sense-making or constructing a mental model in this context [16]. However, it is noteworthy that anthropomorphism is often intentionally incorporated into these interfaces, including the utilization of pronouns, natural-sounding voices, and the emulation of human-human communication models, among other techniques [16, 22, 23, 34, 40].

Overtrust and undertrust in CUIs can have detrimental consequences [31, 39], potentially resulting in discontinued or infrequent use of such systems or conversely, fostering addiction and diminishing user agency. [10, 12, 29]. Specifically, in critical domains such as finance and health, overreliance on voice-based advice may lead to misguided financial decisions [19] or, in the case of health, the inadvertent disclosure of sensitive medical information without a complete understanding of the associated privacy risks, including potential breaches of confidentiality and unauthorized access to personal health data [9]. These concerns may be exacerbated by demographic factors, such as age, where older individuals may be more vulnerable to trust-related issues due to their potentially limited familiarity with AI technologies [34]. Moreover, these challenges can vary across cultural contexts, as highlighted in the Global South, where cultural factors and technological savviness can significantly influence the levels of trust individuals place in these technologies [32]. Furthermore, as the conversational abilities of

these technologies advance, the emotional bonds they develop with users may deepen, intensifying the emotional distress experienced in the event of dysfunction or breakdown [26], further underscoring the need for careful consideration of trust dynamics in CUIs.

Within the landscape of trust calibration, diverse strategies have emerged, ranging from efforts to introduce a sense of ‘distrust’ [33] to persistent calls for transparency, notably advocated by the burgeoning field of Explainable AI (XAI) [45]. However, a comprehensive discussion of trust requires a multifaceted examination from various perspectives, encompassing specialized areas within CHI, such as Human-AI Interaction (HAI), Human-Robot Interaction (HRI), XAI, and CUI. Our objective is to foster collaborations that could help to converge these diverse approaches and ultimately lead to the development of a cohesive framework for a systematic study of trust. In our workshop discussions, we will reflect on:

- **Conceptual Clarity of Trust in CUIs:** The foundation of our exploration involves refining and defining the conceptual boundaries of trust within the context of CUIs. We aim to foster a more informed and unified discourse by establishing a clear and shared understanding of trust.
- **Design Techniques and Trust Calibration:** We will scrutinize design strategies that can inadvertently result in the miscalibration of user trust and raise ethical concerns. An illustrative example is the integration of human-like traits into CUIs, which can impact user perceptions.
- **Interconnections of Trust with Other Values:** Trust can exhibit a delicate balance, be diminished, or even conflict with other values such as privacy, transparency, or usability. This exploration discusses the emergence of tensions in specific scenarios and highlights how design decisions in CUIs can significantly impact these perceptions.
- **Integrative Measurement Methods:** Developing comprehensive methods for measuring trust, combining objective and subjective scales, is important to our inquiry to mitigate overtrust and undertrust.
- **Temporal evolution of trust:** Trust is a dynamic concept that evolves over time and context. We will explore the temporal aspects of trust, examining how it fluctuates in different stages of interaction and adoption.

By exploring these and other emergent topics, we aspire to foster an environment where interdisciplinary discussions on trust can flourish, drawing inspiration from diverse strategies within trust calibration and design. In alignment with the overarching theme of the ACM CUI’24 conference, which focuses on designing trustworthy conversational interfaces, our workshop aims to contribute a deeper understanding of trust in CUIs across all stages—from design to deployment.

2 ORGANIZERS

This workshop is being organized by a team of experienced researchers who bring expertise from diverse fields aligned with our objectives. Collectively, the organizers have a track record of publishing on a wide range of topics, including conversational user interfaces [11, 18, 20, 26, 27, 46, 47], human-robot interaction [8, 30], multi-modal communication [3, 7], and privacy and security [5, 6, 15, 24]. Integrating these disciplines shall facilitate

collaboration among our trans-disciplinary participants and bridge the gap between these research areas.

Smit Desai is a Ph.D. candidate in the School of Information Sciences at the University of Illinois, Urbana-Champaign. His primary research focus centers around comprehending the mental models of users as they engage with conversational agents, utilizing innovative research techniques such as metaphor analysis [17]. He leverages this valuable insight to advance the development of conversational agents in diverse social roles, including educators [13] and storytellers [14]. Smit is a Provocation Papers Co-Chair at the upcoming ACM CUI 2024 conference. His research has yielded publications in esteemed HCI forums like CHI, TOCHI, CSCW, and CUI.

Christina Wei is a Ph.D. student in the School of Information at the University of Toronto, Canada. Her research focuses on designing conversation architecture for AI agents [43] in human-AI collaborative decision-making scenarios. She has previously organized a workshop at ACM CUI 2023, and she is on the organizing committee for the upcoming ACM CUI 2024 conference.

Jaisie Sin is a Post-Doctoral Fellow at the Okanagan Visualization & Interaction Lab at the University of British Columbia. Her research focuses on the inclusive design of conversational interfaces for underrepresented users, with a primary focus on ageing. She is an Accessibility Chair at the upcoming CUI 2024 conference. She has also been a Full Papers Co-Chair at CUI '23, and co-organizer of previous CUI conferences related workshops at CHI '19–'23 (lead in '23), IUI '20–'21, and CSCW '20.

Mateusz Dubiel is a Research Associate in the Department of Computer Science at the University of Luxembourg, where he works on development and evaluation of conversational agents. Specifically, his current research focuses on assessment of cognitive and usability implications of interfaces that feature speech, and exploration of their potential to persuade and inspire positive behavioural change in users. He served as Short Papers Chair for CUI '22 and will be one of General Chairs for the upcoming CUI '24.

Nima Zargham is a Ph.D. student in the Digital Media Lab at the University of Bremen. His research focuses on human-centered approaches for designing speech-based systems that elicit desirable user experiences. Nima has previously organized CUI-related workshops at notable conferences such as ACM/IEEE HRI 2023 and ACM CUI 2023. Additionally, he served as a local chair at the ACM CHI-PLAY 2022 conference. His research efforts have resulted in publications featured in prestigious HCI venues, including CHI, CUI, and CHI-PLAY.

Shashank Ahire is a PhD candidate in Mensch Computer Interaktion group at Leibniz University Hannover. He is interested in developing CUIs to improve the health and well-being of knowledge Workers while at work [2]. Previously, he has also researched on designing and developing CUIs for marginalized populations in the global south [1].

Martin Porcheron is a Senior UX Researcher at Bold Insight UK, based in London. His recent work focuses on the user experience of interactive “AI” systems—from voice interfaces through to robot interfaces. In his Ph.D., he also studied public places like cafés and pubs, to understand how we use devices like smartphones while we are socialising. As well as organising a number of workshops at

international conferences, he was Full Papers Chair for the inaugural CUI '19 conference and Programme Chair for CUI '20 and '21, and was General Chair for CUI '23.

Anastasia Kuzminkyh is an Assistant Professor in Human-Computer Interaction at the University of Toronto in the Faculty of Information and the Director of the Toronto Human-AI Interaction Research School (THAIRS). Her work touches on diverse technologies, including conversational agents, large language models, and AI-based decision support systems, exploring the mechanisms driving the perceptions of different types of AI systems, such as trust, anthropomorphization, or perceived reliability to inform the system and algorithm design requirements for effective, efficient, and ethical AI.

Minha Lee is an assistant professor at the Eindhoven University of Technology in the Department of Industrial Design, with a background in philosophy, digital arts, and HCI. Her research concerns morally relevant interactions with various agents like robots or chatbots, exploring moral concepts like compassion and trust. She organized workshops such as CHI 2022 workshop on CUI Ethics, CSCW 2022 workshop on Activism in Academia, HRI 2021 workshop on Artificial Identity and more, as well as serving as a General Chair for CUI '23, Full papers for Chair CUI '21, and Provocation papers Chair for CUI '20.

Heloisa Candello is a research scientist and a manager of the Human-centered and Responsible technologies group at IBM Research. Her work focuses on human and social aspects of Artificial Intelligence systems, particularly conversational user interfaces. Currently, Heloisa is leading a project that aims to bring “conscious” access to micro-credit by enhancing non-traditional financial practices of low-income small business owners with AI technology in the Global South. Her research resulted in several publications in leading conferences (CHI, CUI, CSCW, DRS) and recognition in the HCI and Design field.

Joel Fischer is Professor of Human-Computer Interaction at the University of Nottingham, and Research Director of two national programmes on Responsible and Trustworthy AI (TAS and RAI UK). Joel's interests include close studies of conversational interaction by adopting a conversation analytic perspective. He has published widely and co-organised many workshops on the topic at CHI, CUI, CSCW, and HRI.

Cosmin Munteanu is an Associate Professor and Schlegel Research Chair in Technology for Healthy Aging at the Department of Systems Design Engineering, University of Waterloo, and Director of the Technologies for Ageing Gracefully lab. Cosmin takes a primarily ethnomethodological approach to study how voice- and language-enabled interfaces should be designed in a safe, effective, inclusive, and ethical manner, in order to empower digitally underrepresented groups such as older adults. Cosmin is the co-founder of the new ACM Conversational User Interfaces conference series and one of the early researchers to bring speech processing and interface design research together. He is also serving as a member of the ACM SIGCHI Ethics Committee where he contributes perspectives on ethical research with underrepresented user groups.

Benjamin Cowan is a Professor at University College Dublin's School of Information and Communication Studies. His research

focuses on understanding user interaction with speech-based conversational user interfaces. He is co-founder of the ACM Conversational User Interfaces conference (ACM CUI) and has published extensively on user interaction issues with speech interfaces within the HCI community. He is an Associate Editor of the International Journal of Human-Computer Studies, has served as CHI AC, and was one of the inaugural SCs of the Understanding People: Quantitative Methods subcommittee at CHI 2023.

3 IN-PERSON AND HYBRID

Based on our past experience in conducting workshops, we are aware of the various limitations attendees may face, (visa issues, financial constraints, and personal commitments) which can prevent them from attending in-person workshops. Therefore, we plan to host both an in-person and a hybrid workshop. Similar to the previous year, we aim to facilitate synchronous discussions for in-person attendees (with dedicated cameras) and online participants. Additionally, for the purpose of brainstorming and idea generation, we intend to synchronize the experiences of both in-person and online attendees by utilizing online platforms such as Miro.

We are committed to supporting the CHI2024 initiative aimed at creating an inclusive environment for in-person and hybrid attendees¹. In line with this commitment, we will kindly request all interested attendees to follow the 'ACM Guide for Accessible Submission²' when preparing their submissions. Additionally, we will encourage them to conduct an accessibility check³ before finalizing their submissions. To ensure an inclusive experience during the workshop, we will provide automatic transcription of keynote, and paper presentations for online participants.

4 ASYNCHRONOUS ENGAGEMENT

To advertise our workshop, we will use our existing channels: email list (cui-announcements@listserv.acm.org), Twitter handle (https://twitter.com/ACM_CUI), Discord channel (<https://discord.gg/rDmctrVb>) and workshop website (<https://cui.acm.org/workshops/CHI2024/>). During the workshop, we will create a sub-channel (CUI@CHI2024) on existing Discord to disseminate information.

We will encourage participants to create 3-minute video presentations, which will be archived on the CUI channel on YouTube. The YouTube video of the presentation will offer automated subtitles to ensure accessibility for the uploaded presentations. We will record the keynote presentation and a summary of the discussions during the workshop. The video recordings will be accessible on Discord for participants to view asynchronously.

5 WORKSHOP ACTIVITIES

We plan for a one-day, in-person/hybrid workshop structured in a series of presentations, activities, and structured discussions to exchange insights and lessons learned related to the building trust in CUIs from design to deployment. We are aiming for approximately 20 participants, which has been the norm in the past CUI workshops. Our tentative schedule is below.

- (1) **Introduction (15 mins):** We start with brief introductions from organizers, highlighting the goals of the workshop to build trustworthy CUIs, and go over plans for the day.
- (2) **Keynote Address (45 mins):** Our keynote address will be delivered by **Dr. S. Shyam Sundar**. He is the founder of the Media Effects Research Laboratory, a leading facility in the USA. He also serves as the Director of the Center for Socially Responsible Artificial Intelligence, an interdisciplinary consortium at Penn State University, USA dedicated to AI development with a strong emphasis on its social and ethical implications, aiming to enhance societal good and mitigate misuse. The speech will kick off the workshop to discuss different aspects of building trust with machines, such as explainable AI, trust calibration, and specific considerations for conversational interfaces. This keynote will set a solid foundation for the upcoming breakout discussions.
- (3) **Introduction & Presentations (60 mins):** Each participant will have 3-5 minutes to introduce themselves and present their accepted paper.
- (4) **Break (15 mins)**
- (5) **Breakout Session 1 - Prototype activity (60 mins):** Participants will be divided into 3 or 4 groups for small group discussions, with a designated facilitator assigned for each group. Each group will be given a scenario to design a trustworthy CUI interaction. For example, designing a CUI to provide financial guidance to help individual stay within their budgets, or designing a CUI as a social companion for older adults. At the end of the session, each group will present their CUI prototype and discuss the elements they selected to incorporate in their design, as well as received feedback on their design from the broader audience.
- (6) **Lunch (60 mins)**
- (7) **Breakout Session 2 - Design Techniques (60 mins):** Participants and facilitators will go back to their small groups to discuss different design techniques for building trust in CUIs related to their given scenario. This includes questions such as: What are the different trust design techniques for CUIs? What are the contextual considerations when applying these design techniques? What improvements can be made to the prototype in Breakout 1 based on these design techniques? A sharing session will follow for each group to present a summary of their discussions as well as proposed design guidelines.
- (8) **Breakout Session 3 - Evaluation Techniques & Methods (60 mins):** The session will be introduced by Dr. Katie Seaborn. They are an Associate Professor at Tokyo Institute of Technology and founder of the Aspirational Computing Lab. Their research focuses on voice UX, AI, and intersectional design. Afterwards, participants and facilitators will go back to their small groups to discuss methods for measuring trust related to their given scenario. This includes questions such as: What are the existing methods used to measure trust with CUIs? What aspects of trust do these methods measure? What are some additional considerations for measuring trust that need to be taken into account for the given scenario? A sharing session will follow for each group to present a summary of their discussions.

¹<https://chi2024.acm.org/for-attendees/accessibility-faq/>

²<https://sigchi.org/conferences/author-resources/accessibility-guide/>

³<https://www.youtube.com/watch?v=G8EYDJ56c6k>

(9) Break (15 mins)

- (10) Closing (60 mins):** The organizers will end the workshop by synthesizing the main points from previous sessions. Participants and organizers will reflect together to develop tangible research topics and potential outcomes for CUI research. This will help in the development of future directions for the CUI community to consider in relation to building trustworthy CUIs, including how to plan for upcoming CUI conference series and workshops by taking into account participants' views.

6 POST-WORKSHOP PLANS

The expected workshop outcomes include:

- Reconnecting, sustaining, and extending the existing CUI community of researchers.
- Invitation of participants to trustworthy design initiatives within the CUI community.
- Identification of major themes in trustworthy design relevant to CUI research.
- A proposal for a special issue at relevant journals to highlight trust building design considerations relevant to CUI research, e.g., ACM interactions magazine, International Journal of Human-Computer Interaction, International Journal of Human-Computer Studies.
- A white paper summarizing outcomes from the workshop to be posted on our workshop's website, as well as socialized through platforms such as Twitter, Discord, and Medium.
- Invitation for a selection of papers to be published in workshop proceedings as outlined below.

7 PLANS TO PUBLISH WORKSHOP PROCEEDINGS

The primary objective of this workshop is to foster collaboration among individuals with diverse backgrounds and perspectives, all of whom share an interest in examining the role of trust in interactions with CUIs and strategies for establishing trust between users and these agents. All accepted contributions from participants will be published on our workshop's website to ensure the broad dissemination of the workshop's findings. Additionally, a selection of the accepted papers will be included in the workshop proceedings, which we aspire to publish on *arXiv*⁴ or *ceur-ws*⁵. Ideas and discussion points will be documented throughout the session on an open online platform, both during and after the symposium. Workshop participants will also be invited to engage in future collaborative projects that emerge from the discussions and talks held during the workshop. Finally, we will strongly encourage all participants to submit extended versions of their contributions for consideration in a special journal issue, which we are committed to facilitating and making a reality.

⁴<https://arxiv.org/>

⁵<https://ceur-ws.org/>

8 CALL FOR PARTICIPATION

How can we advance research on Conversational User Interfaces (CUIs) by including consideration of trust? Since CUIs, such as Apple Siri, Amazon Alexa and Google Home are becoming increasingly used in privacy-sensitive contexts such as self-reflection, tutoring or therapy, investigation of trust and impact of user's reliance and agency becomes ever more pertinent. What is more, proliferation of LLMs, such as OpenAI's ChatGPT and Google Bard and growing role their role in assisting human decision-making also calls for investigation.

Authors across all disciplines are invited to submit position papers to the CHI 2024 workshop on the Building Trust in CUIs—From Design to Deployment. The goal of this edition of the workshop is to encourage discussion on implementing ethical and user-oriented design practices to develop CUIs that are focused on enhancing user's agency while preserving privacy and preventing overreliance. We expect that the workshop will encourage contributions and discussions from participants from both industry and academia. Papers can address the state of CUIs and their design for privacy-sensitive scenarios where trust plays an essential role, ways to practice design of trustworthy CUIs, and how we can move forward in design of CUIs that empower users and enhance their agency.

Submissions are expected to be between 2 to 4 pages in the CHI Publication Format⁶, including references, authors' work thus far, disciplinary background(s), as well as how they relate to the topics and goals of the workshop. We will select submissions based on quality; if accepted, papers will be featured on our website and at least one author must attend the workshop. Papers should be submitted to smitad2@illinois.edu. Those who cannot or do not wish to submit a position paper to the workshop can submit a one-page statement of interest for consideration. This should include why the applicant is interested in participating in the workshop, their relevancy to the workshop, and what the person expects to contribute.

ACKNOWLEDGMENTS

We thank Dr. S. Shyam Sundar for accepting the invitation to present the keynote address at the workshop. We also thank Dr. Katie Seaborn for contributing to this workshop proposal.

REFERENCES

- [1] Shashank Ahire. 2023. Designing a Smart Speaker for Emergent Users: Human Plus AI Response. In *Proceedings of the 13th Indian Conference on Human-Computer Interaction* (Hyderabad, India) (*IndiaHCI '22*). Association for Computing Machinery, New York, NY, USA, 67–72. <https://doi.org/10.1145/3570211.3570217>
- [2] Shashank Ahire, Michael Rohs, and Simon Benjamin. 2022. Ubiquitous Work Assistant: Synchronizing a Stationary and a Wearable Conversational Agent to Assist Knowledge Work. In *Proceedings of the 1st Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Durham, NH, USA) (*CHIWORK '22*). Association for Computing Machinery, New York, NY, USA, Article 3, 9 pages. <https://doi.org/10.1145/3533406.3533420>
- [3] Vino Avanesi, Johanna Rockstroh, Thomas Mildner, Nima Zargham, Leon Reicherts, Maximilian A. Friehs, Dimosthenis Kontogiorgos, Nina Wenig, and Rainer Malaka. 2023. From C-3PO to HAL: Opening The Discourse About The Dark Side of Multi-Modal Social Agents. In *Proceedings of the 5th International Conference on Conversational User Interfaces* (Eindhoven, Netherlands) (*CUI '23*). Association for Computing Machinery, New York, NY, USA, Article 62, 7 pages. <https://doi.org/10.1145/3571884.3597441>

⁶<https://chi2024.acm.org/submission-guides/chi-publication-formats/>

- [4] Tita Alissa Bach, Amna Khan, Harry Hallock, Gabriela Beltrão, and Sonia Sousa. 2022. A Systematic Literature Review of User Trust in AI-Enabled Systems: An HCI Perspective. *International Journal of Human-Computer Interaction* 0, 0 (2022), 1–16. <https://doi.org/10.1080/10447318.2022.2138826>
- [5] Mehrdad Bahrini, Nima Zargham, Johannes Pfau, Stella Lemke, Karsten Sohr, and Rainer Malaka. 2020. Good vs. Evil: Investigating the Effect of Game Premise in a Smart Home Security Educational Game. In *Extended Abstracts of the 2020 Annual Symposium on Computer-Human Interaction in Play* (Virtual Event, Canada) (CHI PLAY '20). Association for Computing Machinery, New York, NY, USA, 182–187. <https://doi.org/10.1145/3383668.3419887>
- [6] Mehrdad Bahrini, Nima Zargham, Alexander Wolff, Dennis-Kenji Kipker, Karsten Sohr, and Rainer Malaka. 2022. It's Long and Complicated! Enhancing One-Pager Privacy Policies in Smart Home Applications. In *Nordic Human-Computer Interaction Conference* (Aarhus, Denmark) (NordCHI '22). Association for Computing Machinery, New York, NY, USA, Article 73, 13 pages. <https://doi.org/10.1145/3546155.3546657>
- [7] Michael Bonfert, Nima Zargham, Florian Saade, Robert Porzel, and Rainer Malaka. 2021. An Evaluation of Visual Embodiment for Voice Assistants on Smart Displays. In *CUI 2021-3rd Conference on Conversational User Interfaces*. ACM, New York, NY, USA, 1–11.
- [8] Andriana Boudouraki, Joel E Fischer, Stuart Reeves, and Sean Rintel. 2023. "Being in on the Action" in Mobile Robotic Telepresence: Rethinking Presence in Hybrid Participation. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 63–71.
- [9] Eugene Cho. 2019. Hey Google, Can I Ask You Something in Private?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (CHI '19). ACM, New York, NY, USA, 258:1–258:9. <https://doi.org/10.1145/3290605.3300488> event-place: Glasgow, Scotland Uk.
- [10] Minji Cho, Sang-su Lee, and Kun-Pyo Lee. 2019. Once a Kind Friend is Now a Thing: Understanding How Conversational Agents at Home are Forgotten. In *Proceedings of the 2019 on Designing Interactive Systems Conference* (DIS '19). Association for Computing Machinery, New York, NY, USA, 1557–1569. <https://doi.org/10.1145/3322276.3322332>
- [11] Leigh Clark, Nadia Pantidi, Orla Cooney, Philip Doyle, Diego Garaialde, Justin Edwards, Brendan Spillane, Emer Gilmartin, Christine Murad, Cosmin Munteanu, et al. 2019. What makes a good conversation? Challenges in designing truly conversational agents. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–12.
- [12] Benjamin R. Cowan, Nadia Pantidi, David Coyle, Kellie Morrissey, Peter Clarke, Sara Al-Shehri, David Earley, and Natasha Bandeira. 2017. "What can i help you with?": infrequent users' experiences of intelligent personal assistants. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '17). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3098279.3098539>
- [13] Smit Desai and Jessie Chin. 2023. OK Google, Let's Learn: Using Voice User Interfaces for Informal Self-Regulated Learning of Health Topics among Younger and Older Adults. 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 847, 21 pages. <https://doi.org/10.1145/3544548.3581507>
- [14] Smit Desai, Morgan Lundy, and Jessie Chin. 2023. "A Painless Way to Learn:" Designing an Interactive Storytelling Voice User Interface to Engage Older Adults in Informal Health Information Learning. In *Proceedings of the 5th International Conference on Conversational User Interfaces* (Eindhoven, Netherlands) (CUI '23). Association for Computing Machinery, New York, NY, USA, Article 5, 16 pages. <https://doi.org/10.1145/3571884.3597141>
- [15] Smit Desai, Tanusree Sharma, and Pratyasha Saha. 2023. Using ChatGPT in HCI Research—A Trioethnography. In *Proceedings of the 5th International Conference on Conversational User Interfaces*. 1–6.
- [16] Smit Desai and Michael Twidale. 2022. Is Alexa like a computer? A search engine? A friend? A silly child? Yes.. In *Proceedings of the 4th Conference on Conversational User Interfaces* (CUI '22). Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3543829.3544535>
- [17] Smit Desai and Michael Twidale. 2023. Metaphors in Voice User Interfaces: A Slippery Fish. *ACM Trans. Comput.-Hum. Interact.* 30, 6, Article 89, 37 pages. <https://doi.org/10.1145/3609326>
- [18] Mateusz Dubiel, Alessandra Cervone, and Giuseppe Riccardi. 2019. Inquisitive mind: a conversational news companion. In *Proceedings of the 1st International Conference on Conversational User Interfaces*. 1–3.
- [19] Mateusz Dubiel, Sylvain Daronnat, and Luis A. Leiva. 2022. Conversational Agents Trust Calibration: A User-Centred Perspective to Design. In *Proceedings of the 4th Conference on Conversational User Interfaces* (CUI '22). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3543829.3544518>
- [20] Mateusz Dubiel, Martin Halvey, Pilar Oplustil Gallegos, and Simon King. 2020. Persuasive synthetic speech: Voice perception and user behaviour. In *Proceedings of the 2nd Conference on Conversational User Interfaces*. 1–9.
- [21] Justin Edwards and Elaheh Sanoubari. 2019. A need for trust in conversational interface research. In *Proceedings of the 1st International Conference on Conversational User Interfaces* (CUI '19). Association for Computing Machinery, New York, NY, USA, 1–3. <https://doi.org/10.1145/3342775.3342809>
- [22] Dimosthenis Kontogiorgos, Andre Pereira, Olle Andersson, Marco Koivisto, Elena Gonzalez Rabal, Ville Vartiainen, and Joakim Gustafson. 2019. The Effects of Anthropomorphism and Non-Verbal Social Behaviour in Virtual Assistants. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents* (IVA '19). Association for Computing Machinery, New York, NY, USA, 133–140. <https://doi.org/10.1145/3308532.3329466> event-place: Paris, France.
- [23] Anastasia Kuzminykh, Jenny Sun, Nivetha Govindaraju, Jeff Avery, and Edward Lank. 2020. Genie in the Bottle: Anthropomorphized Perceptions of Conversational Agents. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376665>
- [24] Effie Lai-Chong Law, Asbjørn Følstad, Jonathan Grudin, and Björn Schuller. 2022. Conversational Agent as Trustworthy Autonomous System (Trust-CA)(Dagstuhl Seminar 21381). In *Dagstuhl Reports*, Vol. 11. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [25] John D. Lee and Katrina A. See. 2004. Trust in automation: designing for appropriate reliance. *Human Factors* 46, 1 (2004), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- [26] Minha Lee, Lily Frank, Yvonne De Kort, and Wijnand IJsselstein. 2022. Where is Vincent? Expanding our emotional selves with AI. In *Proceedings of the 4th Conference on Conversational User Interfaces* (CUI '22). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3543829.3543835>
- [27] Minha Lee, Lily Frank, and Wijnand IJsselstein. 2021. Brokerbot: A cryptocurrency chatbot in the social-technical gap of trust. *Computer Supported Cooperative Work* (CSCW) 30, 1 (2021), 79–117.
- [28] Q. Vera Liao and Jennifer Wortman Vaughan. 2023. AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. arXiv:2306.01941 (Aug. 2023). <https://doi.org/10.48550/arXiv.2306.01941> arXiv:2306.01941 [cs].
- [29] Ewa Luger and Abigail Sellen. 2016. "Like Having a Really Bad PA": The Gulf between User Expectation and Experience of Conversational Agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16). Association for Computing Machinery, New York, NY, USA, 5286–5297. <https://doi.org/10.1145/2858036.2858288>
- [30] Donald McMullan, Razan Jaber, Benjamin R. Cowan, Joel E. Fischer, Bahar Iran, Ronald Cumbal, Nima Zargham, and Minha Lee. 2023. Human-Robot Conversational Interaction (HRCI). In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) (HRI '23). Association for Computing Machinery, New York, NY, USA, 923–925. <https://doi.org/10.1145/3568294.3579954>
- [31] Bonnie M Muir. 1987. Trust between humans and machines, and the design of decision aids. *International journal of man-machine studies* 27, 5-6 (1987), 527–539.
- [32] Chinasa T. Okolo, Srujana Kamath, Nicola Dell, and Aditya Vashistha. 2021. "It cannot do all of my work": Community Health Worker Perceptions of AI-Enabled Mobile Health Applications in Rural India. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (CHI '21). Association for Computing Machinery, New York, NY, USA, 1–20. <https://doi.org/10.1145/3411764.3445420>
- [33] Claudio S. Pinhanez. 2021. Expose Uncertainty, Instill Distrust, Avoid Explanations: Towards Ethical Guidelines for AI. arXiv:2112.01281 (Nov. 2021). <https://doi.org/10.48550/arXiv.2112.01281> arXiv:2112.01281 [cs].
- [34] Alisha Pradhan, Leah Findlater, and Amanda Lazar. 2019. "Phantom Friend" or "Just a Box with Information": Personification and Ontological Categorization of Smart Speaker-based Voice Assistants by Older Adults. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 214:1–214:21. <https://doi.org/10.1145/3359316>
- [35] Alisha Pradhan, Amanda Lazar, and Leah Findlater. 2020. Use of intelligent voice assistants by older adults with low technology use. *ACM Transactions on Computer-Human Interaction* (TOCHI) 27, 4 (2020), 1–27.
- [36] Minjin Rheu, Ji Youn Shin, Wei Peng, and Jina Huh-Yoo. 2021. Systematic Review: Trust-Building Factors and Implications for Conversational Agent Design. *International Journal of Human-Computer Interaction* 37, 1 (Jan. 2021), 81–96. <https://doi.org/10.1080/10447318.2020.1807710>
- [37] Alex Sciuto, Arnita Saini, Jodi Forlizzi, and Jason I. Hong. 2018. "Hey Alexa, What's Up?": A Mixed-Methods Studies of In-Home Conversational Agent Usage. In *Proceedings of the 2018 Designing Interactive Systems Conference* (DIS '18). ACM, New York, NY, USA, 857–868. <https://doi.org/10.1145/3196709.3196772> 00000 event-place: Hong Kong, China.
- [38] Jaisie Sin, Cosmin Munteanu, Numrita Ramanand, and Yi Rong Tan. 2021. VUI influencers: How the media portrays voice user interfaces for older adults. In *Proceedings of the 3rd Conference on Conversational User Interfaces*. 1–13.
- [39] Linda J Skitka, Kathleen L Mosier, and Mark Burdick. 1999. Does automation bias decision-making? *International Journal of Human-Computer Studies* 51, 5 (1999), 991–1006.

- [40] Clara Strathmann, Jessica Szczuka, and Nicole Krämer. 2020. She Talks to Me as If She Were Alive: Assessing the Social Reactions and Perceptions of Children toward Voice Assistants and Their Appraisal of the Appropriateness of These Reactions. In *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents (IVA '20)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3383652.3423906> event-place: Virtual Event, Scotland, UK.
- [41] Matthias Söllner, Axel Hoffmann, and Jan Marco Leimeister. 2016. Why different trust relationships matter for information systems users. *European Journal of Information Systems* 25, 3 (May 2016), 274–287. <https://doi.org/10.1057/ejis.2015.17>
- [42] Takane Ueno, Yuto Sawa, Yeongdae Kim, Jacqueline Urakami, Hiroki Oura, and Katie Seaborn. 2022. Trust in Human-AI Interaction: Scoping Out Models, Measures, and Methods. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3491101.3519772>
- [43] Christina Ziyang Wei, Young-Ho Kim, and Anastasia Kuzminykh. 2023. The Bot on Speaking Terms: The Effects of Conversation Architecture on Perceptions of Conversational Agents. In *Proceedings of the 5th International Conference on Conversational User Interfaces (CUI '23)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3571884.3597139>
- [44] Katharina Weitz, Dominik Schiller, Ruben Schlagowski, Tobias Huber, and Elisabeth André. 2019. “Do you trust me?”: Increasing User-Trust by Integrating Virtual Agents in Explainable AI Interaction Design. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents (IVA '19)*. Association for Computing Machinery, New York, NY, USA, 7–9. <https://doi.org/10.1145/3308532.3329441>
- [45] Katharina Weitz, Dominik Schiller, Ruben Schlagowski, Tobias Huber, and Elisabeth André. 2021. “Let me explain!”: exploring the potential of virtual agents in explainable AI interaction design. *Journal on Multimodal User Interfaces* 15, 2 (June 2021), 87–98. <https://doi.org/10.1007/s12193-020-00332-0>
- [46] Nima Zargham, Michael Bonfert, Robert Porzel, Tanja Doring, and Rainer Malaka. 2021. Multi-Agent Voice Assistants: An Investigation Of User Experience. In *20th International Conference on Mobile and Ubiquitous Multimedia (Leuven, Belgium) (MUM 2021)*. Association for Computing Machinery, New York, NY, USA, 98–107. <https://doi.org/10.1145/3490632.3490662>
- [47] Nima Zargham, Leon Reicherts, Michael Bonfert, Sarah Theres Voelkel, Johannes Schoening, Rainer Malaka, and Yvonne Rogers. 2022. Understanding Circumstances for Desirable Proactive Behaviour of Voice Assistants: The Proactivity Dilemma. In *Proceedings of the 4th Conference on Conversational User Interfaces (Glasgow, United Kingdom) (CUI '22)*. Association for Computing Machinery, New York, NY, USA, Article 3, 14 pages. <https://doi.org/10.1145/3543829.3543834>