



King's Research Portal

DOI:

[10.1145/3571884.3603760](https://doi.org/10.1145/3571884.3603760)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Seymour, W., Zhan, X., Côté, M., & Such, J. (2023). Who are CUIs Really For? Representation and Accessibility in the Conversational User Interface Literature. In *Proceedings of the 5th International Conference on Conversational User Interfaces, CUI 2023* Article 26 (Proceedings of the 5th International Conference on Conversational User Interfaces, CUI 2023). <https://doi.org/10.1145/3571884.3603760>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Who are CUIs Really For? Representation and Accessibility in the Conversational User Interface Literature

William Seymour

william.1.seymour@kcl.ac.uk
King's College London
London, UK

Mark Coté

mark.cote@kcl.ac.uk
King's College London
London, UK

Xiao Zhan

xiao.zhan@kcl.ac.uk
King's College London
London, UK

Jose Such

jose.such@kcl.ac.uk
King's College London
London, UK

ABSTRACT

The theme for CUI 2023 is ‘designing for inclusive conversation’, but who are CUIs really designed for? The field has its roots in computer science, which has a long acknowledged diversity problem. Inspired by studies mapping out the diversity of the CHI and voice assistant literature, we set out to investigate how these issues have (or have not) shaped the CUI literature. To do this we reviewed the 46 full-length research papers that have been published at CUI since its inception in 2019. After detailing the eight papers that engage with accessibility, social interaction, and performance of gender, we show that 90% of papers published at CUI with user studies recruit participants from Europe and North America (or do not specify). To complement existing work in the community towards diversity we discuss the factors that have contributed to the current status quo, and offer some initial suggestions as to how we as a CUI community can continue to improve. We hope that this will form the beginning of a wider discussion at the conference.

CCS CONCEPTS

- **Security and privacy** → **Usability in security and privacy**;
- **Human-centered computing** → **HCI theory, concepts and models**; *Natural language interfaces*.

KEYWORDS

CUI, Inclusivity, Diversity, Conversational Assistants, Voice Assistants

ACM Reference Format:

William Seymour, Xiao Zhan, Mark Coté, and Jose Such. 2023. Who are CUIs Really For? Representation and Accessibility in the Conversational User Interface Literature. In *ACM conference on Conversational User Interfaces (CUI '23)*, July 19–21, 2023, Eindhoven, Netherlands. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3571884.3603760>

1 INTRODUCTION

Conversational assistants and chatbots continue to increase in popularity, leveraging advancements in machine learning to improve reasoning, natural language processing, and voice recognition. Since their inception, an interdisciplinary research community has grown around voice and conversational technologies, beginning in computer science with the likes of ELIZA [26]. It has since expanded to encompass a wide array of disciplines, and this is the fifth year that we have congregated under the CUI banner.

It has long been recognised that both computer science as a discipline and technology as a sector have major diversity problems, which may even be worsening over time¹ through a vicious cycle of underrepresentation and underenrollment that is compounded by structural inequality, bias, and harassment. The multiplicative power of design means that this white male homogeneity extends to the products we use every day, leading to technology that isn't designed for the marginalised majority. It is unsurprising then that in the home we see that the installation and abuse [16] of connected devices has a significant gender component, and that video doorbells reinforce existing racism.²

But what about research? An analysis of CHI papers from 2016–2020 found that 73% of CHI studies recruited Western participant samples representing less than 12% of the world's population; this reflects the organisations that submit to CHI, as over 80% of papers used participants from the country of authors' institutions [14]. Inspired by this we have previously investigated these issues in the context of voice assistants and found a similar story. As part of a larger systematic review of ethical concerns with voice assistants, we coded papers for participant demographics and analysed the community's response to issues of equity including accessibility and performance of gender [23]. In this provocation we extend this analysis to CUI papers to address an important aspect of the conference theme ‘designing for inclusive conversation’. To do this we briefly outline the findings from the voice assistant review before discussing those findings in relation to the papers published at CUI from the inaugural conference in 2019 through to 2022.

CUI '23, July 19–21, 2023, Eindhoven, Netherlands

© 2023 Copyright held by the owner/author(s).

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM conference on Conversational User Interfaces (CUI '23)*, July 19–21, 2023, Eindhoven, Netherlands, <https://doi.org/10.1145/3571884.3603760>.

¹<https://www.wired.com/story/computer-science-graduates-diversity>

²<https://www.vice.com/en/article/qvyvzd/amazons-home-security-company-is-turning-everyone-into-cops>

2 THE SURVEY

We conducted a systematic literature review of the literature on ethical concerns with voice assistants, analysing 117 items published since 2012 [23]. We adopted a broad view of ethics and related concerns, using the definition of “what a design object ought to be based on ethical and moral codes”, in contrast with its purpose/function (reason), and visual values/presentation (aesthetics) [9]. The scope of the review included journal articles and conference papers indexed by the ACM Digital Library, IEEE Explore, Web of Science, and DBLP.

The majority of manuscripts that were in-scope for the review, and thus its major findings, were centred around privacy, social interaction, accessibility, performance of gender, and social order. We briefly summarise these here, going into more depth for topics of particular relevance to the provocation:

Privacy. Users’ privacy concerns related to voice assistants, including data collection by vendors, sharing of devices between users, and access to devices from those outside the household/home

Social Interaction. Often built on anthropomorphism and the Computers as Social Actors paradigm, these papers explored whether voice assistants are ontologically distinct from people and/or machines, how they conceptualise them, as well as the emotional connections and relationships users form with them

Social Order. How voice assistants affect the dynamics of domestic hierarchies, including parent-child relationships

Accessibility. CUIs are interestingly situated in terms of accessibility. On the one hand, their text/language based conversational nature makes them accessible to the blind and partially sighted [1, 20]. On the other hand, for verbal systems those with speech, language, and hearing difficulties are less adequately served [18], as well as those who have difficulties formulating requests in the manner and speed required [7, 15, 18, 22]. Speech recognition accuracy varies across languages, accents, gender, and other demographics, although this is often framed as a challenge of engineering rather than social inequality [13].

Performance of Gender. In many products and in popular culture, the default persona for a conversational user interface might be characterised as a white middle-class American woman, with vendors justifying this based on user acceptability with studies suggesting that people perceive women’s voices and names as sounding more gentle, kind, and caring than men’s [5, 21]. It has been known since the 90’s that people apply gender and other stereotypes to computers [17], and as a research community we have long accepted that we have a duty to push back against harmful preconceptions around gender roles in society. At the same time, attempts to create ‘genderless’ presentations often do not stop users gendering conversational interfaces [25] and can (ironically) erase or ignore those outside the gender binary [19].

2.1 Additional Analysis

We also coded papers for the country of residence of their participants (shown in Table 1) and the specific groups demographics they recruited (shown in Table 2) where applicable. Of the 85 papers with human participants, **94% drew solely or mostly from North**

Table 1: Participant country of residence for papers with human participants.

Country	Papers	Country	Papers
USA	47	Netherlands	2
UK	11	Ireland	1
Germany	10	South Korea	1
Italy	4	Spain	1
India	2	Sweden	1
Brazil	2	Switzerland	1
Canada	2		

Table 2: Number of papers recruiting specific groups, written as reported.

Sub-population	# of Papers
None	55
Children	10
Blind and visually impaired	3
People with Dysarthria	2
Older adults	2
Trans and non-binary	1
Users with disabilities	1
Users with motor impairments	1
Portuguese Speakers	1
Young Adults (18-36)	1

America and Europe. Of the 22 papers that explored the experiences of a particular group of people, 13 were focused on age, 7 on (dis)ability, and 1 each on trans/non binary and Portuguese speakers. Together these represented just over one quarter of the studies that recruited human participants. Within the papers concerned with age there was a diversity of concerns across anthropomorphism (5), accessibility (4), social order (2), and accountability (1). In contrast, all of the papers that recruited based on ability were focused on accessibility concerns. Studies looking at specific groups mainly took qualitative approaches (70%), suggesting a focus on lived experiences rather than the gathering of a high level overview.

3 BUT IT’S A CHI PROBLEM, RIGHT?

We found that 94% of participants in papers *investigating ethics* were from Europe or North America. That’s bad, but those papers were mainly from CHI and CSCW; only 3 of the 117 papers were published at CUI. It’s a fair criticism that the timing of the survey ruled out papers from CUI 2022, and that the only captured papers on voice assistants—this ignores the fact that CUI is broader than just voice assistants, with the conference encompassing a host of other conversational interfaces.

To reflect and address this, we conducted a (less rigorous) sift of the 46 full papers that have been published at CUI since its inception. Using the conference proceedings in the ACM Digital Library, we analysed paper abstracts, introductions, and methodologies, comparing them against the themes and demographics from the original review paper. In total, we found eight papers (17%) exploring topics

Table 3: Majority participant region of residence for CUI papers. For one paper at each of CUI 2020 and CUI 2022 it was not possible to ascertain the participants' region.

CUI	Europe	N. America	S. America	Other	Multiple
2019	6	2	1	0	0
2020	6	2	1	0	0
2021	4	6	0	0	1
2022	5	5	0	0	1
Total	21	15	2	0	2

around accessibility [4, 8, 27], social interaction [6, 11, 12], and the performance of gender by CUIs [10, 25]. By contrast, around half of the provocation papers at CUI 2022 explored questions related to ethics.

For accessibility, CUI papers cover similar themes to the voice assistant review, focusing on enabling children to interact with CUIs on their own terms [4, 8]. They also analysed command construction in native and non-native speakers [27]. All accessibility papers used voice as the mode of interaction. The papers on social interaction were also split between categories encountered during the voice assistant literature survey. Two focused on expanding communication with CUIs to encompass speech acts and emotions [11, 12], and one on the effect of accented synthetic speech on peoples' responses [6]. Finally, the two papers on performance of gender explored sensitivity when designing gender for voice interfaces [25] and user acceptability of speech with varying gender-linked characteristics [10].

In terms of participants, **90% of CUI papers with user studies recruited participants solely residing in Europe and North America or did not report where participants were recruited from** (38 of 42 papers), with the two papers studying participants from South America sharing common authors (the only studies focusing on locations outside of Europe and North America) [2, 3]. The breakdown of participant location by year of the conference is given in Table 3. As a general rule, the geographic location of participants matched that of paper authors. The size of the sample meant that there were no meaningful trends in terms of distribution, apart from a balancing of papers from Europe and the US in more recent years of the conference.

The objects of study in CUI papers are often (implicitly) WEIRD (Western, Educated, Industrialised, Rich, and Democratic)—where commercially available technologies formed the basis of experiments these were overwhelmingly produced in Europe or North America, and studied for English language tasks. Whilst conducting the literature survey we began to compile a list of global voice assistant launches to demonstrate the diversity of the field (inevitably incomplete, shown in table 4); despite the existence of many non-WEIRD CUIs, they remain dramatically unstudied by CUI papers.

In sum we see that CUI as a conference/community is one that is willing to engage with the ethical aspects of the technologies being studied, but less willing to consider CUI design and evaluation outside of its own geographic context. Provocations are much more likely to address issues of ethics, but these do not seem to be further developed into full papers.

Table 4: Global timeline of voice assistant releases.

Name	Language(s)	Launch
Siri	Various (20)	2011
Alexa	Various (9)	2013
Cortana	Various (9)	2014
NUGU	Korean	2016
Mycroft	English	2016
Clova	Japanese & Korean	2017
Bixby	Various (8)	2017
Kakao Mini	Korean	2017
Xiao AI	Chinese (Mandarin & Cantonese)	2017
Xiaoyi	Chinese (further details unknown)	2019
Celia	Various (4)	2020

4 WHAT CAN WE DO

There are many forces that shape the design of CUIs and who they serve. Vendors are motivated by profit at the expense of accessibility, data protection regulations designed to protect individuals from organisations indirectly limit multi-user usage, design teams are more aware of issues faced by people like themselves, and the lack of high level perspective results in decisions around e.g. performance of gender that perpetuate existing societal biases. We might not be able to single handedly bring down capitalism, but we *can* alter the research landscape that feeds into industrial research and development and will inspire the next generation of researchers. A key aspect of diversity of geography, culture, gender, ability, and others in the research process is how it shapes *who CUIs are for*. The lack of visibility and consideration for marginalised groups feeds back into who uses and benefits from CUIs, as well as who goes on to research and design them.

Looking at the causes of the status quo, the situation is understandable from the perspective of an individual researcher. The contextual nature of many of the topics studied by the community (e.g. social interactions and social order) mean that participants must often be drawn from the same culture; widening participation is therefore not simply a matter of balancing participant demographics, studies need to replicated (and potentially also localised) across countries. This conflicts with the pressure to publish—if a paper meets the bar for publication with a single participant group, there is little benefit to replicating it in terms of the metrics by which academics are judged.

Looking back at our own papers when writing this, we found ourselves coming up with excuses. There were prototypes that need to be tested in person, topics that were deeply affected by culture, or language barriers our team simply was not equipped to handle and that local research teams would be. But the reviews by Linxen et al. and ourselves of CHI/voice assistant publications show us that these arguments apply to at least *three quarters of the community*. We can't rely on authors in non-WEIRD countries to generate the diverse research that we need because those authors *aren't submitting to CUI* (or CHI, or CSCW). Much easier, perhaps, to use provocation papers to point out the inequities that we cannot or will not address in our research (exactly as this paper does)?

At a conference level this begins to more clearly conflict with the stated values of the community. To facilitate attendance from a more diverse group of researchers, CUI offers a 25% reduction in the conference fee for attendees who live in economically developing countries and allows for remote attendance (CHI similarly offers 35% and 80% reductions and virtual attendance). Travel costs, however, are often considerably more than the conference registration, meaning that facilitating virtual attendance is likely to be more effective in terms of diversifying attendance. At the same time, virtual attendance does not present the same quality of conference experience, particularly around informal discussions, networking opportunities, and community building, which risks creating a two tier system that does not address the underlying structural problem. CUI attendees can make use of Gary Marsden Travel Awards to fund in-person travel, but only *once every three years* for physical attendance. More than two thirds of recipients so far from 2023 have been from Europe or North America.

So what else can we do, given that the theme of CUI 2023 is ‘designing for inclusive conversation’? The surveys suggest that measures focusing on the authors of publications will in turn increase the diversity of participants, as researchers tend to recruit the people around them. Additional financial support for registration and travel for attendees outside WEIRD countries is a key part of this, which could be supported via increased professional registration fees. Other potential options include having parallel submission tracks for more diverse/underrepresented research, as well as to encourage the replication of previously published studies across communities and cultures. This may better align the incentives of researchers, who are pushed to publish, with those of the conference and larger community, who recognise the value of more diverse research and community.

We recognise that these are not perfect answers; to an extent these problems are systemic and a result of the global academic system prioritising North America and Europe (CUI, for example, meets in Europe and uses English as its working language). We are interested to hear suggestions at the conference as to how we might work against this, particularly in ways that align with our other values like sustainability.

5 DISCUSSION AND REFLECTION

The discussion and reflections presented here draw from reviews received by the paper, in order to offer a starting point for further discussion at the conference. As per ACM policy we do not directly report the content of reviews, but rather general arguments and suggestions from our reviewers.³

5.1 Diversity Is More Than an Institution

A key issue with trying to ‘measure’ diversity using information available in research papers is that data provided on author affiliations and reported participant demographics is limited and often not directly related to what many would consider to be diversity in this context. Using e.g. the country of an author’s institution as a proxy for diversity of culture, background, or ethnicity reduces and shifts the scope of what is implied to be diverse. Analysing participant country of residence similarly glosses over the many

different types of diversity that exist within countries, portraying them as homogeneous groups and at a higher level implying that research carried out in a WEIRD context is inherently not diverse. One of the issues at play is the lack of precision around the different kinds of diversity being explored; this can lead to the conflation of apparently similar definitions as well as the reporting of very specific concepts as all-encompassing (e.g. ‘majority country of residence of the participants on the platform used by a study’ as just ‘diversity of participants’).

5.2 Diversity Is Not a New Problem

A recurring comment in the reviews was that, as we point out in Section 1, while these issues might not be well represented *at CUI* these kinds of diversity are not a new or unknown problem for HCI and its associated fields of study. There is a rich history of work exploring scientific colonialism and subsequent decolonisation movements that are already present in adjacent conferences like CHI. Indeed, since the original submission of this provocation there has been the CUI@CHI '23 workshop which focused on the inclusive design of CUIs, including submissions relating to global inclusivity [24]. We are delighted to see these discussions beginning to spread from venues where they are already established.

5.3 Diversity is a Complex Issue

It was rightly pointed out that this is a very complicated issue with many different components, and that our attempt to offer suggestions in Section 4 smoothed over many of these complexities. Fostering global inclusivity for our research community and those who participate in its research also has political aspects that we did not address: e.g. the structure of the peer-review process inherently makes it harder to publish work that does not conform to WEIRD agendas and methodologies, as the reviewer pool is primarily made up of WEIRD reviewers. Many of the barriers to inclusivity at the level of a conference are often bureaucratic or budgetary. While things may be changing at the top of SIGCHI and ACM, the pace of that change remains very slow. However, this does not mean that we should just wait for change to manifest, and we hope that despite its imperfections this work will be able to provoke discussion and ideas for change at the conference.

6 CONCLUSION

Inspired by work that has mapped out the diversity of the CHI and voice assistant literature, we conducted an informal review of the CUI literature in order to explore how the community approaches issues of ethics and diversity in terms of research, researchers, and participants. While we found a variety of papers covering accessibility, social interaction, and social order, the participants recruited for user studies were overwhelmingly from Europe and North America. We outline contributing factors to the problem of diversity in CUI research and in HCI more broadly, suggesting some starting points for improvement with the aim of seeding a larger discussion at the conference.

ACKNOWLEDGMENTS

This research was funded by the UK Engineering and Physical Sciences Research Council under grant EP/T026723/1.

³<https://www.acm.org/publications/policies/roles-and-responsibilities>

REFERENCES

- [1] Ali Abdolrahmani, Ravi Kuber, and Stacy M. Branham. 2018. "Siri Talks at You": An Empirical Investigation of Voice-Activated Personal Assistant (VAPA) Usage by Individuals Who Are Blind. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (Galway, Ireland) (ASSETS '18). Association for Computing Machinery, New York, NY, USA, 249–258. <https://doi.org/10.1145/3234695.3236344>
- [2] Fabricio Barth, Heloisa Candello, Paulo Cavalin, and Claudio Pinhanez. 2020. Intentions, Meanings, and Whys: Designing Content for Voice-Based Conversational Museum Guides. In *Proceedings of the 2nd Conference on Conversational User Interfaces* (Bilbao, Spain) (CUI '20). Association for Computing Machinery, New York, NY, USA, Article 8, 8 pages. <https://doi.org/10.1145/3405755.3406128>
- [3] Heloisa Candello, Claudio Pinhanez, Mauro Pichiliani, Marisa Vasconcelos, and Haylla Conde. 2019. Can Direct Address Affect User Engagement with Chatbots Embodied in Physical Spaces?. In *Proceedings of the 1st International Conference on Conversational User Interfaces* (Dublin, Ireland) (CUI '19). Association for Computing Machinery, New York, NY, USA, Article 3, 9 pages. <https://doi.org/10.1145/3342775.3342787>
- [4] Fabio Catania, Micol Spitale, Giulia Cosentino, and Franca Garzotto. 2020. What is the Best Action for Children to "Wake Up" and "Put to Sleep" a Conversational Agent? A Multi-Criteria Decision Analysis Approach. In *Proceedings of the 2nd Conference on Conversational User Interfaces* (Bilbao, Spain) (CUI '20). Association for Computing Machinery, New York, NY, USA, Article 4, 10 pages. <https://doi.org/10.1145/3405755.3406129>
- [5] Pedro Costa and Luísa Ribas. 2019. AI becomes her: Discussing gender and artificial intelligence. *Technoetic Arts* 17, 1-2 (2019), 171–193.
- [6] Benjamin R. Cowan, Philip Doyle, Justin Edwards, Diego Garaialde, Ali Hayes-Brady, Holly P. Branigan, João Cabral, and Leigh Clark. 2019. What's in an Accent? The Impact of Accented Synthetic Speech on Lexical Choice in Human-Machine Dialogue. In *Proceedings of the 1st International Conference on Conversational User Interfaces* (Dublin, Ireland) (CUI '19). Association for Computing Machinery, New York, NY, USA, Article 23, 8 pages. <https://doi.org/10.1145/3342775.3342786>
- [7] Yao Du, Kerri Zhang, Sruthi Ramabadrán, and Yusa Liu. 2021. "Alexa, What is That Sound?" A Video Analysis of Child-Agent Communication From Two Amazon Alexa Games (IDC '21). Association for Computing Machinery, New York, NY, USA, 513–520. <https://doi.org/10.1145/3459990.3465195>
- [8] Layne Hubbard, Shanli Ding, Vananh Le, Pilyoung Kim, and Tom Yeh. 2021. Voice Design to Support Young Children's Agency in Child-Agent Interaction. In *Proceedings of the 3rd Conference on Conversational User Interfaces* (Bilbao (online), Spain) (CUI '21). Association for Computing Machinery, New York, NY, USA, Article 9, 10 pages. <https://doi.org/10.1145/3469595.3469604>
- [9] Rikke Hagensby Jensen, Yolande Strengers, Jesper Kjeldskov, Larissa Nicholls, and Mikael B. Skov. 2018. *Designing the Desirable Smart Home: A Study of Household Experiences and Energy Consumption Impacts*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3173574.3173578>
- [10] Iris Jestin, Joel Fischer, Maria Jose Galvez Trigo, David Large, and Gary Burnett. 2022. Effects of Wording and Gendered Voices on Acceptability of Voice Assistants in Future Autonomous Vehicles. In *Proceedings of the 4th Conference on Conversational User Interfaces* (Glasgow, United Kingdom) (CUI '22). Association for Computing Machinery, New York, NY, USA, Article 24, 11 pages. <https://doi.org/10.1145/3543829.3543836>
- [11] Minha Lee. 2020. Speech Acts Redux: Beyond Request-Response Interactions. In *Proceedings of the 2nd Conference on Conversational User Interfaces* (Bilbao, Spain) (CUI '20). Association for Computing Machinery, New York, NY, USA, Article 13, 10 pages. <https://doi.org/10.1145/3405755.3406124>
- [12] 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* (Glasgow, United Kingdom) (CUI '22). Association for Computing Machinery, New York, NY, USA, Article 19, 11 pages. <https://doi.org/10.1145/3543829.3543835>
- [13] Lanna Lima, Vasco Furtado, Elizabeth Furtado, and Virgilio Almeida. 2019. Empirical Analysis of Bias in Voice-Based Personal Assistants. In *Companion Proceedings of The 2019 World Wide Web Conference* (San Francisco, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 533–538. <https://doi.org/10.1145/3308560.3317597>
- [14] Sebastian Linxen, Christian Sturm, Florian Brühlmann, Vincent Cassau, Klaus Opwis, and Katharina Reinecke. 2021. *How WEIRD is CHI?* Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3411764.3445488>
- [15] Silvia Lovato and Anne Marie Piper. 2015. "Siri, is This You?": Understanding Young Children's Interactions with Voice Input Systems. In *Proceedings of the 14th International Conference on Interaction Design and Children* (Boston, Massachusetts) (IDC '15). Association for Computing Machinery, New York, NY, USA, 335–338. <https://doi.org/10.1145/2771839.2771910>
- [16] Dana McKay and Charlynn Miller. 2021. Standing in the Way of Control: A Call to Action to Prevent Abuse through Better Design of Smart Technologies. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 332, 14 pages. <https://doi.org/10.1145/3411764.3445114>
- [17] Clifford Nass, Jonathan Steuer, and Ellen R. Tauber. 1994. Computers Are Social Actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) (CHI '94). Association for Computing Machinery, New York, NY, USA, 72–78. <https://doi.org/10.1145/191666.191703>
- [18] Alisha Pradhan, Kanika Mehta, and Leah Findlater. 2018. "Accessibility Came by Accident": Use of Voice-Controlled Intelligent Personal Assistants by People with Disabilities. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3174033>
- [19] Cami Rincón, Os Keyes, and Corinne Cath. 2021. Speaking from Experience: Trans/Non-Binary Requirements for Voice-Activated AI. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 132 (apr 2021), 27 pages. <https://doi.org/10.1145/3449206>
- [20] Sergio Sayago and Mireia Ribera. 2020. Apple Siri (Input) + Voice Over (Output) = a de Facto Marriage: An Exploratory Case Study with Blind People. In *9th International Conference on Software Development and Technologies for Enhancing Accessibility and Fighting Info-Exclusion* (Online, Portugal) (DSAI 2020). Association for Computing Machinery, New York, NY, USA, 6–10. <https://doi.org/10.1145/3439231.3440603>
- [21] Florian Schneider. 2021. Recommended by Google Home. In *International Conference on Human-Computer Interaction*. Springer, 485–493.
- [22] Alex Sciuto, Armita 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* (Hong Kong, China) (DIS '18). Association for Computing Machinery, New York, NY, USA, 857–868. <https://doi.org/10.1145/3196709.3196772>
- [23] William Seymour, Xiao Zhan, Mark Cote, and Jose Such. 2023. A Systematic Review of Ethical Concerns with Voice Assistants. In *Accepted to the 2023 AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*. <https://arxiv.org/abs/2211.04193>
- [24] Jaisie Sin, Heloisa Candello, Leigh Clark, Benjamin R. Cowan, Minha Lee, Cosmin Munteanu, Martin Porcheron, Sarah Theres Völkel, Stacy Branham, Robin N. Brewer, Ana Paula Chaves, Razan Jabbar, and Amanda Lazar. 2023. CUI@CHI: Inclusive Design of CUIs Across Modalities and Mobilities. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI EA '23). Association for Computing Machinery, New York, NY, USA, Article 341, 5 pages. <https://doi.org/10.1145/3544549.3573820>
- [25] Selina Jeanne Sutton. 2020. Gender Ambiguous, Not Genderless: Designing Gender in Voice User Interfaces (VUIs) with Sensitivity. In *Proceedings of the 2nd Conference on Conversational User Interfaces* (Bilbao, Spain) (CUI '20). Association for Computing Machinery, New York, NY, USA, Article 11, 8 pages. <https://doi.org/10.1145/3405755.3406123>
- [26] Joseph Weizenbaum. 1966. ELIZA-A Computer Program for the Study of Natural Language Communication between Man and Machine. *Commun. ACM* 9, 1 (Jan. 1966), 36–45. <https://doi.org/10.1145/365153.365168>
- [27] Yunhan Wu, Martin Porcheron, Philip Doyle, Justin Edwards, Daniel Rough, Orla Cooney, Anna Bleakley, Leigh Clark, and Benjamin Cowan. 2022. Comparing Command Construction in Native and Non-Native Speaker IPA Interaction through Conversation Analysis. In *Proceedings of the 4th Conference on Conversational User Interfaces* (Glasgow, United Kingdom) (CUI '22). Association for Computing Machinery, New York, NY, USA, Article 10, 12 pages. <https://doi.org/10.1145/3543829.3543839>