

Speech AI for All: Promoting Accessibility, Fairness, Inclusivity, and Equity

Shaomei Wu
AImpower.org
Mountain View, California, USA
shaomei@aimpower.org

Kimi Wenzel
Human-Computer Interaction
Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
kwenzel@andrew.cmu.edu

Jingjin Li
AImpower.org
Mountain View, California, USA
jingjin@aimpower.org

Qisheng Li
AImpower.org
Seattle, Washington, USA
qishengli@aimpower.org

Alisha Pradhan
Department of Informatics, Ying Wu
College of Computing
New Jersey Institute of Technology
Newark, New Jersey, USA
alisha.pradhan@njit.edu

Raja Kushalnagar
Gallaudet University
Washington, District of Columbia
USA
raja.kushalnagar@gallaudet.edu

Colin Lea
Apple
Cupertino, California, USA
colin_lea@apple.com

Allison Koenecke
Cornell University
Ithaca, New York, USA
koenecke@cornell.edu

Christian Vogler
Technology Access Program
Gallaudet University
Washington, District of Columbia
USA
christian.vogler@gallaudet.edu

Mark Hasegawa-Johnson
The Grainger College of Engineering
University of Illinois
Urbana-Champaign
Urbana-Champaign, Illinois, USA
jhasegaw@illinois.edu

Norman Makoto Su
Department of Computational Media
University of California, Santa Cruz
Santa Cruz, California, USA
normsu@ucsc.edu

Nan Bernstein Ratner
Department of Hearing and Speech
Sciences
University of Maryland, College Park
College Park, Maryland, USA
Neuroscience and Cognitive Science
University of Maryland, College Park
College Park, Maryland, USA
nratner@umd.edu

Abstract

Trained and optimized for typical and fluent speech, speech AI works poorly for people with speech diversities, often interrupting them and misinterpreting their speech. The increasing deployment of speech AI in automated phone menus, AI-conducted job interviews, and everyday devices poses tangible risks to people with speech diversities. To mitigate these risks, this workshop aims to build a multidisciplinary coalition and set the research agenda for fair and accessible speech AI. Bringing together a broad group of academics and practitioners with diverse perspectives, including HCI, AI, and other relevant fields such as disability studies, speech language pathology, and law, this workshop will establish a shared

understanding of the technical challenges for fair and accessible speech AI, as well as its ramifications in design, user experience, policy, and society. In addition, the workshop will invite and highlight first-person accounts from people with speech diversities, facilitating direct dialogues and collaboration between speech AI developers and the impacted communities. The key outcomes of this workshop include a summary paper that synthesizes our learnings and outlines the roadmap for improving speech AI for people with speech diversities, as well as a community of scholars, practitioners, activists, and policy makers interested in driving progress in this domain.

Permission to make digital or hard copies of all or part 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 '25, Yokohama, Japan

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1395-8/25/04

<https://doi.org/10.1145/3706599.3706746>

CCS Concepts

• **Human-centered computing** → **Accessibility**; **Human computer interaction (HCI)**; • **Computing methodologies** → **Artificial intelligence**; • **Social and professional topics** → **Computing / technology policy**.

Keywords

AI FATE, automatic speech recognition, speech technology, disability, accessibility, speech diversity

ACM Reference Format:

Shaomei Wu, Kimi Wenzel, Jingjin Li, Qisheng Li, Alisha Pradhan, Raja Kushalnagar, Colin Lea, Allison Koenecke, Christian Vogler, Mark Hasegawa-Johnson, Norman Makoto Su, and Nan Bernstein Ratner. 2025. Speech AI for All: Promoting Accessibility, Fairness, Inclusivity, and Equity. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25), April 26–May 01, 2025, Yokohama, Japan*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3706599.3706746>

1 Motivation

Speech AI technologies have evolved and proliferated at an accelerating pace, transforming how we communicate and interact with others and the world. While advancements in automatic speech recognition (ASR) and speech synthesis have unleashed a wide range of applications, such as smart speakers, automated phone agents, AI interviewers, and in-car speech interfaces, the widespread adoption of these technologies also introduces barriers and fairness issues for people with speech diversities. Trained and optimized for “typical”¹ speech, existing ASR models have significant difficulty in processing speech with diverse patterns such as stuttering [13, 15], deaf speech [7], speech patterns of individuals experiencing age-related changes [19, 25], aphasia [11], second language speech [10, 27], as well as regional vernaculars and ethnic dialects [12]. As a result, ASR-powered systems often misinterpret the speech of people with speech diversities, interrupt them prematurely, or fail to provide accurate responses and transcriptions. The inability of speech AI systems to work with diverse speech not only creates additional barriers for people with speech diversities to interact with popular products and services, like personal voice assistance and automated phone menus, but may also lead to more serious psychological harms [5, 26] and reduced economic opportunities [16, 28].

In addition to the lack of inclusivity in recognizing and understanding diverse speech patterns of end users, researchers are drawing attention to the inherent inequalities within the current design of speech-based interfaces due to the voices these technologies represent: speech-based interfaces can reinforce “whiteness as the norm” through its default voice [4, 18, 20]. Artificial, yet, human-like speech, can encode attributes like gender, race, age, and class [24]. When these attributes are centered solely around dominant groups [6], as is often the case with commodity speech interfaces, they can have significant ramifications at a societal level.

The workshop will examine and address the accessibility and fairness of speech AI (both foundation speech models and the systems powered by them) for people with speech differences as well as encourage discussions on the equitable design of artificial speech. We aim to bring together researchers, practitioners, policy makers, and community advocates to form a cross-sector coalition for fair and accessible speech AI.

¹In this workshop, we define typical speech to include fluent speech, speech with a dominant accent (often embodied by those with higher socioeconomic classes), and dominant vernaculars.

1.1 Workshop Themes and Objectives

The workshop will center the lived experiences and expectations of those most affected by disparities in speech AI performance. Based on this grounding, participants will then have opportunities to present and discuss the technical and design challenges and opportunities for fair and accessible speech AI, as well as the norms and public policies that underpin these challenges and opportunities.

- **Understanding the Lived Experience**—We recognize the inherent relational nature of AI disparities and the epistemic privileges held by marginalized communities to identify and address harm [3]. Our workshop aims to elevate the voices of those most affected through our organizing team, invited speakers, and attendees. Accordingly, we will discuss the disparate cognitive and emotional burdens of speech AI [17, 26, 27, 29, 30], its social impacts on stigmatizing and suppressing speech diversities [9, 14], and its historical roots in structural inequalities [2, 23]. A further objective under this theme is to outline opportunities for allyship.
- **Showcasing Solutions**—Grounded in the experiences of affected users, our workshop will capture incremental solutions to speech AI inequities. This will include both technical and design approaches proposed by the research, non-profit, and industry communities. For example, we ask: What metrics for speech recognition systems are the most respectful and representative of the experiences of marginalized users? What is the most equitable approach to engage the affected communities in data collection and solution development? How do we balance long-term capacity building versus short-term band-aid solutions? Through presentations, demos, and discussion, this workshop will offer human-centered technical, measurement, and design recommendations.
- **Unpacking Norms and Policy**—While many developed countries have laws and regulations to protect marginalized groups against systematic inaccessibility and discrimination, current legal and policy frameworks have also fallen behind with the rapid development of AI technologies and new challenges they bring [22, 31]. Our workshop will attend to the structures and norms that underlie the challenges and progress in fair and accessible speech AI. We will discuss and propose policy recommendations - such as assessments and requirements for ASR performance parity for both typical and diverse speech, as well as advocacy avenues - such as academic conference accessibility guidelines that accommodate speech diversity.

2 Organizers

Shaomei Wu is a person who stutters and the founder of AIm-power.org – a tech non-profit researching and co-developing inclusive videoconferencing and speech AI technologies **for, with, and by** the stuttering community. Her research focuses on fairness and ethics in mainstream technologies, including social media, AI, and videoconferencing. She has co-organized workshops at CSCW, Meta, and National Stuttering Association’s annual conference and served as a workshop co-chair for ICWSM.

Kimi Wenzel is a PhD student at Carnegie Mellon University’s Human-Computer Interaction Institute. Her research centers on

understanding the downstream representational harms of speech AI.

Jingjin Li is a research fellow at AImpower.org where she leads research on co-designing inclusive Speech AI and videoconferencing tools with the stuttering community. She has co-organized panels and served on the organizing committee for CSCW and DIS.

Qisheng Li is a research engineer at Meta Reality Labs and a research fellow at AImpower.org where she leads the community-led Chinese stuttered speech dataset creation and benchmarking project. Her research interests lie in HCI and AI, with a emphasis on AI for social good, crowdsourcing, and evaluation. She has served on the program committee at ASSETS and CSCW.

Alisha Pradhan is an Assistant Professor in the Department of Informatics at the New Jersey Institute of Technology. Her research has examined design and use of conversational voice technologies by older adults, in particular identifying the accessibility benefits and barriers posed by these technologies, and engaging older adults in design of equitable voice technologies. She has co-organized workshops and panels at conferences including, ACM CHI and UbiComp.

Raja Kushalnagar is a Deaf Professor and Director of the Artificial Intelligence, Accessibility and Sign languages Center with deaf-accented speech. His research explores the accuracy and usability of conversational voice technologies by deaf and hard of hearing people, and fairness, equity and inclusion in AI models and platforms. He has served on the program committee at ASSETS, ICCHP and CSUN.

Colin Lea is a research scientist and manager at Apple. His group focuses on making interactive technologies — especially speech — more inclusive for people with disabilities. His work is at the intersection of machine learning, HCI, and accessibility and emphasizes data collection/curation and ML modeling.

Allison Koenecke is a mainstream American English speaker, and an Assistant Professor of Information Science at Cornell University. Her research on algorithmic fairness includes auditing disparities in ASR system performance, especially among underrepresented speech populations including African American English, d/Dhh, and aphasic speakers.

Christian Vogler is a deaf person who speaks English with an accent that derives both from his German roots and his deafness. He is the director of the Technology Access Program at Gallaudet University and has led numerous grants and projects on accessible technologies for the DHH. Some of his most recent work focuses on both text-to-speech and speech-to-text technologies for DHH people. He has served on the organizing committee of conferences and workshops, including ASSETS, Gesture Workshops, AI-related workshops, and others.

Mark Hasegawa-Johnson is a Professor of Electrical and Computer Engineering at University of Illinois Urbana-Champaign. He is a Fellow of the Acoustical Society of America, a Fellow of ISCA, and a Fellow of the IEEE for contributions to speech processing of under-resourced languages. Professor Hasegawa-Johnson is Editor-in-Chief-elect of the IEEE Transactions on Audio, Speech, and Language, member of the ISCA Diversity Committee, and Technical Program Chair of IEEE ASRU 2025.

Norman Makoto Su is an associate professor in the Department of Computational Media at UC Santa Cruz. His research interests lie in

human-computer interaction (HCI) and computer-supported cooperative work (CSCW). He directs the Authentic User Experience Lab (AUX Lab), where they integrate empirical and humanistic methods in HCI to study and design with subcultures. He has published work on collective action and around new forms of and challenges with AI work and the techlash. He has co-organized workshops at CHI, CSCW, and GROUP.

Nan Bernstein Ratner is a professor in the Department of Hearing and Speech Sciences at the University of Maryland, College Park. Her primary areas of research are fluency development and disorder, psycholinguistics, and child language development. Nan is a co-founder and co-manager of FluencyBank [21], a corpus of annotated disfluent speech that has been highly influential in both fluency research and speech AI development. She is a longstanding organizer of sessions for the Annual Meeting of the American Association for the Advancement of Science and the recipient of multiple NIH Conference grants (R15).

3 Pre-Workshop Plans

3.1 Target Audience

One overarching goal of this workshop is to bring together researchers, practitioners, and stakeholders from various disciplines and backgrounds to jointly develop a research agenda and a cross-sector coalition. To achieve this, we will leverage the extended networks of the workshop organizers to broadly advertise the workshop to a wide range of potential audiences, including academic institutions, tech companies, community organizations, professional organizations, and policy/legal agencies. The workshop's Call for Participation will be shared in emails, social media, and offline events.

We have also created a website at speechai4all.org to host relevant information about the workshop, including the organizers' bios and background literature on this topic, allowing interested participants to familiarize themselves with the workshop team and the topics we aim to explore.

Interested participants will be invited to submit an application form, sharing their disciplinary backgrounds and motivation to participate, as well as an optional position paper, technical paper, tech demo, design artifact, and/or policy brief. Participants will be selected based on the relevance of their disciplinary backgrounds, the synergy of their motivations, and the quality of submitted materials, if any.

Expected size of attendance: we expect 20 to 30 in-person attendees and around 10 remote attendees.

3.2 Plans to Publish Workshop Proceedings

As the workshop aims to build a cross-sector, interdisciplinary coalition around fair and accessible speech AI, we encourage and anticipate the participation of stakeholders from non-HCI and non-academic fields. Therefore, we do not require the submission of a paper and expect only a small number of papers presented at the workshop. Participants will be encouraged to share their papers and artifacts on open access platforms such as arXiv and GitHub. With permission, accepted papers and artifacts will also be published on the workshop website, along with recorded presentations.

4 Workshop Structure

4.1 Workshop Format

4.1.1 In-person and Hybrid Plan. We plan for a one-day, in-person and hybrid workshop comprised of varied activities, including micro-keynote provocations, community panel discussions, presentations of position papers, and themed small group discussions to explore the challenges and opportunities for fair and accessible speech AI. Most of the planned activities will be conducted in person to facilitate in-depth interactions and networking among participants. However, as one of our goals is to build a movement and coalition for inclusive speech AI, we will also invite people with speech diversities and domain experts from non-HCI areas (e.g., speech language pathologists, policy makers) as “micro-keynote” speakers and panelists. While we will seek sponsorship for travel grants to assist micro-keynote speakers and panelists to attend the workshop in person, we will also provide the option for them to participate remotely to maximize inclusivity.

The remote speakers and panelists will participate through the Zoom videoconferencing platform and real-time communication channels such as Slack. Zoom meeting links and Slack invitations will be sent out to all workshop attendees at least seven days prior to the workshop to provide opportunities for the attendees to introduce themselves in the Slack channel before the workshop.

4.1.2 Accessibility. Given the workshop’s emphasis on accessibility, fairness, inclusivity, and equity in speech AI, we will welcome and expect attendees with speech diversities, such as deaf and hard of hearing (DHH) individuals [8] and people who stutter. Sign language interpreters and live transcriptions will be requested and arranged to facilitate the full participation of all attendees in our workshop. Additionally, we will implement the following measures to accommodate diverse communication needs [1]: regulate turn-taking using both physical hand-raising and the virtual hand-raise button to prevent interruptions; reach out to all participants in advance to assess their specific needs and preferences; establish clear ground rules at the start, emphasizing respectful communication, speaking one at a time, and utilizing non-verbal channels like chat, emojis, or written notes to ensure inclusive engagement for both online and in-person participants.

4.1.3 Asynchronous Engagement. Asynchronous materials and content generated from workshop activities will be posted and regularly updated on the workshop website to allow remote access and serve as an informational portal after the workshop. A dedicated Slack workspace will be set up at least one week prior to the workshop to support asynchronous interactions before, during, and after the workshop.

4.2 Workshop Activities

4.2.1 Introduction. We will start the day by introducing the structure and goals of the workshop and engaging the participants in an icebreaker to get to know each other. The icebreaker activities will be designed with consideration for the participants’ backgrounds and accommodation needs, including the option to participate non-verbally.

4.2.2 Micro-keynotes. Understanding the challenge with speech AI is not only a technical problem. We will bring in the perspectives from stakeholders and experts outside of HCI for three 20-minute “micro-keynote” provocation talks, followed by discussions and Q&A. We plan to invite experts from relevant fields such as disability studies, speech therapy, and policy to facilitate cross-pollination and ground the follow-up discussion in holistic approaches.

4.2.3 Community Panel. Different from the conventional, expert-led model in most academic workshops, this workshop will actively invite and highlight voices from the communities most impacted by speech AI technologies. People with speech diversities will share their first-person accounts of speech AI experiences in the community panel as the foundation for potential technical, design, and policy interventions proposed at the workshop. A list of questions for panelists will be pre-solicited from workshop participants before the workshop. Ample time will also be reserved for participants to have direct conversations with the panelists as well.

4.2.4 Paper Presentation. Participants who submitted position papers, technical papers, tech demos, design artifacts, and policy papers will have the opportunity to present their work during the workshop. While we expect fewer than 20 papers for the two paper presentation sessions, we are also prepared to switch to a poster presentation format if we receive a large number of high-quality submissions.

4.2.5 Themed Small Group Discussion. We will prepare a list of discussion topics in advance, along with interesting topics that emerge during the workshop, to facilitate the formation of small group discussions. Each group will select a moderator to lead a one-hour, structured discussion on one or two specific topics assigned to the group and then reconvene with the entire workshop to share their insights. The learnings and insights from these group discussions will form the basis of a white paper that summarizes our collective knowledge on this domain and proposes a research agenda and design guidelines for fair and accessible speech AI. Workshop participants will be invited to co-author the white paper after the workshop.

A structured timeline for proposed activities can be found in Table 1.

5 Post-Workshop Plans

We will close the workshop with a discussion that reflects our collective learnings and builds a plan to create momentum for fair and inclusive speech AI for all.

The participants will be encouraged to publish their optional papers on the workshop’s website for archival and visibility purposes. They will have the opportunity to co-author a white paper that summarizes the insights from the workshop and provides a research roadmap for this topic. The white paper will be published on the workshop’s website and submitted to ACM Interactions. The participants will also be invited to form a cross-sector coalition, organized through a mailing list as well as regular virtual meet-ups, to drive forward the roadmap we built at the workshop.

Collectively, the coalition will also carry out follow-up activities such as identifying funding and collaboration opportunities (e.g. joint grant proposals, shared projects), planning for additional

Table 1: Proposed structure of the workshop

Time	Format	Activity
9:00 - 9:30 AM	In-person	Welcome & Introduction: Organizers will introduce the workshop, outline its goals, and engage participants in icebreaker activities.
9:30 - 11:00 AM	Hybrid	Expert Micro-Keynote Provocations: Three 20-minute keynotes from experts beyond the HCI community, followed by a 30-minute Q&A session.
11:00 - 12:00 PM	In-person	Paper Presentations: Participants will present accepted papers, followed by Q&A.
12:00 - 1:00 PM		Lunch Break
1:00 - 2:00 PM	In-person	Paper Presentations: Continued presentations, followed by Q&A.
2:00 - 3:00 PM	Hybrid	Community Panel: First-person testimony and direct engagement with communities impacted by speech AI.
3:00 - 3:15 PM		Break
3:15 - 4:30 PM	In-person	Themed Discussion: Participants first form small groups around interested topics solicited from workshop participants (15 mins), then engage in in-depth, structured discussions within groups (1 hour).
4:30 - 5:00 PM	In-person	Group Sharing & Wrap-Up: Participants form different small groups and share their discussion highlights with the whole workshop.
5:00 - 5:30 PM	In-person	Closing Remarks

workshops on fair and accessible speech AI at CHI or other venues (e.g. NeurIPS, CSCW, FAccT), organizing special issues in relevant journals, and establishing the infrastructure for industry-academic-community partnerships.

6 Call for Participation

This one-day, in-person, and hybrid CHI 2025 workshop invites researchers, practitioners, policy makers, and community members interested in fair and inclusive speech AI technologies to explore the challenges, impact, and opportunities of speech AI for people with speech diversities.

Participants will have the opportunity to directly engage with and learn from impacted communities and experts from technical and non-technical fields to build a deeper understanding of the challenges and opportunities for fair and inclusive speech AI. Leveraging the collective knowledge shared during the workshop, participants will co-create a roadmap for fair and inclusive speech AI, driven by a cross-sector coalition of scholars and stakeholders formed through the workshop.

Participants are encouraged, but not required, to submit position papers, tech demo papers, technical research papers, policy papers, or experience reports and briefs from other fields, under the general theme of understanding and improving speech AI technologies for and with people with speech diversities. Accepted papers will be presented during the workshop as oral or poster presentations, and published on the workshop's website.

Interested participants can apply by completing a digital form to share their background and motivation for participation, with an option to attach a paper of up to ten pages (plus references) in the ACM single-column format.

The workshop organizing team will select participants based on the alignment between the participants' backgrounds and the workshop themes, while striving to assemble a diverse group across a range of disciplines, methodologies, and seniorities.

Acknowledgments

This work is partially supported by the National Science Foundation under Grant No. 2427710. We thank Apple Inc. for sponsoring this workshop at CHI 2025.

References

- [1] Almpower.org. 2023. Inclusive Video Conferencing Recommendations. https://aimpower.org/wp-content/uploads/2023/08/inclusive_videoconferencing_recommendations.pdf Accessed: 2024-09-23.
- [2] A. Baker-Bell. 2020. *Linguistic Justice: Black Language, Literacy, Identity, and Pedagogy*. Taylor & Francis. <https://books.google.com/books?id=uTXfDwAAQBAJ>
- [3] Abeba Birhane. 2021. Algorithmic injustice: a relational ethics approach. *Patterns* 2, 2 (2021).
- [4] Robin N. Brewer, Christina Harrington, and Courtney Heldreth. 2023. Envisioning Equitable Speech Technologies for Black Older Adults. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 379–388. doi:10.1145/3593013.3594005
- [5] Geoffrey A. Coalson, Alexis Crawford, Shanley B. Treleaven, Courtney T. Byrd, Lauren Davis, Lillian Dang, Jillian Edgerly, and Alison Turk. 2022. Microaggression and the adult stuttering experience. *Journal of Communication Disorders* 95 (2022), 106180. doi:10.1016/j.jcomdis.2021.106180
- [6] Sheena Erete, Aarti Israni, and Tawanna Dillahunt. 2018. An intersectional approach to designing in the margins. *Interactions* 25, 3 (April 2018), 66–69. doi:10.1145/3194349
- [7] Abraham Glasser. 2019. Automatic Speech Recognition Services: Deaf and Hard-of-Hearing Usability. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3290607.3308461
- [8] Abraham Glasser, Kesavan Kushalnagar, and Raja Kushalnagar. 2017. Deaf, Hard of Hearing, and Hearing Perspectives on Using Automatic Speech Recognition in Conversation. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (Baltimore, Maryland, USA) (ASSETS '17). Association for Computing Machinery, New York, NY, USA, 427–432. doi:10.1145/3132525.3134781
- [9] Christina N. Harrington, Radhika Garg, Amanda Woodward, and Dimitri Williams. 2022. "It's Kind of Like Code-Switching": Black Older Adults' Experiences with a Voice Assistant for Health Information Seeking. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 604, 15 pages. doi:10.1145/3491102.3501995
- [10] Seung-Eun Kim, Bronya R. Chernyak, Olga Seleznova, Joseph Keshet, Matthew Goldrick, and Ann R. Bradlow. 2024. Automatic recognition of second language speech-in-noise. *JASA Express Letters* 4, 2 (02 2024), 025204. doi:10.1121/10.0024877 arXiv:https://pubs.aip.org/asa/jel/article-pdf/doi/10.1121/10.0024877/19669029/025204_1_10.0024877.pdf

- [11] Allison Koenecke, Anna Seo Gyeong Choi, Katelyn X. Mei, Hilke Schellmann, and Mona Sloane. 2024. Careless Whisper: Speech-to-Text Hallucination Harms. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 1672–1681. doi:10.1145/3630106.3658996
- [12] Allison Koenecke, Andrew Nam, Emily Lake, Joe Nudell, Minnie Quartey, Zion Mengesha, Connor Toups, John R Rickford, Dan Jurafsky, and Sharad Goel. 2020. Racial disparities in automated speech recognition. *Proceedings of the national academy of sciences* 117, 14 (2020), 7684–7689.
- [13] Colin Lea, Zifang Huang, Jaya Narain, Lauren Tooley, Dianna Yee, Dung Tien Tran, Panayiotis Georgiou, Jeffrey P Bigham, and Leah Findlater. 2023. From User Perceptions to Technical Improvement: Enabling People Who Stutter to Better Use Speech Recognition. 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 361, 16 pages. doi:10.1145/3544548.3581224
- [14] Jingjin Li, Shaomei Wu, and Gilly Leshed. 2024. Re-envisioning Remote Meetings: Co-designing Inclusive and Empowering Videoconferencing with People Who Stutter. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) (DIS '24). Association for Computing Machinery, New York, NY, USA, 1926–1941. doi:10.1145/3643834.3661533
- [15] Qisheng Li and Shaomei Wu. 2024. Towards Fair and Inclusive Speech Recognition for Stuttering: Community-led Chinese Stuttered Speech Dataset Creation and Benchmarking. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems* (CHI EA '24). Association for Computing Machinery, New York, NY, USA, Article 331, 9 pages. doi:10.1145/3613905.3650950
- [16] Joshua L Martin and Kelly Elizabeth Wright. 2022. Bias in Automatic Speech Recognition: The Case of African American Language. *Applied Linguistics* 44, 4 (12 2022), 613–630. doi:10.1093/applin/amac066 arXiv:https://academic.oup.com/applij/article-pdf/44/4/613/51053335/amac066.pdf
- [17] Zion Mengesha, Courtney Heldreth, Michal Lahav, Juliana Sublewski, and Ellyse Tuennerman. 2021. “I don’t think these devices are very culturally sensitive.”—Impact of automated speech recognition errors on African Americans. *Frontiers in Artificial Intelligence* 4 (2021), 725911.
- [18] Alisha Pradhan and Amanda Lazar. 2021. Hey Google, Do You Have a Personality? Designing Personality and Personas for Conversational Agents. In *Proceedings of the 3rd Conference on Conversational User Interfaces* (Bilbao (online), Spain) (CUI '21). Association for Computing Machinery, New York, NY, USA, Article 12, 4 pages. doi:10.1145/3469595.3469607
- [19] Alisha Pradhan, Amanda Lazar, and Leah Findlater. 2020. Use of Intelligent Voice Assistants by Older Adults with Low Technology Use. *ACM Trans. Comput.-Hum. Interact.* 27, 4, Article 31 (Sept. 2020), 27 pages. doi:10.1145/3373759
- [20] Yolanda A. Rankin and Kallayah K. Henderson. 2021. Resisting Racism in Tech Design: Centering the Experiences of Black Youth. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 192 (April 2021), 32 pages. doi:10.1145/3449291
- [21] Nan Bernstein Ratner and Brian MacWhinney. 2018. Fluency Bank: A new resource for fluency research and practice. *Journal of fluency disorders* 56 (2018), 69–80.
- [22] Adam Satariano and Cecilia Kang. 2023. How Nations Are Losing a Global Race to Tackle A.I.’s Harms. <https://www.nytimes.com/2023/12/06/technology/ai-regulation-policies.html>.
- [23] Joshua St. Pierre. 2012. The Construction of the Disabled Speaker: Locating Stuttering in Disability Studies. *Canadian Journal of Disability Studies* 1, 3 (Aug. 2012), 1–21. doi:10.15353/cjds.v1i3.54
- [24] Selina Jeanne Sutton, Paul Foulkes, David Kirk, and Shaun Lawson. 2019. Voice as a Design Material: Sociophonetic Inspired Design Strategies in Human-Computer Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3290605.3300833
- [25] Ravichander Vipperla, Steve Renals, and Joe Frankel. 2010. Ageing Voices: The Effect of Changes in Voice Parameters on ASR Performance. *EURASIP Journal on Audio, Speech, and Music Processing* (2010). <https://asmp-urasipjournals.springeropen.com/articles/10.1155/2010/525783>
- [26] Kimi Wenzel, Nitya Devireddy, Cam Davison, and Geoff Kaufman. 2023. Can Voice Assistants Be Microaggressors? Cross-Race Psychological Responses to Failures of Automatic Speech Recognition. 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 109, 14 pages. doi:10.1145/3544548.3581357
- [27] Kimi Wenzel and Geoff Kaufman. 2024. Designing for Harm Reduction: Communication Repair for Multicultural Users’ Voice Interactions. 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 879, 17 pages. doi:10.1145/3613904.3642900
- [28] Meredith Whittaker, Meryl Alper, Cynthia L. Bennett, Sara Hendren, Elizabeth Kazianas, Mara Mills, Meredith Ringel Morris, Joy Lisi Rankin, Emily Rogers, Marcel Salas, and Sarah Myers West. 2019. Disability, Bias & AI Report. *AI Now Institute* (20 11 2019).
- [29] Shaomei Wu. 2023. “The World is Designed for Fluent People”: Benefits and Challenges of Videoconferencing Technologies for People Who Stutter. 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 782, 17 pages. doi:10.1145/3544548.3580788
- [30] Shaomei Wu, Jingjin Li, and Gilly Leshed. 2024. Finding My Voice over Zoom: An Autoethnography of Videoconferencing Experience for a Person Who Stutters. 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 916, 16 pages. doi:10.1145/3613904.3642746
- [31] Esmat Zaidan and Imad Antoine Ibrahim. 2024. AI Governance in a Complex and Rapidly Changing Regulatory Landscape: A Global Perspective. *Humanities and Social Sciences Communications* 11, 1 (2024), 1–18.