Lauren G. Wilcox, PhD
lgw231@acm.org

Bio | LinkedIn | Google Scholar | TASC at Google Research | Hx Lab at Georgia Tech

Summary

A leading voice and pioneer in human–AI interaction research and leadership, advancing sociotechnical approaches to AI safety and responsibility through novel interaction paradigms, evaluation frameworks, and participatory methods. Led advancements spanning foundational model development, healthcare AI, and digital wellbeing at Google, focusing on human feedback mechanisms and community-centered approaches to AI development. Advanced state-of-the-art in human feedback elicitation for large language models, including key contributions to Google's Bard/Gemini community testing program. Published extensively on human–AI collaboration and sociotechnical safety approaches, with multiple best papers in top HCI and AI ethics venues (ACM CHI, CSCW, FAccT).

Earned Degrees

Ph.D. Computer Science, Columbia University	Oct 2013
M.S. Computer Science, Columbia University	May 2006
B.S. Computer Science, Columbia University	May 2004

Experience

Senior Director and Distinguished Applied Scientist, eBay (leading eBay's Office of Responsible AI)	San Jose, CA Sep 2023–
Technology, AI, Society & Culture, Google Research Senior Staff Research Scientist and Group Manager	Mountain View, CA Nov 2022–Sep 2023
People & AI Research (PAIR), Google Research Staff Research Scientist and Research Manager	Mountain View, CA Aug 2021–Nov 2022
Google Medical Brain→ Health, Digital Wellbeing Research Lead	Mountain View, CA Feb 2019–Aug 2021
Georgia Institute of Technology Adjunct Associate Professor, School of Interactive Computing Associate Professor, School of Interactive Computing (tenure) Assistant Professor, School of Interactive Computing	Atlanta, GA Jan 2022– Aug 2020–Jan 2022 Nov 2013–Aug 2020
Columbia University Graduate Research Assistant, Department of Computer Science	New York, NY Sep 2007–Aug 2013

Microsoft Research Research Intern, Computational User Experiences	Redmond, WA Jun 2010–Aug 2010 Jun 2009–Aug 2009
IBM Research Research Co-op, Intelligent Information Interaction	Hawthorne, NY Sep 2007–May 2009
IBM Software Group Executive Briefing Center, Solutions Experience Lab Staff Software Engineer	Austin, TX Jul 2006–Aug 2007

Selected Honors and Awards

Women of Influence, Silicon Valley Business Journal	2024
Best Paper Award, ACM conference on Fairness, Accountability, and Transparency (FAccT) 2024	2024
Best Paper Honorable Mention, ACM conference on Human Factors in Computing Systems (ACM CHI)	2024
ACM Distinguished Member	2023
Best Paper Honorable Mention & Methods Recognition ACM conference on Computer-Supported Cooperative Work And Social Computing (ACM CSCW)	2023
Best Paper Award, ACM conference on Human Factors in Computing Systems (ACM CHI)	2023
Google Research Inclusive Culture Award	2022
Best Case Study, ACM conference on Human Factors in Computing Systems (ACM CHI)	2021
Google Research Excellence in Leadership Award (Responsible AI and Human-Centered Technology)	2021
ACM Senior Member	2020
Best Paper Honorable Mention, ACM conference on Human Factors in Computing Systems (CHI)	2020
Best Paper Award, ACM conference on Designing Interactive Systems (DIS)	2019
Editor’s Choice Paper, Journal of the American Medical Association (JAMIA)	2018

ACM Future of Computing Academy	2017
NSF CAREER Award	2017
High Value Patent Award, International Business Machines (IBM)	2010
Early Tenure Inventor Award, IBM Software Group	2008
Master of Science (MS) Teaching Assistant Fellowship, Columbia University, Department of Computer Science (fully-funded MS degree)	2005–2006

Employment (Selected)

eBay

Senior Director and Distinguished Applied Scientist, Responsible AI
(leading eBay's Office of Responsible AI), Sep 2023– current

Led formation and scaling of eBay's Office of Responsible AI, bringing deep expertise in human–AI interaction and AI safety to shape technical strategy and governance across the company's AI initiatives.

Established comprehensive frameworks for:

- Technical Safety & Evaluation: Developed policies and associated methodologies for dataset construction, transformation, and model evaluation to enable human value alignment and safety testing
- Product Integration: Created scalable approaches to evaluate and enhance AI features across the product portfolio, focusing on human-centered interaction design and safety
- Research Leadership: Advanced state-of-the-art in AI safety evaluation while enabling rapid product development, resulting in industry-leading practices
- Strategic Impact: Built and led two-tier governance structure (VP/SVP) driving company-wide AI strategy, significantly elevating eBay's position in responsible AI development

Outcomes included acceleration of AI product launches while maintaining safety, establishment of new industry standards for responsible AI in e-commerce, and successful integration of human-centered evaluation approaches across the organization.

Google (Summary)

Research Center for Responsible AI and Human-Centered Technology
Google Research, July 2021–Sep 2023

Technology, AI, Society and Culture (TASC) (Team Lead)

Jan 2022–Sept 2023 (note: leadership overlapped with PAIR group co-leadership, below)
Senior Staff Research Scientist and Group Manager

- Led development of novel feedback elicitation methods for LLMs
- Pioneered participatory approaches to AI system development
- Advanced human-centered evaluation frameworks
- Shaped community testing methodologies

Group manager for 17+ person research team pioneering foundational advances in human-centered AI safety and evaluation:

- Human-AI Interaction & Safety: Led development of novel interaction paradigms and evaluation frameworks for generative AI systems, including key contributions to Google's community testing program for Bard/Gemini. Advanced methods for human feedback elicitation and alignment.

- Responsible Data & Evaluation: Developed novel datasets and transparency frameworks for AI system evaluation, including methodological innovations in:
  - Human feedback collection for reinforcement learning
  - Adversarial testing protocols for foundation models
  - Scalable evaluation frameworks for model capabilities and risks
  - Novel approaches to human annotation and data quality

- Participatory AI Research: Created new methodologies for community engagement in AI development and evaluation:

- Culturally-situated evaluation frameworks
  - Community-centered approaches to model testing
  - Novel benchmarking methods incorporating diverse perspectives
  - Stakeholder-informed safety criteria

Research directly shaped Google's AI safety evaluation practices, and advanced industry standards for model and AI application evaluations, novel approaches to human alignment, and frameworks for identifying and mitigating societal risks

Spearheaded new Research Awards for Collective and Society-Centered AI with Technology & Society team, to fund selected research efforts that exemplify a society-centered approach and include community partners throughout research and development.

Recognized by company with **Google Research Inclusive Culture Award (2022)**

People + AI Research (PAIR)

July 2021–Sept 2022. Staff Research Scientist and group co-lead under Director Meredith R. Morris until TASC team graduation in 22Q4.

Recognized by organization with **Excellence in Leadership Award (2021)**

Google Digital Wellbeing (now Google Wellbeing Lab)
Dec 2019–July 2021. In addition to shaping internal strategy, frameworks:

- First-of-its-kind mobile study platform for Google launched in Play store
- Pioneered human–AI interaction and wellbeing research shaping launch of Digital Wellbeing product experience toolkit
- Pioneered a new funding area in our COVID-19 AI Award portfolio. Sponsored multiple faculty awards related to health equity and health AI fairness**,** which advanced strategies to reach vulnerable populations with vaccine information.
- Advanced fundamental research in human-centered AI for health (resulting in multiple influential publications)

Google AI Healthcare (was Medical Brain, then Google Health)

Feb 2019–Dec 2019. Set the human-centered AI research agenda for incubated AI products. Drove research management processes and established first-of-their-kind applications of clinical AI in partnership with external clinical partners for research and software pilot studies, shaping product strategy.

Contributed to org-wide initiatives related to AI in healthcare, including:

- Co-lead on efforts to **promote safety, fairness, and equity in Product and Research groups** (co-founded Health Equity at Google)

- First published studies examining the use of a deep learning-based AI system in patient care, used as an **exemplar for clinical trial designs** and reporting that include AI in healthcare and contributed to critical discussions that led to new clinical trial design and reporting standards.

  (Full employment record including Georgia Institute of Technology available on LinkedIn)

Refereed Publications

Conference Presentations with Proceedings (Refereed)

b2c1. A Rothschild, D Wang, NJ Vilvanathan, L Wilcox, C DiSalvo, B DiSalvo. The Problems with Proxies: Making Data Work Visible through Requester Practices. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society 7, 1255-1268.

b2c2. M Tahaei, D Wilkinson, A Frik, M Muller, R Abu-Salma, L Wilcox. Surveys Considered Harmful? Reflecting on the Use of Surveys in AI Research, Development, and Governance. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society 7, 1416-1433.

b2c3. Michael Madaio, Shivani Kapania, Rida Qadri, Ding Wang, Andrew Zaldivar, Remi Denton, and Lauren Wilcox. Learning about Responsible AI On-The-Job: Learning Pathways, Orientations, and Aspirations, ACM FAccT 2024, **Best Paper Award**

b2c4. Zijie J. Wang, Chinmay Kulkarni, Lauren Wilcox, Michael Terry, Michael Madaio. Farsight: Fostering Responsible AI Awareness During AI Application Prototyping, CHI Conference on Human Factors in Computing Systems (CHI), 2024. (Github project) **Best Paper Honorable Mention**

b2c5. Allison Woodruff, Renee Shelby, Patrick Gage Kelley, Steven Rousso-Schindler, Jamila Smith-Loud, Lauren Wilcox. How Knowledge Workers Think Generative AI Will (Not) Transform Their Industries. CHI Conference on Human Factors in Computing Systems (CHI), 2024.

b2c1. Lauren Wilcox, Renee Shelby, Rajesh Veeraraghavan, Oliver Haimson, Gabi Erickson, Michael Turken, Beka Gulotta. Infrastructuring Care: How Trans and Non-Binary People Meet Health and Well-Being Needs through Technology. Proc. *ACM Conference on Human Factors in Computing Systems (ACM CHI) 2023* **Best Paper Award**

b2c2. Qiaosi (Chelsea) Wang, Michael Madaio, Shivani Kapania, Shaun Kane, Michael Terry, Lauren Wilcox. Designing Responsible AI: Adaptations of UX Practice to Meet Responsible AI Challenges. ACM Conference on Human Factors in Computing Systems (ACM CHI) 2023. **Editor's Choice Selection**

b2c3. Ned Cooper, Tiffanie Horne, Gillian R Hayes, Courtney Heldreth, Michal Lahav, Jess Holbrook, Lauren Wilcox. A Systematic Review and Thematic Analysis of Community-Collaborative Approaches to Computing Research. Proc. *ACM CHI Conference on Human Factors in Computing Systems, 2022*

b2c4. Negar Rostamzadeh, Diana Mincu, Subhrajit Roy, Andrew Smart, Lauren Wilcox, Mahima Pushkarna, Jessica Schrouff, Razvan Amironesei, Nyalleng Moorosi, Katherine Heller. Healthsheet: Development of a Transparency Artifact for Health Datasets. Proc. *ACM Conference on Fairness, Accountability, and Transparency (FAccT), 2022*.

b2c5. Ferran Altarriba Bertran, Alexandra Pometko, Muskan Gupta, Lauren Wilcox, Reeta Banerjee, Katherine Isbister. The playful potential of shared mealtime: a speculative catalog of playful technologies for day-to-day social eating experiences. Proc ACM on Human–Computer Interaction, *ACM CHI Play 2021*

b2c6. Mark Schurgin, Mark Schlager, Laura Vardoulakis, Laura Pina, Lauren Wilcox. Isolation in Coordination: Challenges of Informal Caregivers in the USA. Proc. *ACM CHI Conference on Human Factors in Computing Systems, 2021*

b2c7. Carrie J. Cai, Samantha Winter, David Steiner, Lauren Wilcox, Michael Terry. Onboarding Materials as Cross-functional Boundary Objects for Developing AI

Assistants. Proc. Proc. *ACM CHI Conference on Human Factors in Computing Systems, 2021*. **Best Case Study**

b2c8. Qiaosi Wang, Shan Jing, David Joyner, Lauren Wilcox, Hong Li, Thomas Ploetz and Betsy DiSalvo. Sensing Affect to Empower Students: Learner Perspectives on Affect-Sensitive Technology in Large Educational Contexts. *Proc. ACM Learning at Scale 2020*. Atlanta, GA, USA. **Best Student Paper Award**

b2c9. Emma Beede, Elizabeth Elliott Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, Laura Vardoulakis. A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy. *Proc. ACM CHI '20,* Proc. *ACM CHI Conference on Human Factors in Computing Systems, 2020*. Honolulu, Hawai'i, USA. 1–12. DOI: https://doi.org/10.1145/3313831.3376718
Best Paper Honorable Mention (Top 5%, Acceptance rate 24%)
(among top-most downloaded research articles in history of ACM CHI)

b2c10. Matthew Hong, Udaya Lakshmi, Kimberly Do, Sampath Prahalad, Thomas Olson, Rosa Arriaga, Lauren Wilcox. Using Diaries to Probe the Illness Experiences of Adolescent Patients and Parental Caregivers. *Proc. ACM CHI '20, ACM CHI Conference on Human Factors in Computing Systems, 2020*. Honolulu, Hawai'i, USA. 1–16. DOI: https://doi.org/10.1145/3313831.3376426 (Acceptance rate 24%)

b2c11. Wilcox, L., DiSalvo, B., Henneman, R., Wang, Q. Design in the HCI Classroom: Setting a Research Agenda. Proc. *ACM 2019 Conf. on Designing Interactive Systems (ACM DIS '19),* ACM, New York, NY, USA, 871-883.
Best Paper Award (Top 1% of papers)

(Full list includes work in data privacy, participatory design, and digital health, see Google Scholar)

Selected Published and Accepted Journal Articles

b1j1. Sheena Erete, Eric Corbett, Natasha Smith-Walker, Jay L. Cunningham, Erin Gatz, Tina M. Park, Tam Perry, Lauren Wilcox, Remi Denton. Towards Equitable Community-Industry Collaborations: Understanding the Experiences of Nonprofits' Collaborations with Tech Companies. Proc. ACM Hum.-Comput. Interact. CSCW 2025. *Accepted for presentation and publication.*

b1j2. Lauren WIlcox, Robin Brewer, Fernando Diaz. Speculating Consent Futures: Case Study on Voice Data Collection with Clinicians. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW 2, Article 316 (October 2023). **Best Paper Honorable Mention, Methods Recognition**

b1j3. Ferran Altarriba Bertran, Alexandra Pometko, Muskan Gupta, Lauren Wilcox, Reeta Banerjee, and Katherine Isbister. 2022. Designerly Tele-Experiences: A New Approach to Remote Yet Still Situated Co-Design. ACM Trans. Comput.-Hum. Interact. 29, 5, Article 44 (October 2022)

b1j4. Cai, C., Winter, S., Steiner, D., Wilcox, L., Terry. M. "Hello AI": Uncovering Needs of Medical Practitioners for Human–AI Collaborative Decision-Making. Proc *ACM on Human-Computer Interaction. 3, CSCW, Article 104 (November 2019) (ACM CSCW 2019) Austin, TX, USA.*

(Full list includes work in data privacy, digital wellbeing, and digital health, see Google Scholar)

Selected Patents (Issued)

c1p1. Morris, D., Tan, D., Wilcox, L., Smith, G., Karlson, A., Roseway, A. Automatic Generation of an Executive Summary for a Medical Event in an Electronic Medical Record. Issued Oct. 2017, USPTO: 9,805,160.

c1p2. Wilcox, L., Lu, J., Lai, J. Method and Apparatus for Run-Time User Control of System-Generated Content. Issued Aug. 2017, USPTO: 20110029472.

c1p3. Wilcox, L., Lamb, M., Lauridsen, C., Ong, M., Creating and Managing Reference Elements of Deployable Web Archive Files. Issued Nov. 2013, USPTO: 8,583,658.

c1p4. Lauridsen, C., Wilcox, L. Applying User-Generated Deployment Events to a Grouping of Deployable Portlets. Issued July 2013, USPTO: 8,495,048.

c1p5. Lauridsen, C., Ong, M., Wilcox, L. Summarizing Portlet Usage in a Portal Page. Issued May 2012, USPTO: 8,191,002.

c1p6. Lauridsen, C., Ong, M., Wilcox, L. Summarizing Portlet Usage Captured Responsive to Trigger Events in a Portal Page. Issued Mar. 2011, USPTO: 7,904,818 (IBM High Value Patent Award).

Selected Presentations

Invited and Juried Presentations (Selected Seminar, Panel Invitations, and Keynotes)

d1t1. Panelist: Dreamforce 2024 Managing Data Privacy, Security, and AI Ethics in High Tech

d1t2. Panelist: Partnership on AI, AI Policy Forum

d1t3. Panelist: Reuters AI Momentum, Breakfast Panel, June 2024

d1t4. Keynote for Nordic AI Meet, 2023

d1t5. Panelist: Community-collaborative Visions for Computing Research, ACM CSCW 2023

d1t6. Panelist: Race and Racism in Digital Health and Information Technology, U. Maryland

d1t7. Panelist: Community-collaborative Visions for Computing Research, ACM FAccT Craft 2023

d1t8. Panelist: User Engagement in Algorithm Testing and Auditing, ACM FAccT 2023

d1t9. Panelist: Tales from the Trenches, Symposium on Artificial Intelligence for Learning Health Systems (SAIL) 2023

d1t10. Panelist: Societal Bias in Data, Machine Learning, Health, and Science, AAAS 2023

d1t11. Wilcox, L."Participatory Approaches to Health AI", Columbia University, Department of Biomedical Informatics, April 2023.

d1t12. Panelist: Building Models and Solutions Equitably and Inclusively in Health and Healthcare, 2022 Health Datapalooza and National Health Policy Conference

d1t13. Wilcox, L. “Participatory Approaches to AI for Digital Health and Wellbeing” AI and Social Impact Series, Harvard University, Mar 2021.

d1t14. Wilcox, L. “Participatory Approaches to AI for Digital Health and Wellbeing” The Donald Bren School of Information and Computer Sciences, University of California, Irvine, Mar 2021.

d1t15. Invited Digital Wellbeing panel at the Global Experience Summit, a large, international, three-part online festival covering every element of the user/customer experience process with hundreds of attendees.

d1t16. Wilcox, L. “Participatory Approaches to AI for Digital Health and Wellbeing” Digital Culture Speaker Series, School of Arts, Media, and Engineering, Arizona State University, Oct 2020.

Other Professional Activities

Co-author of blog post receiving views from tens of thousands of unique visitors, proposing changes to the peer review process to mitigate negative impacts of computing research. The proposal influenced the submission and peer-review process at NeurIPS 2020 and onward.

Hecht, B., Wilcox, L., Bigham, J.P., Schöning, J., Hoque, E., Ernst, J., Bisk, Y., De Russis, L., Yarosh, L., Anjum, B., Contractor, D. and Wu, C. 2018. It’s Time to Do Something: Mitigating the Negative Impacts of Computing Through a Change to the Peer Review Process. *ACM Future of Computing Blog*. https://acm-fca.org/2018/03/29/negativeimpacts/.

Selected Professional Service

Organization and Chairpersonship of Technical Sessions, Conferences, Committees

2023 Associate Chair, Technical Program Committee for ACM FAccT 2024
2023 Partnership on AI, Global Task Force for Inclusive AI
2021 Subcommittee Chair, ACM CHI 2022
2020 Computing Innovation Fellows 2020 Reviewer
2020 Associate Chair, Technical Program Committee for ACM CHI 2021
2019 Associate Chair, Technical Program Committee for ACM CHI 2020