Anna Kawakami

annakawakami.com ♦ akawakam@andrew.cmu.edu

EDUCATION

Carnegie Mellon University 09/2021 – 05/2026 (Expected)
School of Computer Science
Pittsburgh, PA
PhD student in Human-Computer Interaction
Advisors: Ken Holstein and Haiyi Zhu

Wellesley College 09/2017 – 06/2021
Wellesley, MA
Bachelor of Arts in Computer Science (Honors), Minor in Economics
Advisor: Eni Mustafaraj

University of Oxford 09/2019 – 06/2020
Oxford, England
Visiting student in Computer Science and Mathematics

RESEARCH POSITIONS

Carnegie Mellon University, Human-Computer Interaction Institute Fall 2021 –
Graduate Researcher advised by Drs. Ken Holstein and Haiyi Zhu

Microsoft Research, FATE Team in New York City Summer 2023
Fairness, Accountability, Transparency, and Ethics in AI (FATE) group
PhD Research Intern mentored by Dr. Alex Chouldechova

Microsoft Research, FATE Team in Montreal Summer 2022
PhD Research Intern mentored by Drs. Koustuv Saha, Vera Liao, Jina Suh, Alexandra Olteanu, and Shamsi Iqbal

Wellesley College, Computer Science Department 2018 – 2021
Undergraduate Research Assistant advised by Dr. Eni Mustafaraj

Microsoft Research, Aether and Adaptive Systems & Interactions Group in Redmond Summer 2020
Research Intern advised by Drs. Mihaela Vorvoreanu, Ece Kamar, and Ben Zorn
Undergraduate Research Internship in Computing Program

University of Oxford, Centre for Technology and Global Affairs 2019 – 2020
Undergraduate Research Assistant advised by Dr. Takis Metaxas

Carnegie Mellon University, Software, Systems, and Society Department Summer 2019
Undergraduate Researcher advised by Drs. Lujo Bauer and Limin Jia
NSF REU in Software Engineering Program

CONFERENCE & JOURNAL PUBLICATIONS

An asterisk (*) indicates co-first authorship. A tilde (~) indicates co-senior authorship.

P17 **Kawakami, A.,** Zhu, H., Holstein, K. "AI Measurement as a Participatory Design Practice: A Multidisciplinary Framework and Research Agenda." *Under review.* 2026

P16 Lin, H. **Kawakami, A.,** D'Ignazio, C., Holstein, K., Gajos, K. "Funding AI For Good: A Call for Meaningful Engagement" Forthcoming at *ACM CHI Conference on Human Factors in Computing Systems,* CHI '26. 2026

P15 **Kawakami, A.,** Taylor, J., Fox, S., Zhu, H., Holstein, K. "*AI Failure Loops* in Devalued Work: The Confluence of Overconfidence in AI and Underconfidence in Worker Expertise." Forthcoming at *Big Data & Society.* 2026

P14 **Kawakami, A.,** Wilkinson, D., Chouldechova, A. "Do Responsible AI Artifacts Help Advance Stakeholder Goals? Four Key Barriers Perceived by Legal and Civil Stakeholders." In *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society,* AIES '24. 2024

P13 **Kawakami, A.,** Coston, A., Heidari, H. ~, Holstein, K.~, Zhu, H.~. "Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions on AI Design and Use" In *ACM SIGCHI Conference on Computer-Supported Cooperative Work & Social Computing*, CSCW '24. 2024

P12 **Kawakami, A.,** Coston, A., Zhu, H.~, Heidari, H. ~, Holstein, K~. "The *Situate AI* Guidebook: Co-designing a Toolkit to Support Multi-stakeholder, Early-stage Deliberations Around Public Sector AI Proposals" In *ACM CHI Conference on Human Factors in Computing Systems,* CHI '24. 2024

P11 **Kawakami, A.,** Guerdan, L., Cheng, Y., Lee, M., Carter, S., Arechiga, N., Glazko, K., Zhu, H.~, Holstein, K~. "Training Towards *Critical Use:* Learning to Situate AI Predictions Relative to Human Knowledge." In *ACM Collective Intelligence, CI '23.* 2023

P10 **Kawakami, A.,** Chowdhary S., Iqbal S., Liao V., Olteanu A., Suh J., and Saha, J. "Sensing Wellbeing in the Workplace, Why and For Whom? Envisioning Impacts with Organizational Stakeholders." In *ACM SIGCHI Conference on Computer-Supported Cooperative Work & Social Computing*, CSCW '23. 2023

P09 Chowdhary, S., **Kawakami, A.,** Gray, M. L., Suh, J., Olteanu A., and Saha, J. "Can Workers Meaningfully Consent to Workplace Wellbeing Technologies?" In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23. 2023

P08 Coston, A., **Kawakami, A.,** Zhu, H., Holstein, K. and Heidari, H. "A Validity Perspective on Evaluating the Justified Use of Data-driven Decision-making Algorithms." In *IEEE Conference on Secure and Trustworthy Machine Learning*, IEEE SaTML '23. **Best Paper Award (top 1% of submissions)** 2023

P07 **Kawakami, A.,** *, Sivaraman, V.*, Cheng, H.F., Stapleton, L., Perer, A., Wu, Z.S., Zhu, H. and Holstein, K. "'Why Do I Care What's Similar?' Probing Challenges in AI-Assisted Child Welfare Decision-Making through Worker-AI Interface Design Concepts." In *ACM SIGCHI Conference on Designing Interactive Systems,* DIS '22. 2022

P06 **Kawakami, A.,** Sivaraman, V., Cheng, H.F., Stapleton, L., Cheng, Y., Qing, D., Perer, A., Wu, Z.S., Zhu, H. and Holstein, K. "Improving Human-AI Partnerships in Child Welfare: Understanding Worker Practices, Challenges, and Desires for Algorithmic Decision Support." In *ACM CHI Conference on Human Factors in Computing Systems,* CHI '22. **Best Paper Honorable Mention Award (top 5% of submissions)** 2022

P05 Cheng, H.F.*, Stapleton, L.*, **Kawakami, A.,** Sivaraman, V., Cheng, Y., Qing, D., Perer, A., Holstein, K., Wu, Z.S. and Zhu, H. "How Child Welfare Workers Reduce Racial Disparities in Algorithmic Decisions." In *ACM CHI Conference on Human Factors in Computing Systems,* CHI '22. 2022

P04 **Kawakami, A.,** Umarova, K., Huang, D. and Mustafaraj, E**.** "The 'Fairness Doctrine' lives on? Theorizing about the Algorithmic News Curation of Google's Top Stories" In *ACM Conference on Hypertext and Social Media*, HT '20. 2020

P03 **Kawakami, A.,** Umarova, K., and Mustafaraj, E. "The Media Coverage of the 2020 US Presidential Election Candidates through the Lens of Google's Top Stories." In *AAAI Conference on Web and Social Media*, ICWSM '20. 2020

P02 Lerner A., He H., **Kawakami, A.**, Zeamer, S., Hoyle, R. "Privacy and Activism in the Transgender Community." In *ACM CHI Conference on Human Factors in Computing Systems*, CHI '20. 2020

P01 Cobb, C., Surbatovich, M., **Kawakami, A.,** Sharif, M., Bauer, L., Das, A. and Jia, L. "How Risky are Real Users' IFTTT Applets?" In *Sixteenth Symposium on Usable Security and* Privacy, SOUPS '20. 2020

TUTORIALS ORGANIZED

T01 **Kawakami, A.,** Blodgett, S. L., Barocas, S., Chouldechova, A., Jacobs, A., Sheng, E., Vaughan, J. W., Wallach, H., Winecoff, A., Wang, A., Zhu, H., Holstein, K. "Translation Tutorial: AI Measurement as a Stakeholder-Engaged Design Practice." *ACM Conference on Fairness, Accountability, and Transparency,* FAccT '25. 2025

WORKSHOPS ORGANIZED

W04 Saxena, D., Moon, E., **Kawakami, A.,** Jackson, C., Lee, M., Fox, S., Guha, S., Erete, S., Hardy, J. "Participation, Procurement, & Proof of Impact in Public Sector AI Innovation." Forthcoming workshop at the *ACM CHI Conference on Human Factors in Computing Systems,* CHI '26. 2026

W03 Ming, J., Pei, L., Varansai, R., **Kawakami, A.,** Verdezoto, N., Cheon, E. "Labor, Visibility, and Technology: Weaving Together Academic Insights and On-Ground Realities." Workshop at the *ACM Conference on Computer-Supported Cooperative Work and Social Computing,* CSCW '24. 2024

W02 Wan, R., Alvarado Garcia, A., Saxena, D., Vajiac, C., **Kawakami, A.,** Stapleton, L., Zhu, H., Holstein, K., Candello, H., Badillo-Urquiola, K. "Community-driven AI: Empowering people through responsible data-driven decision-making." Workshop at the *ACM Conference on Computer-Supported Cooperative Work and Social Computing,* CSCW '23. 2023

W01 Stapleton, L., Saxena, D., **Kawakami, A.,** Nguyen, T., Flügge, A., Eslami, M., Møller, N., Lee, M. K., Guha, S., Holstein, K. "Who has an interest in 'Public Interest Technology'? Critical questions for working with local government & impacted communities." Workshop at the *ACM Conference on Computer-Supported Cooperative Work and Social Computing,* CSCW '22. 2022

WORKSHOP PAPERS & POSTERS

S05 **Kawakami, A.,** Taylor, J., Fox, S., Zhu, H., Holstein, K. "*AI Failure Loops* in Feminized Labor: Understanding the Interplay of Workplace AI and Occupational Devaluation." In *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society,* AIES '24 (non-archival). 2024

S04 **Kawakami, A.,** Coston, A., Heidari, H., Holstein, K, Zhu, H. "Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions Around AI Design and Use." In *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '23 (non-archival). 2023

S03 **Kawakami, A.,** Coston, A., Heidari, H., Zhu, H., Holstein, K. "Recentering Validity Considerations through Early-stage Deliberations around AI and Policy Design." In Workshop on Designing Technology and Policy Simultaneously at *ACM Conference on Human Factors in Computing Systems,* CHI '23. 2023

S02 Coston, A., **Kawakami, A.,** Zhu, H., Holstein, K., Heidari, H. "A validity perspective on the justified use of data-driven decision-making algorithms." In *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '22 (non-archival). 2022

S01 **Kawakami, A.,** Guerdan, L., Cheng, Y., Sun, A., Hu, A., Glazko, K., Arechiga, N., Lee, M., Carter, S., Zhu, H., Holstein, K. "Towards a Learner-centered Explainable AI: Lessons from the Learning Sciences." In Workshop on Human-Centered Explainable AI (HCXAI) at *ACM Conference on Human Factors in Computing Systems,* CHI '22. 2022

POLICY WRITING

Po01 **Kawakami, A.,** Zhu, H., Holstein, K. "Expanding State and Local Government Capacity for AI Procurement and Use via a Federal Center of Excellence." *Policy memo with Federation of American Scientists.* 2025

OTHER WRITING & REPORTS

R02 **Kawakami, A.,** Vorvoreanu, M., Zorn, B., Liles, N., Kamar, E. "AI Reliability & Safety: Practices and Challenges at Microsoft." *Microsoft-internal technical report.* 2021

R01 **Kawakami, A.** "The News We See When Searching: Exploring Users' Mental Models of Google's Top Stories." *Honors senior thesis.* 2021

GRANTS

G01 **Participatory AI Measurement Design: Supporting Early Stakeholder Engagement in the Design of AI Measurements** **Awarded $75,000**
Google Academic Research Award 2024
Co-written with Ken Holstein and Haiyi Zhu

OTHER GRANT WRITING EXPERIENCE

Defining Pro-Worker AI for Workers, by Workers: Participatory Design of Worker-Centered Measurements for AI and Policy
Block Center for Technology and Society, Co-written with Ken Holstein and Haiyi Zhu

Defining Pro-Worker AI for Workers, by Workers: Participatory Design of Worker-Centered Measurements for AI and Policy
The Washington Center for Equitable Growth, Co-written with Ken Holstein and Haiyi Zhu

SELECTED AWARDS & HONORS

NSF Graduate Research Fellow ($147,000), National Science Foundation 2023

Best Paper Award Recipient ($1,000), IEEE Conference on Secure and Trustworthy Machine – *awarded to the top 1% of submissions* 2023

K&L Gates Ethics and AI Presidential Fellow ($60,000)**,** K&L Gates Endowment for Ethics and Computational Technologies at Carnegie Mellon University – *awarded annually to 3 – 4 doctoral students at Carnegie Mellon University whose studies relate to the field of ethics and computational science and technology* 2023

CASMI PhD Fellow, Center for Advancing Safety of Machine Intelligence at Northwestern University – *awarded to PhD students researching topics around the design, development, deployment, and evaluation of safe AI technologies* 2022

Best Paper Honorable Mention Award Recipient, ACM CHI Conference on Human Factors in Computing Systems (CHI '22) – *awarded to the top 5% of submissions* 2022

Trustee Scholar, Wellesley College Board of Trustees – *awarded to four graduating seniors at Wellesley College by the Wellesley Board of Trustees* 2021

Academic Excellence in Computer Science Award, Wellesley College Computer Science Department – *awarded to two graduating Computer Science majors by the Wellesley Computer Science Department* 2021

OTHER RECOGNITION

Travel Award, GeorgiaTech Summit on Responsible Computing, AI, and Society 2025

Travel Award, Human-Computer Interaction Consortium (HCIC) Annual Meeting 2025

Two Special Recognitions for Outstanding Reviews, ACM CHI Conference on Human Factors in Computing Systems (CHI '25) 2024

Travel Award ($900), AAAI HCOMP Conference on Human Computation and Crowdsourcing 2024

Travel Award ($1,000)**,** AAAI/ACM Conference on AI, Ethics, and Society 2024

Day One Project ($500)**,** Federation of American Scientists – *policy idea for supporting AI governance was selected for further development into a policy memo* 2024

Summer School on Designing Technology for Social Impact Participant ($1,000)**,** Cornell University (*deferred to 2025*) 2024

Travel Award ($1,700)**,** AAAI HCOMP Conference on Human Computation and Crowdsourcing and ACM CI Conference on Collective Intelligence 2023

Special Recognition for Outstanding Review, ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW '23) 2022

Invited Participant (€2,500), International Symposium on Augmented Intelligence at Work, hosted by University of Cambridge, the RWTH Aachen University, and the Volkswagen Stiftung – *awarded to 50 researchers across the globe, working in the field of augmented intelligence at work* 2022

Honors in Computer Science, Wellesley College Computer Science Department – *awarded to students who successfully complete an honors undergraduate thesis* 2021

Durant Scholar Magna Cum Laude, Wellesley College – awarded to students with an average GPA of 3.75 of higher 2021

Sigma Xi Scientific Research Honor Society Nomination, Wellesley College Sigma Xi Chapter – *nominated by a Wellesley College professor* 2021

INVITED TALKS & EVENTS

Invited Speaker, AI Transformation in Government Panel, hosted by European Political Science Society. 2026

Invited Speaker, Annual VU Women in Tech Conference, hosted by Villanova University. *Impacts of AI on Feminized Labor.* 2026

Invited Expert, Workshop on AI and the Future of Work for Young People, hosted by Microsoft AI Futures. 2025

Invited Speaker, hosted by Harvard University, Center for Human-centered AI Research and Methods. *Designing Positive Futures for AI in Socially Complex Work With and For Workers.* 2025

Invited Speaker, hosted by Wellesley College, Computer Science Department. *Designing Positive Futures for AI in Socially Complex Work With and For Workers.* 2025

Invited Speaker, Trust in the Times of Generative AI, Academic Fringe Festival, hosted by TU Delft. *Designing Positive Futures for AI in Socially Complex Work With and For Workers.* 2025

Invited Speaker, Engineering Design, hosted by The Ellis Upper School. 2025

Guest Lecturer, AI-Assisted Decision-Making, hosted by Brown University, Introduction to Socio-technical Systems and HCI. *AI-Assisted Social Decision-Making in the Workplace.* 2025

Invited Speaker, AI and the Future of Work Conference, hosted by The Wharton School of the University of Pennsylvania. 2025

Invited Expert, Expert Convening on Predictive Analytics in Child Welfare, hosted by University of Bristol. 2025

Invited Expert, Policy Summit, hosted by Federation of American Scientists. 2024

Invited Expert, Expert Convening on Public Sector AI, hosted by Roosevelt Institute. 2024

Guest Lecturer, Automation and AI in the Workplace, hosted by Carnegie Mellon University, Human Factors Course. 2024

Invited Speaker, Open Science Community Event, hosted by Cohere for AI. 2024

Guest Lecturer, AI in Social Decision-Making, hosted by Carnegie Mellon University, Human-AI Interaction Course. *Understanding and Designing for AI-Assisted Social Decision-Making in the Workplace.* 2024

Guest Lecturer, Automation and AI in the Workplace Session, hosted by Carnegie Mellon University, Human Factors Course. *Supporting Worker-AI Complementarity for Complex Decision-Making in the Real-World.* 2023

Invited Speaker, Computer Science Colloquium, hosted by Wellesley College. *Supporting Responsible AI Practices in Public Sector Contexts.* 2023

Invited Speaker, Human-AI Interaction Session, hosted by Carnegie Mellon University, User Experience Association. 2022

Invited Speaker, Human-AI Teams Session, hosted by Institute for Operations Research and Management Sciences (INFORMS) Annual Meeting. *Understanding human-AI decision-making in the real-world: From observational studies to theoretical models.* 2022

Invited Participant, International Symposium on Augmented Intelligence at Work, hosted by University of Cambridge, the RWTH Aachen University, the Volkswagen Stiftung. 2022

Invited Panelist, "Future of HCI" Panel, hosted by Carnegie Mellon University, Introduction to Human-Computer Interaction Course. 2022

Invited Panelist, Graduate Student Panel, hosted by Carnegie Mellon University, REU in Software Engineering Program. 2022

Invited Panelist, Academia Tech Talk Series, hosted by Wellesley College, Developer Student Club. 2021

POLICY CITATIONS

Center for Democracy and Technology, A framework for assessing AI transparency in the public sector. 12/2025

Center for Democracy and Technology, Fostering responsible tech use: Balancing the benefits and risks among public child welfare agencies. 05/2024

Office of Human Services Policy, Avoiding racial bias in child welfare agencies' use of predictive risk modeling. 10/2022

SELECTED PRESS

Techtonic Justice, Inescapable AI: The ways AI decides how low-income people work, live, learn, and survive. 11/2024

Logic(s) Magazine, The Allegheny Family Screening Tool's overestimation of utility and risk. 12/2023

Center for Advancing Safety of Machine Intelligence, Empowering Human Knowledge for More Effective AI-Assisted Decision-Making 11/2023

Center for Advancing Safety of Machine Intelligence, Award-winning research highlights challenges and opportunities for more reliable human-AI decision-making. 09/2023

Data & Society, How recidivism risk assessments impede judgments. 04/2023

ACLU, The devil is in the details: Interrogating values embedded in the Allegheny Family Screening Tool. 03/2023

Children's Bureau Express, Use of artificial intelligence-based decision tools in child welfare. 11/2022

Associated Press News, White House unveils artificial intelligence 'Bill of Rights.' 10/2022

The Gavel, A Look at Child Welfare in the Digital Age. 07/2022

The Takeaway, The biases behind predictive algorithms for child welfare tracking. 06/2022

Marketplace Tech, Should an algorithm play a role in child welfare decisions? 06/2022

The Imprint, Oregon officials phase out use of artificial intelligence tool in child welfare cases. 06/2022

Techdirt, Oregon state officials dump AI tool used to initiate child welfare investigations. 06/2022

PBS News Hour, Oregon dropping AI tool used to help decide child abuse cases. 06/2022

Engadget, Oregon is shutting down its controversial child welfare AI in June. 06/2022

Associated Press (AP News), Oregon dropping AI tool used in child abuse cases. 06/2022

Pittsburgh Tribune-Review, Editorial: Can algorithms help or hurt child protective services? 05/2022

Open Minds, The power & problems of algorithm-driven decision-making in child welfare. 05/2022

Scary Mommy, Data isn't objective: An algorithm is flagging black homes for child welfare investigations. 05/2022

InnovateLI, The slippery slope of child abuse algorithms. 05/2022

PBS News Hour, How an algorithm that screens for child neglect could harm racial disparities. 04/2022

Associated Press News, An algorithm that screens for child neglect raises concerns. 04/2022

RESEARCH MENTORSHIP

Ashveen Banga, Carnegie Mellon University, HCI Independent Study 06/2025–

Chloe Qianhui Zhao, Carnegie Mellon University, Research Assistant 06/2025–

Adrian Shin, Carnegie Mellon University, HCI Independent Study 05/2024–05/2025

Harrison Leon, Carnegie Mellon University, MSPPM Capstone Project 01–05/2024

Jonathan Ho, Carnegie Mellon University, HCI Independent Study 02–05/2023

Shreya Chowdhary, Olin College, Microsoft Undergrad Research Intern 06–09/2023

Yang Cheng, Carnegie Mellon University, MHCI Independent Study 02–10/2022

Diana Qing, UC Berkeley, Research Assistant 09/2021–04/2022

Anita Sun, Carnegie Mellon University, MHCI Independent Study 09–12/2021

Alison Hu, Carnegie Mellon University, HCI Independent Study 09–12/2021

Matthew Ok, Carnegie Mellon University, HCI Independent Study 09–12/2021

Mia Tang, Carnegie Mellon University, HCI Independent Study 09–12/2021

TEACHING POSITIONS

Human Factors, Carnegie Mellon University, Teaching Assistant Spring 2024

Programming User Interfaces, Carnegie Mellon University, Teaching Assistant Fall 2024

SELECTED ACADEMIC & COMMUNITY SERVICE

CHI, Associate Chair	2025 –
CHI, CSCW, FAccT, DIS, Reviewer	2022 –
Graduate Application Support Program, Carnegie Mellon University, Mentor	2022–2024
AI Fairness, Equity, Accountability, Transparency Reading Group, Carnegie Mellon University, Co-Organizer	2024
REU in HCI, Carnegie Mellon University, Admissions Committee	2024
Northeast HCI Meeting, Local Logistics Volunteer	2024
HCII Women and Non-binary Weekly Lunch, Carnegie Mellon University, Co-Organizer	2022–2024
REU in Software Engineering, Carnegie Mellon University, Admissions Committee	2022–2024
HCII Four-Minute Madness, Carnegie Mellon University, Student Co-Organizer	2022–2023
OurCS Conference, Carnegie Mellon University, Co-Mentor	2022
HCII Open House, Carnegie Mellon University, Student Co-Leader	2022
Computer Science Club, Wellesley College, Executive Board Member	2018–2019