

RESEARCH INTERESTS

I work to create more equitable, just, and empowering natural language processing (NLP) technologies by centering people—who are inseparable from language—in technologies’ development and evaluation. I take an interdisciplinary and sociotechnical perspective; I draw from and collaborate across disciplines, including NLP, machine learning and AI, responsible AI, human-computer interaction, (socio)linguistics, and science & technology studies, and I use a variety of methods. My recent work has focused on making visible and interrogating values, assumptions, and practices shaping NLP technologies; principled evaluation and measurement, particularly of generative AI capabilities, behaviours, and impacts; understanding how people perceive and use generative AI; and supporting AI and NLP practitioners in their ethical work.

EDUCATION

University of Massachusetts Amherst

Ph.D., Computer Science	2021
M.S., Computer Science	2018

Wellesley College

B.A., Mathematics	2015
-------------------	------

RESEARCH EXPERIENCE

Principal Researcher. Microsoft Research Montréal.	2025 – 2026
Senior Researcher.	2022 – 2025
Postdoctoral Researcher.	2020 – 2022
Fairness, Accountability, Transparency, & Ethics (FATE) Group.	

Research Intern. Microsoft Research New York City.	Summer 2019
Mentors: Solon Barocas, Hal Daumé III, Hanna Wallach.	

Visiting Appointments

Independent Visiting Researcher. Mila - Québec Artificial Intelligence Institute.	2026 –
Host: Jackie Chi Kit Cheung	

HONORS AND AWARDS

ACL Senior Area Chair Highlight Award ($\leq 3\%$ of accepted main conference papers)	2025, 2026
100 Brilliant Women in AI Ethics	2022
ACL Outstanding Reviewer Award	2021
EMNLP Outstanding Reviewer Award	2020
ACL Outstanding Reviewer Award	2020
NSF Computing Innovation Fellowship (declined)	2020
Rising Stars in EE/CS Workshop Participant (fully funded travel & arrangements)	2019
ICWSM Best Reviewer Award	2019
Microsoft Accomplishments in Search & Artificial Intelligence Award for UMass CICS	2019
UMass CICS Outstanding Synthesis Award	2019
New England Future Faculty Workshop Participant	2018
Passed Ph.D. Qualifying Exam (Portfolio) with Distinction	2018
NSF Graduate Research Fellowship	2017
W-NUT Best Paper Award	2017

CRA-Women Grad Cohort Workshop Participant (fully funded travel & arrangements)	2016
UMass CICS Paul Utgoff Memorial Scholarship (for one student in machine learning)	2016

PUBLICATIONS

Full Conference Publications (Peer-Reviewed)

- [31] **Learning by Chatting? Investigating the Impact of Generative AI on Information Seeking and Learning.** Shravika Mittal, Su Lin Blodgett, and Q. Vera Liao. *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2026.
- [30] **Illusions of the Gold Standard: A Large-scale Analysis of Human Evaluation Protocols for Long-form Text Generation.** Katelyn X. Mei, Yi-Li Hsu, Minjoon Choi, Zongwan Cao, Chenjun Xu, Bingbing Wen, Su Lin Blodgett, and Lucy Lu Wang. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2026. 👩 Senior Area Chair Highlight.
- [29] **AI Automatons: AI Systems Intended to Imitate Humans.** Alexandra Olteanu, Solon Barocas, Su Lin Blodgett, Lisa Egede, Alicia DeVrio, and Myra Cheng. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2026.
- [28] **“A natural language interface for everyday life”: The social and political functions of AI capabilities discourses.** Justine Zhang, Su Lin Blodgett, and Nina Markl. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2026.
- [27] **A Framework to Characterize Reporting on Generative AI Use.** Agathe Balayn, Varun Nagaraj Rao, Su Lin Blodgett, Aylin Caliskan, and Solon Barocas. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2026.
- [26] **From Use to Oversight: How Mental Models Influence User Behavior and Output in AI Writing Assistants.** Shalaleh Rismani, Su Lin Blodgett, Q. Vera Liao, Alexandra Olteanu, and AJung Moon. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2026.
- [25] **Rigor in AI: Doing Rigorous AI Work Requires a Broader, Responsible AI-Informed Conception of Rigor.** Alexandra Olteanu, Su Lin Blodgett, Agathe Balayn, Angelina Wang, Fernando Diaz, Flavio du Pin Calmon, Margaret Mitchell, Michael Ekstrand, Reuben Binns, and Solon Barocas. *Conference on Neural Information Processing Systems (NeurIPS)*, 2025.
- [24] **Dehumanizing Machines: Mitigating Anthropomorphic Behaviors in Text Generation Systems.** Myra Cheng, Su Lin Blodgett, Alicia DeVrio, Lisa Egede, and Alexandra Olteanu. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2025. 👩 Senior Area Chair Highlight.
- [23] **Understanding and Meeting Practitioner Needs When Measuring Representational Harms Caused by LLM-Based Systems.** Emma Harvey, Emily Sheng, Su Lin Blodgett, Alexandra Chouldechova, Jean Garcia-Gathright, Alexandra Olteanu, and Hanna Wallach. *Findings of the Association for Computational Linguistics*, 2025.
- [22] **Position: Evaluating Generative AI Systems is a Social Science Measurement Challenge.** Hanna Wallach, Meera Desai, A. Feder Cooper, Angelina Wang, Chad Atalla, Solon Barocas, Su Lin Blodgett, Alexandra Chouldechova, Emily Corvi, P. Alex Dow, Jean Garcia-Gathright, Alexandra Olteanu, Nicholas J. Pangakis, Stefanie Reed, Emily Sheng, Dan Vann, Jennifer Wortman Vaughan, Matthew Vogel, Hannah Washington, and Abigail Z. Jacobs. *International Conference on Machine Learning (ICML)*, 2025.
- [21] **Measuring Machine Learning Harms from Stereotypes Requires Understanding Who Is Being Harmed by Which Errors in What Ways.** Angelina Wang, Xuechunzi Bai, Solon Barocas, and Su Lin Blodgett. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2025.

- [20] **A Taxonomy of Linguistic Expressions That Contribute To Anthropomorphism of Language Technologies.** Alicia DeVrio, Myra Cheng, Lisa Egede, Alexandra Olteanu, and Su Lin Blodgett. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2025.
- [19] **“It was 80% me, 20% AI”: Seeking Authenticity in Co-Writing with Large-Language Models.** Angel Hsing-Chi Hwang, Q. Vera Liao, Su Lin Blodgett, Alexandra Olteanu, and Adam Trischler. *ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2025.
- [18] **Metrics for What, Metrics for Whom: Assessing Actionability of Bias Evaluation Metrics in NLP.** Pieter Delobelle, Giuseppe Attanasio, Debora Nozza, Su Lin Blodgett, and Zeerak Talat. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [17] **ECBD: Evidence-Centered Benchmark Design for NLP.** Yu Lu Liu, Su Lin Blodgett, Jackie Chi Kit Cheung, Q. Vera Liao, Alexandra Olteanu, and Ziang Xiao. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [16] **Understanding the Impacts of Language Technologies’ Performance Disparities on African American Language Speakers.** Jay L. Cunningham, Su Lin Blodgett, Hal Daumé III, Christina Harrington, Hanna Wallach, and Michael A. Madaio. *Findings of the Association for Computational Linguistics*, 2024.
- [15] **“One-Size-Fits-All?” Examining Expectations about What Constitute “Fair” or “Good” NLG System Behaviors.** Li Lucy, Su Lin Blodgett, Milad Shokouhi, Hanna Wallach, and Alexandra Olteanu. *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2024.
- [14] **The Perspectivist Paradigm Shift: Assumptions and Challenges of Capturing Human Labels.** Eve Fleisig, Su Lin Blodgett, Dan Klein, and Zeerak Talat. *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2024.
- [13] **Responsible AI Considerations in Text Summarization Research: A Review of Current Practices.** Yu Lu Liu, Meng Cao, Su Lin Blodgett, Jackie Chi Kit Cheung, Alexandra Olteanu, and Adam Trischler. *Findings of the Association for Computational Linguistics: EMNLP*, 2023.
- [12] **FairPrism: Evaluating Fairness-Related Harms in Text Generation.** Eve Fleisig, Aubrie N. Amstutz, Chad Atalla, Su Lin Blodgett, Hal Daumé III, Alexandra Olteanu, Emily Sheng, Dan Vann, and Hanna Wallach. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [11] **It Takes Two to Tango: Navigating Conceptualizations of NLP Tasks and Measurements of Performance.** Arjun Subramonian, Xingdi Yuan, Hal Daumé III, and Su Lin Blodgett. *Findings of the Association for Computational Linguistics*, 2023.
- [10] **This Prompt is Measuring <MASK>: Evaluating Bias Evaluation in Language Models.** Seraphina Goldfarb-Tarrant, Eddie L. Ungless, Esma Balkır, and Su Lin Blodgett. *Findings of the Association for Computational Linguistics*, 2023.
- [9] **Taxonomizing and Measuring Representational Harms: A Look at Image Tagging.** Jared Katzman, Angelina Wang, Morgan Scheuerman, Su Lin Blodgett, Kristen Laird, Hanna Wallach, and Solon Barocas. *Annual AAAI Conference on Artificial Intelligence (AAAI)*, 2023.
- [8] **Examining Responsibility and Deliberation in AI Impact Statements and Ethics Reviews.** David Liu, Priyanka Nanayakkara, Sarah Sakha, Grace Abuhamad, Su Lin Blodgett, Nicholas Diakopoulos, Jessica Hullman, and Tina Eliassi-Rad. *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2022.
- [7] **Deconstructing NLG Evaluation: Evaluation Practices, Assumptions, and Their Implications.** Kaitlyn Zhou, Su Lin Blodgett, Adam Trischler, Hal Daumé III, Kaheer Suleman, and Alexandra Olteanu. *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2022.

- [6] **Stereotyping Norwegian Salmon: An Inventory of Pitfalls in Fairness Benchmark Datasets.** Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2021.
- [5] **A Survey of Race, Racism, and Anti-Racism in NLP.** Anjalie Field, Su Lin Blodgett, Zeerak Tatalat, and Yulia Tsvetkov. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2021.
- [4] **Language (Technology) is Power: A Critical Survey of “Bias” in NLP.** Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [3] **Twitter Universal Dependency Parsing for African-American and Mainstream American English.** Su Lin Blodgett, Johnny Wei, and Brendan O’Connor. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.
- [2] **Monte Carlo Syntax Marginals for Exploring and Using Dependency Parses.** Katherine Keith, Su Lin Blodgett, and Brendan O’Connor. *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2018.
- [1] **Demographic Dialectal Variation in Social Media: A Case Study of African-American English.** Su Lin Blodgett, Lisa Green, and Brendan O’Connor. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.

Lightly Peer-Reviewed Publications

- [6] **“I Was a Young AI”: On Probing the Effectiveness of Intervening on Anthropomorphic AI System Outputs.** Su Lin Blodgett, Myra Cheng, and Alexandra Olteanu. *Beyond Alignment: Transdisciplinary Conversations on Human-AI Futures Workshop (BATCH)*, 2026.
- [5] **“I Am the One and Only, Your Cyber BFF”: Understanding the Impact of GenAI Requires Understanding the Impact of Anthropomorphic AI.** Myra Cheng, Alicia DeVrio, Lisa Egede, Su Lin Blodgett, and Alexandra Olteanu. *International Conference on Learning Representations (ICLR) Blogposts*, 2025.
- [4] **Examining Political Rhetoric with Epistemic Stance Detection.** Ankita Gupta, Su Lin Blodgett, Justin H. Gross, and Brendan O’Connor. *Workshop on Natural Language Processing and Computational Social Science (NLP+CSS)*, 2022.
- [3] **A Dataset and Classifier for Recognizing Social Media English.** Su Lin Blodgett, Johnny Wei, and Brendan O’Connor. *Workshop on Noisy User-generated Text (W-NUT)*, 2017. 🏆 **Best Paper Award.**
- [2] **Racial Disparity in Natural Language Processing: A Case Study of Social Media African-American English.** Su Lin Blodgett and Brendan O’Connor. *Fairness, Accountability, and Transparency in Machine Learning Workshop (FAT/ML)*, 2017.
- [1] **Visualizing Textual Models with In-Text and Word-as-Pixel Highlighting.** Abram Handler, Su Lin Blodgett, and Brendan O’Connor. *Workshop on Human Interpretability in Machine Learning (WHI)*, 2016.

Non-Archival Presentations (Lightly Peer-Reviewed)

- [8] **The predatory fantasy of worker empowerment in AI marketing.** Justine Zhang, Su Lin Blodgett, and Nina Markl. *AI x Crisis: Tracing New Directions Beyond Deployment and Use Workshop (Aarhus)*, 2025.
- [7] **Evaluating Generative AI Systems is a Social Science Measurement Challenge.** Hanna Wallach, Meera Desai, Nicholas Pangakis, A. Feder Cooper, Angelina Wang, Solon Barocas, Alexandra Chouldechova, Chad Atalla, Su Lin Blodgett, Emily Corvi, P. Alex Dow, Jean Garcia-Gathright, Alexandra Olteanu, Stefanie Reed, Emily Sheng, Dan Vann, Jennifer Wortman Vaughan, Matthew Vogel, Hannah Washington, and Abigail Z. Jacobs. *Evaluating Evaluations Workshop (NeurIPS)*, 2024.

- [6] **Gaps Between Research and Practice When Measuring Representational Harms Caused by LLM-Based Systems.** Emma Harvey, Emily Sheng, Su Lin Blodgett, Alexandra Chouldechova, Jean Garcia-Gathright, Alexandra Olteanu, and Hanna Wallach. *Evaluating Evaluations Workshop (NeurIPS)*, 2024.
- [5] **How different mental models of AI-based writing assistants impact writers’ interactions with them.** Shalaleh Rismani, Su Lin Blodgett, Alexandra Olteanu, Q. Vera Liao, and AJung Moon. *Workshop on Intelligent and Interactive Writing Assistants (CHI)*, 2024.
- [4] **Representational Harms in Image Captioning and Labeling.** Jared Katzman, Solon Barocas, Su Lin Blodgett, Kristen Laird, and Hanna Wallach. *Beyond Fairness: Towards a Just, Equitable, and Accountable Computer Vision Workshop (CVPR)*, 2021.
- [3] **Debunking Debiasing: A Critique of Bias Measurement in NLP.** Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. *Women in Machine Learning Workshop (NeurIPS)*, 2019.
- [2] **Debunking Debiasing: A Critique of Bias Measurement in NLP.** Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. *Conference on New Directions in Analyzing Text as Data (TADA)*, 2019.
- [1] **Natural Language Processing Methods to Learn about Social Determinants of Health from Electronic Health Records.** David L. Chin, Vishal A. Shetty, Su Lin Blodgett, and Brendan T. O’Connor. *AcademyHealth Annual Research Meeting*, 2019.

Book Chapters

- [3] **Ethics and Bias in NLP.** Zeerak Talat and Su Lin Blodgett. In *International Encyclopedia of Language and Linguistics, 3rd Edition*, Hilary Nesi and Petar Milin (Eds.), Elsevier, 2025.
- [2] **Evaluating the Social Impact of Generative AI Systems in Systems and Society.** Irene Solaiman, Zeerak Talat, William Agnew, Lama Ahmad, Dylan Baker, Su Lin Blodgett, Canyu Chen, Hal Daumé III, Jesse Dodge, Isabella Duan, Felix Friedrich, Avijit Ghosh, Usman Gohar, Sara Hooker, Yacine Jernite, Ria Kalluri, Alberto Lusoli, Alina Leidingner, Michelle Lin, Xiuzhu Lin, Sasha Luccioni, Jennifer Mickel, Margaret Mitchell, Jessica Newman, Anaelia Ovalle, Marie-Therese Png, Shubham Singh, Andrew Strait, Lukas Struppek, and Arjun Subramonian. In *The Oxford University Press Handbook of Generative AI*, Philip Hacker, Andreas Engel, Sarah Hammer, and Brent Mittelstadt (Eds.), Oxford University Press, 2025.
- [1] **Beyond “Fairness”: Structural Injustice Lenses On AI for Education.** Michael Madaio, Su Lin Blodgett, Elijah Mayfield, and Ezekiel Dixon-Román. In *The Ethics of Artificial Intelligence in Education: Current Challenges, Practices and Debates*, Wayne Holmes and Kaśka Porayska-Pomsta (Eds.), Routledge, 2022.

Other

- [4] **LLMs produce racist output when prompted in African American English.** Su Lin Blodgett and Zeerak Talat. *Nature News and Views*, 2024.
- [3] **Fairness and Sequential Decision Making: Limits, Lessons, and Opportunities.** Samer B. Nashed, Justin Svegliato, and Su Lin Blodgett. *arXiv*, 2023.
- [2] **Risks of AI Foundation Models in Education.** Su Lin Blodgett and Michael Madaio. *arXiv*, 2021.
- [1] **How to Write a Bias Statement: Recommendations for Submissions to the Workshop on Gender Bias in NLP.** Christian Hardmeier, Marta R. Costa-jussà, Kellie Webster, Will Radford, and Su Lin Blodgett. *arXiv*, 2021.

WORKSHOPS AND TUTORIALS ORGANIZED

- [8] **Fourth Workshop on Bridging Human–Computer Interaction and Natural Language Processing.** Su Lin Blodgett, Amanda Cercas Curry, Sunipa Dev, Siyan Li, Michael Madaio, Jack Wang, Sherry Wu, Ziang Xiao, and Diyi Yang. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.
- [7] **Translation Tutorial: AI Measurement as a Stakeholder-Engaged Design Practice.** Anna Kawakami, Solon Barocas, Su Lin Blodgett, Alex Chouldechova, Abigail Jacobs, Emily Sheng, Jenn Wortman Vaughan, Hanna Wallach, Amy Winecoff, Angelina Wang, Haiyi Zhu, and Ken Holstein. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2025.
- [6] **Tutorial: Human-Centered Evaluation of Language Technologies.** Su Lin Blodgett, Jackie Chi Kit Cheung, Q. Vera Liao, and Ziang Xiao. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [5] **Third Workshop on Bridging Human–Computer Interaction and Natural Language Processing.** Su Lin Blodgett, Amanda Cercas Curry, Sunipa Dev, Michael Madaio, Ani Nenkova, Ziang Xiao, and Diyi Yang. *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2024.
- [4] **Panel: Responsible Language Technologies: Foreseeing and Mitigating Harms.** Su Lin Blodgett, Q. Vera Liao, and Alexandra Olteanu. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2022.
- [3] **Second Workshop on Bridging Human–Computer Interaction and Natural Language Processing.** Su Lin Blodgett, Hal Daumé III, Michael Madaio, Ani Nenkova, Brendan O’Connor, Hanna Wallach, and Qian Yang. *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2022.
- [2] **First Workshop on Bridging Human–Computer Interaction and Natural Language Processing.** Su Lin Blodgett, Michael Madaio, Brendan O’Connor, Hanna Wallach, and Qian Yang. *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, 2021.
- [1] **Tutorial: The Meaning and Measurement of Bias: Lessons from NLP.** Abigail Z. Jacobs, Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. *ACM Conference on Fairness, Accountability, and Transparency (FAT*)*, 2020.

INVITED TALKS AND PANELS

(Upcoming) Talk. Interactions between Formal and Computational Linguistics (ILFC) Seminar	2026
(Upcoming) Keynote. 10th Workshop on Online Abuse and Harms (WOAH), EMNLP	2026
“Towards Responsible System Design and Development with Agentic AI.” Spring School on Agentic AI for Software Engineering, University of Calgary	2026
Invited panelist. “The Human Stakes of Language Data” Panel, 4th Annual Data, Ethics & Policy Conference, Keller Center, Harris School of Public Policy, University of Chicago	2026
“Centering People in the NLP Pipeline: A Sociotechnical Perspective on NLP System Design and Evaluation.” International Laboratory on Learning Systems, École de technologie supérieure	2026
Keynote. “What Can We Learn from Perspectives on Noisy User-Generated Text?” Workshop on Noisy and User-generated Text, NAACL	2025
Keynote. “Engaging with Sociolinguistic Context in Evaluating Language Technologies.” Human-centered Evaluation and Auditing of Language Models Workshop, CHI	2025

Invited panelist. Panel on the AI Evaluation Landscape, Evaluating Evaluations: Examining Best Practices for Measuring Broader Impacts of Generative AI Workshop, NeurIPS	2024
Keynote. “Towards Equitable Language Technologies.” International Web Information Systems Engineering Conference (WISE)	2024
Invited panelist. Panels on AI for Good and on Industry, International Web Information Systems Engineering Conference (WISE)	2024
“Measurement and Disagreement in Equitable NLP.” NLP Seminar Series, Columbia University	2024
“Two Challenges for Equitable Language Technologies.” EQUAL Lab, McGill University	2024
_____. Responsible AI Seminar Series, Nokia Bell Labs	2024
“Two Challenges for Equitable Language Technologies.” Workshop on AI Driven Online Advertising: Market Design, Generative AI, and Ethics, WWW	2024
Invited panelist. Panel on the Future of Work in the Age of LLMs, Workshop on Trust and Reliance in Evolving Human-AI Workflows, CHI	2024
Invited panelist. Panel on Generative AI and Biases, Queer in AI Workshop, NeurIPS	2023
“Towards Equitable Language Technologies.” AI-GAP: Workshop on Algorithmic Biases in Artificial Intelligence from Interdisciplinary Perspectives, University of L’Aquila	2023
_____. Social, Ethical and Legal Issues in Computing Lecture Series, Queens University	2023
_____. Language Technologies Institute Colloquium, Carnegie Mellon University	2023
Invited panelist. Gender & Tech Event, University of Cambridge	2023
“Measurement and Disagreement in Equitable NLP.” Linguistics Colloquium, University of Québec at Montréal	2023
Keynote. “Challenges for Subjectivity in NLP.” 7th Workshop on Online Abuse and Harms (WOAH), ACL	2023
Invited panelist. Panel on algorithmic justice, SPUI25	2023
Keynote. “Towards Equitable Language Technologies.” Workshop on Algorithmic Injustice, University of Amsterdam	2023
Invited panelist. CRAFT Session on User Engagement in Algorithm Testing and Auditing: Exploring Opportunities and Tensions between Practitioners and End Users, FAccT	2023
“Examining How We Examine Language Technologies.” Software Engineering for Machine Learning Applications (SEMLA), Polytechnique Montréal	2023
Invited panelist. Panel on ChatGPT and Education: Hype versus Reality. University of Massachusetts Amherst	2023
“Towards Equitable Language Technologies.” School of Information Colloquium, University of Arizona	2022
_____. NSERC Responsible AI Seminar Series, Toronto Metropolitan University	2022
Keynote. “Prospects for Participatory Approaches in NLP.” 1st Workshop on Perspectivist Approaches to NLP, LREC	2022
“Fairness-Related Harms in Language Technologies.” Emerging Problems: New Challenges in FAccT from Research, to Practice, to Policy, FAccT CRAFT Session	2022
Keynote. “Towards Equitable Language Technologies.” Second Workshop on Speech and Language Technologies for Equality, Diversity, and Inclusion (LT-EDI), ACL	2022
_____. Computational Linguistics Lunch Seminar Series, University of Pennsylvania	2022
_____. Young Achievers Symposium, Center for Socially Responsible Artificial Intelligence, Pennsylvania State University	2022
_____. Linguistics Colloquium Series, University of Washington	2022

“Towards Building Equitable Language Technologies.” Cultural Assessment Working Group, Digital Library Federation	2021
“Towards Equitable Language Technologies.” Computer Science Colloquium, Williams College	2021
“Language and Justice: Social Implications of Language Technologies.” Bocconi University	2021
“Computational Approaches to Measuring Stereotyping in Language.” Workshop on Generics and Stereotypes in Language, Institute for Language, Logic and Computation, University of Amsterdam	2021
“Towards Building Equitable Language Technologies.” Responsible Data Science and AI Speaker Series, University of Illinois at Urbana-Champaign	2021
Invited panelist. Workshop on Foundation Models, Stanford Institute for Human-Centered AI (HAI)	2021
“Towards Building Equitable Language Technologies.” xAI Discussion Group, Aggregate Intellect	2021
Invited panelist, Ethics in AI Panel. Level 3 AI Assistant Conference (L3-AI)	2021
“Towards Building Equitable Language Technologies.” Annual Linguistics Student Association Spring Colloquium, San Diego State University	2021
_____. Artificial Intelligence in the Liberal Arts, Amherst College	2021
“Language and Justice: Reconsidering Harms in Language Systems and Practices.” Seminar Series, Naver Labs	2021
_____. NLP/Text-as-Data Speaker Series, New York University	2021
_____. AI Ethics Group, MIT	2021
_____. NLP Seminar, University of California, Berkeley	2021
Invited panelist, Voices of Data Science at UMass Amherst	2021
“Language and Justice: Reconsidering Harms in Language Systems and Practices.” NLP Seminar, Stanford University	2021
“Language and Justice: Understanding Harms in NLP Systems and Practices.” Computer Science & Engineering Colloquia Series, Washington University in St. Louis	2020
“Language and Justice: A Collection of Thoughts on the Social Implications of Language Technologies.” NLP with Friends	2020
“Language and Justice: Understanding Harms in NLP Systems and Practices.” Educational Testing Service	2020
“Sociolinguistically Driven Approaches for Justice in Natural Language Processing Systems.” Microsoft Research New York	2020
“Fairness and Justice in Natural Language Processing.” Womxn in Computer Science, Dartmouth College	2019
“Demographic Dialectal Variation in Social Media.” Computational Social Science Institute Seminar, University of Massachusetts Amherst	2019
_____. Workshop on New Methods in Computational Sociolinguistics, Lorentz Center, Leiden University	2018
_____. Summer Dialect Research Project, Department of Linguistics, University of Massachusetts Amherst	2018
_____. Data Science Research Symposium, University of Massachusetts Amherst	2017
_____. Computational Linguistics Workshop, Department of Linguistics, University of Massachusetts Amherst	2017
_____. Computer Science Colloquium, Wellesley College	2016

TEACHING AND MENTORING

Instructor

Artificial Intelligence (undergraduate), Co-Instructor with Jackie Cheung, McGill University Fall 2026
Ethical Issues Surrounding Artificial Intelligence Systems and Big Data (first-year seminar), Fall 2018
Co-Instructor, University of Massachusetts Amherst

Teaching Assistant

Advanced Natural Language Processing (graduate), University of Massachusetts Amherst Spring 2017

Guest Lectures

“Ethical reasoning throughout the AI lifecycle.” Responsible AI & Human Rights 2024, 2025
Summer School, Mila
“Towards Equitable Language Technologies.” Williams College 2024
“Language, Society, and Language Technologies.” Carnegie Mellon University 2023
“Language, Language Technologies, and Justice.” Georgia Institute of Technology 2022
“Towards Equitable Language Technologies.” University of Massachusetts Amherst 2021
“Language and Justice.” University of Wisconsin-Madison 2020
“NLP and Social Good.” University of Massachusetts Amherst 2020
“Ethics in NLP.” University of Massachusetts Amherst 2020
———. Mount Holyoke College 2020
“Ethics in Machine Learning.” Mount Holyoke College 2020
“Variation in Social Media.” Mount Holyoke College 2017
———. University of Massachusetts Amherst 2017

Interns and Students Mentored or Co-Mentored

Varun Rao, Princeton University, MSR research internship 2025
Shravika Mittal, Georgia Institute of Technology, MSR research internship 2025
Alicia DeVrio, Carnegie Mellon University, MSR research internship 2024
Myra Cheng, Stanford University, MSR research internship 2024
Lisa Egede, Carnegie Mellon University, MSR research internship 2024
Rajiv Movva, Cornell Tech, joint MSR and Experiences & Devices (JEM) internship 2024
Gaurav Verma, Georgia Institute of Technology, MSR research internship 2023
Yu Lu Liu, McGill University, Mila and MSR collaboration 2022 – 2024
Arjun Subramonian, University of California, Los Angeles, MSR research internship 2022
Jules Barbe, McGill University, Mila and MSR collaboration 2021 – 2022
Lucy Li, University of California, Berkeley, MSR research internship 2021
Kaitlyn Zhou, Stanford University, MSR research internship 2021
Jay Cunningham, University of Washington, MSR research internship 2021

Ph.D. Committees

Marta Marchiori Manerba, University of Pisa 2025
Zhila Aghajari, Lehigh University 2025
Akshita Jha, Virginia Tech 2024

FUNDING

Investigating Interaction-Based Harms in AI Systems. AJung Moon, Shalaleh Rismani, Q. Vera Liao, Alexandra Olteanu, Koustuv Saha, and Su Lin Blodgett. MSR & Mila Research Collaboration. Award

to academic partner \$75,000 CAD. 2022.

Mapping FATE Concerns for Text Summarization. Su Lin Blodgett, Jackie C.K. Cheung, Alexandra Olteanu, and Adam Trischler. MSR & Mila Research Collaboration. Award to academic partner \$55,000 CAD. 2021.

SERVICE AND OUTREACH

Conference Chair Roles

General Chair, ACM Conference on Fairness, Accountability, and Transparency (FAccT)	2026
Ethics Chair, International Conference on Computational Linguistics (COLING)	2025
Student Program Committee Chair, AAAI Conference on AI, Ethics, and Society (AIES)	2023
Ethical Advisory Committee Chair, Annual Meeting of the Association for Computational Linguistics (ACL)	2022

Executive and Advisory Roles

Executive Committee, ACM Conference on Fairness, Accountability, and Transparency (FAccT)	2026 –
Advisory Board, UKRI AI Centre for Doctoral Training in Designing Responsible NLP, University of Edinburgh	2025 –
Advisory Board, Socially Responsible Language Modelling Research (SoLaR) Workshop	2023

Conference and Journal Reviewing

Area Chair: AACL, ACL Rolling Review (ARR), EMNLP, FAccT, Black in AI Workshop

Conference Reviewer: ACL, AIES, ARR, CoNLL, EACL, EMNLP, FAccT, ICWSM, NAACL, WWW

Journal Reviewer: ACM Computing Surveys, Critical AI, Nature, PLOS One, TACL

Workshop Reviewer: Data Science with Human in the Loop, Gender Bias in Natural Language Processing, NLP for Positive Impact, Teaching NLP, Women in Machine Learning, Workshop on Noisy User-generated Text, Workshop on Online Abuse and Harms

Ethics Reviewer: ARR, EMNLP, NeurIPS

Funding Application Reviewing

ERC Consolidator Grant, MITACS Accelerate

Public Engagement and Media Mentions

Presenter, Microsoft Research Forum	2025
“How to build a better AI benchmark.” MIT Technology Review	2025
Guest, Unhinged Collaboration Podcast	2023
Guest, McGill AI Podcast	2022
Guest, The Good Robot Podcast	2022
Guest, Rasa Chats Podcast	2021
Guest, Radical AI Podcast	2021
“Microsoft researchers say NLP bias studies must consider role of social hierarchies like racism.” VentureBeat	2020
“AI Programs Are Learning to Exclude Some African-American Voices.” MIT Technology Review	2017
“A.I. Doesn’t Get Black Twitter.” Inverse	2016

Other

Panelist, Application Guidance Panel, Women@Mila	2023
Moderator, Fairness at the Intersections Tutorial Panel (NeurIPS)	2022

Panelist, UMass CICS CS Women	2022
Mentoring group leader, EACL	2021
Mentoring group leader, EMNLP	2020
Ph.D. Applicant Support Program mentor, UMass CICS	2020
Volunteer, Women in Machine Learning Workshop (WiML)	2019
Senior Ph.D. Student Panelist, UMass CICS CS Women	2019
Succeeding in Graduate School Panelist, UMass Graduate Orientation	2018
Volunteer, ACL	2018
Co-organizer, Inaugural Male Allyship Workshop, UMass Computer Science	2017
Panelist, Mount Holyoke Graduate School Panel	2017
Social Media Chair, UMass CS Women	2016 – 2017
Mentoring Committee Member, UMass Graduate Women in STEM	2016 – 2017
Volunteer, EMNLP	2016
Volunteer, UMass Women in Engineering and Computing Career Day	2015