robert.wolfe[at]rutgers.edu
wolferobert3.github.io
github.com/wolferobert3

Robert Wolfe

Employment

2025 – **Rutgers University** – New Brunswick, NJ
Assistant Professor, Tenure Track
School of Communication and Information (SC&I)

Education

2021 – 2025 **University of Washington** – Seattle, WA
Ph.D., Information Science
M.S., Information Science
Dissertation: *Approaches to Epistemic Risk in Generative and General Purpose AI*
Committee: Dr. Bill Howe (Co-Chair), Dr. Alexis Hiniker (Co-Chair), Dr. Tanushree Mitra, Dr. Leilani Battle (GSR)

2019 – 2021 **The George Washington University** – Washington, D.C.
M.S., Computer Science
Thesis: *The Valence-Assessing Semantics Test for Contextualizing Language Models*
Committee: Dr. Aylin Caliskan (Advisor), Dr. Robert Pless, Dr. Abdou Youssef

2012 – 2014 **Georgetown University** – Washington, D.C.
M.A., English Literature
Thesis: *Driven by Difference: The Embodiment of the Western Maryland Initiative*
Committee: Dr. Patricia O'Connor (Advisor), Dr. Matthew Pavesich

2008 – 2012 **University of Maryland** – College Park, MD
B.A., English Literature
Phi Beta Kappa, University Honors

Peer-Reviewed Conference Publications

2025 [c.28] **Building the Beloved Community: Designing Technologies for Neighborhood Safety.**
Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI) 2025.
Ishita Chordia, Robert Wolfe, Carl DiSalvo, Jason Yip, and Alexis Hiniker.
Summary: Study probing the development of justice-oriented safety technologies to support collective welfare rather than individualistic notions of safety. Leverages principles of Transformative Justice to offer a novel conceptualization of safety focused on the basic needs of the community.
Acceptance Rate: 24.9% Overall. **Best Paper Honorable Mention (Top 5.0%)**.

[c.27] **Fragments to Facts: Partial-Information Fragment Inference from LLMs.**
Proceedings of the International Conference on Machine Learning (ICML) 2025.
Lucas Rosenblatt, Bin Han, Robert Wolfe, and Bill Howe.
Summary: Study developing novel methods using partial, unordered information to probe the privacy vulnerabilities of instruction-tuned large language models.
Acceptance Rate: 26.9% Overall.

[c.26] **Do Language Models Mirror Human Confidence? Exploring Psychological Insights to Address Overconfidence in LLMs.**
Findings of the Association for Computational Linguistics (ACL) 2025.
Chenjun Xu, Bingbing Wen, Bin Han, Robert Wolfe, Lucy Lu Wang, and Bill Howe.
Summary: Study developing confidence calibration methods to mitigate humanlike overconfidence patterns in large language models.

[c.25] **Privacy as Social Norm: Systematically Reducing Dysfunctional Privacy Concerns on Social Media.**
ACM SIGCHI Conference on Computer-Supported Cooperative Work & Social Computing (CSCW) 2025.
JaeWon Kim, Soobin Cho*, Robert Wolfe*, Jishnu Hari Nair, and Alexis Hiniker.
Summary: Study developing a theory of dysfunctional fear among adolescents on social media platforms, along with prototypes empowering adolescents to overcome that fear.

[c.24] **Trust-Enabled Privacy: Social Media Designs to Support Adolescent User Boundary Regulation.**
Symposium on Usable Privacy and Security (SOUPS) 2025.
JaeWon Kim, Robert Wolfe, Ramya Bhagirathi Subramanian, Mei-Hsuan Lee, Jessica Colnago, and Alexis Hiniker.
Summary: Study introducing the trust-enabled privacy design framework to support boundary regulation on social media platforms, reframing online privacy as a trust-driven process rather than a control-based tradeoff.

[c.23] **"AI just keeps guessing": Using ARC Puzzles to Help Children Identify Reasoning Errors in Generative AI.**
Interaction Design and Children (IDC) 2025.
Aayushi Dangol, Runhua Zhao*, Robert Wolfe*, Trushaa Ramanan, Julie Kientz, and Jason Yip.
Summary: Study using Abstraction and Reasoning Corpus (ARC) puzzles to help children understand the shortcomings of modern generative AI, Mayer & Moreno's Cognitive Theory of Multimedia Learning.

[c.22] **Children's Mental Models of AI Reasoning: Implications for AI Literacy Education.**
IDC 2025.
Aayushi Dangol, Robert Wolfe, Runhua Zhao, JaeWon Kim, Trushaa Ramanan, Katie Davis, and Julie Kientz.
Summary: Large empirical study situating the implications of advances in AI reasoning models for the field of AI literacy.

[c.21] **"If anybody finds out you are in BIG TROUBLE": Understanding Children's Hopes, Fears, and Evaluations of Generative AI.**
IDC 2025.
Aayushi Dangol, Robert Wolfe, Daeun Yoo, Arya Thiruvillakkat, Ben Chickadel, and Julie Kientz.
Summary: Study working with elementary school students that finds that children envision generative AI as a companion, as a collaborator, and as an automator, with description of children's hopes and fears through these lenses.

[c.20] **Doors, Decisions, and Discovery: An Interactive Smart Door Lock System to Promote Children's Understanding of AI Classification.**
International Society of the Learning Sciences (ISLS) 2025.
Aayushi Dangol, Robert Wolfe, Rotem Landesman, Jason Yip, and Julie Kientz.
Summary: Study introducing a CLIP-based system for helping children learn about zero-shot multimodal AI classification, along with implications for fairness.

2024 [c.19] **The Impact and Opportunities of Generative AI in Fact-Checking**
ACM Conference on Fairness, Transparency, and Accountability (FAccT) 2024.
Robert Wolfe and Tanushree Mitra.
Summary: Interview study with N=38 fact-checkers that catalogues the in-use, in-progress, and envisioned uses of generative AI in fact-checking, along with epistemic challenges preventing further use. Introduces the dimension of Verification to the design space of generative AI, and takes a value-sensitive approach to mapping tensions between generative AI and fact-checking.
Acceptance Rate: 24% Overall, 28% Users and Experiences Track.

[c.18] **The Implications of Open Generative Models in Human-Centered Data Science Work: A Case Study with Fact-Checking Organizations.**
AI Ethics and Society (AIES) 2024.
Robert Wolfe and Tanushree Mitra.
Summary: Study of open generative models in fact-checking organizations, contextualizing the motivations and challenges of fact-checkers within the data science pipelines used by these organizations.
Acceptance Rate: 31.8% Overall.

[c.17] **Laboratory-Scale AI: Open-Weight Models are Competitive with ChatGPT Even in Low-Resource Settings.**
ACM FAccT 2024.
Robert Wolfe, Isaac Slaughter, Bin Han, Bingbing Wen, Yiwei Yang, Lucas Rosenblatt, Bernease Herman, Eva Brown, Zening Qu, Nic Weber, and Bill Howe.
Summary: Study of the viability of small, open generative models as an alternative to large, proprietary, closed models. Finds that open models are competitive given a relatively small amount of data, and offer benefits in the form of cost-efficiency, differential privacy, and tunable abstention properties (reducing hallucination).
Acceptance Rate: 24% Overall, 22% Systems Track.

[c.16] **Dataset Scale and Societal Consistency Mediate Facial Impression Bias in Vision-Language AI.**
AIES 2024.
Robert Wolfe, Aayushi Dangol, Bill Howe, and Alexis Hiniker.
Summary: Study of the factors affecting the presence of facial impression bias in 43 multimodal CLIP models, as well as the reproduction of facial impression biases by generative multimodal models such as Stable Diffusion.
Acceptance Rate: 31.8% Overall.

[c.15] **Representation Bias of Adolescents in AI: A Bilingual, Bicultural Study.**
AIES 2024.
Robert Wolfe, Aayushi Dangol, Bill Howe, and Alexis Hiniker.
Summary: Study comparing biases about adolescents learned by AI to similar biases identified in traditional and news media sources in both the U.S. and Nepal. Conducts workshops with 13 U.S. teenagers and 18 Nepalese teenagers to understand how teenagers themselves view fair representation in media and AI.
Acceptance Rate: 31.8% Overall.

[c.14] **ML-EAT: A Multilevel Embedding Association Test for Interpretable and Transparent Social Science.**
AIES 2024.
Robert Wolfe, Alexis Hiniker, and Bill Howe.
Summary: Research introducing the Multilevel Embedding Association Test (ML-EAT), a method designed to address issues of ambiguity and difficulty in interpreting the traditional EAT measurement by quantifying bias at several levels of increasing granularity.
Acceptance Rate: 31.8% Overall.

[c.13] **Label-Efficient Group Robustness via Out-of-Distribution Concept Curation.**
IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024.
Yiwei Yang, Anthony Zhe Liu, Robert Wolfe, Aylin Caliskan, and Bill Howe.
Summary: Introduces a Concept Distributively Robust Optimization (DRO) framework that takes curated sets of images for a given concept to estimate group labels, and uses those labels to train with a SOTA DRO objective, significantly reducing classifier biases with relatively small, manually curated sets of images.
Acceptance Rate: 24% Overall.

[c.12] **Mediating Culture: Cultivating Socio-cultural Understanding of AI in Children through Participatory Design.**
ACM Conference on Designing Interactive Systems (DIS) 2024.
Aayushi Dangol, Michelle Newman, Robert Wolfe, Jin Ha Lee, Jason Yip, Julie Kientz, and Caroline Pitt.
Summary: Introduces participatory approach to co-designing AI with kids in ways that facilitate an understanding of AI as a mediator of culture. My involvement included building a prototype vision-language AI system, as well as helping to run human subjects sessions.
Acceptance Rate: 27% Overall.

[c.11] **"Sharing, Not Showing Off": How BeReal Encourages Authentic Self-Presentation on Social Media Through Its Design.**
CSCW 2024.
JaeWon Kim, Robert Wolfe, Ishita Chordia, Katie Davis, and Alexis Hiniker.
Summary: Introduces a set of design guidelines for creating social media platforms that support authentic self-presentation online, such as scaffolding reciprocity and expanding beyond spontaneous photo-sharing to allow users to more accurately and comfortably portray themselves.

2023 [c.10] **Contrastive Language-Vision AI Models Pretrained on Web-Scraped Multimodal Data Exhibit Sexual Objectification Bias.**
ACM FAccT 2023.
Robert Wolfe, Yiwei Yang, Bill Howe, and Aylin Caliskan.
Summary: Mixed methods study of sexually objectifying biases that attend representations of women and girls in multimodal AI models. Traces biases from the embedding space of 9 CLIP models, and in the output of generative text-to-image models such as VQGAN-CLIP and Stable Diffusion.
Acceptance Rate: 25% Overall.

[c.9] **Evaluating Biased Attitude Associations of Language Models in an Intersectional Context.**
AIES 2023.
Shiva Omrani Sabbaghi, Robert Wolfe, and Aylin Caliskan.
Summary: Introduces an SVM-based method for learning subspaces corresponding to human attitudes and concepts in causal and bidirectional transformer language models, and analyzes the human attitudes and biases reflected in the models in an intersectional context.
Acceptance Rate: 28% Overall.

2022 [c.8] **Contrastive Multimodal Pretraining Magnifies the Semantics of Natural Language Representations.**
Proceedings of the Association for Computational Linguistics (ACL) 2022.
Robert Wolfe and Aylin Caliskan.
Summary: Intrinsic evaluation of the surprising properties of contextualized word embeddings and sentence embeddings formed by the CLIP text encoder in comparison with those formed by GPT-2.
Acceptance Rate: 21% Overall. Selected for **Oral Presentation** (Top 8% of papers).

[c.7] **VAST: The Valence-Assessing Semantics Test for Contextualizing Language Models.**
Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI) 2022.
Robert Wolfe and Aylin Caliskan.
Summary: Method for intrinsic evaluation of contextualized word embeddings using human-rated psycholinguistic measurements. Originally a master's thesis at GWU.
Acceptance Rate: 15% Overall.

[c.6] **Evidence for Hypodescent in Visual Semantic AI.**
ACM FAccT 2022.
Robert Wolfe, Mahzarin R. Banaji, and Aylin Caliskan.
Summary: Study of biases in CLIP as they pertain to the perception of multiracial individuals. Used a generative adversarial network to replicate an experiment from experimental psychology, and showed that CLIP has learned an analogue of the rule of hypodescent, or one-drop rule.
Acceptance Rate: 26% Overall, 19% for Track.

[c.5] **Markedness in Visual Semantic AI.**
ACM FAccT 2022.
Robert Wolfe and Aylin Caliskan.
Summary: Study of biases in the multimodal language-and-image AI model CLIP. Examines the proclivity of CLIP to mark the race, gender, and age of some individuals while leaving it unmarked for dominant social groups.
Acceptance Rate: 26% Overall, 23% for Track.

[c.4] **American==White in Multimodal Language-and-Image AI**.
AIES 2022.
Robert Wolfe and Aylin Caliskan.
Summary: Study of biases in three multimodal language-and-image AI models: CLIP, SLIP, and BLIP. Shows that language-and-image AI learns statistically veridical information about state-level demographic distributions. Also demonstrates that some regions, such as the U.S., become associated in AI with a dominant social group – in this case, White individuals.
Acceptance Rate: 34%. Selected to **Open the Conference.**

[c.3] **Gender Bias in Word Embeddings: A Comprehensive Overview of Syntax, Frequency, and Semantics.**
AIES 2022.
Aylin Caliskan, Pimparkar Parth Ajay, Tessa Charlesworth, Robert Wolfe, and Mahzarin R. Banaji.
Summary: Study of gender biases in GloVe and fastText word embeddings taking into account the many properties of language often overlooked when examining word embedding bias. Wrote the first draft of the paper, contributed code and data for the project, mentored junior scientist.
Acceptance Rate: 34% Overall.

[c.2] **Detecting Emerging Associations and Behaviors Using Regional and Diachronic Word Embeddings.**
International Conference on Semantic Computing (ICSC) 2022.
Robert Wolfe and Aylin Caliskan.
Summary: Methods based on the word embedding association test (WEAT) for detecting changes in semantics over time in Dirichlet-smoothed word embeddings trained on low-resource Twitter corpora.

2021 [c.1] **Low-Frequency Names Exhibit Bias and Overfitting in Contextualizing Language Models.**
Proceedings of the ACL: Empirical Methods in Natural Language Processing (EMNLP) 2021.
Robert Wolfe and Aylin Caliskan.
Summary: Study of bias in contextualized word embeddings based on the correspondence of frequency, intrinsic bias, and self-similarity.
Acceptance Rate: 26% Overall.

Workshop Publications

2025 [w.5] **Partial-Information Fragment Inference from LLMs.**
Theory and Practice of Differential Privacy (TPDP) 2025.
Lucas Rosenblatt, Bin Han, Robert Wolfe, and Bill Howe.
Summary: Workshop paper introducing the partial-information fragment paradigm for LLM privacy.

[w.4] **Challenges and Opportunities for Participatory Design of Conversational Agents for Young People's Wellbeing.**
ACM IDC 2025 Workshop on Designing AI for Children's Wellbeing.
Natalia Kucirkova, Alexis Hiniker, Megumi Ishikawa, Sho Tsuji, Aayushi Dangol, and Robert Wolfe.
Summary: Workshop paper outlining the challenges and opportunities of studies of young people and conversational agents across four countries.

2024 [w.3] **Expertise Fog on the GPT Store: Deceptive Design Patterns in User-Facing Generative AI.**

ACM CHI 2024 Workshop on Mobilizing Research and Regulatory Action on Dark Patterns and Deceptive Design Practices.

Robert Wolfe and Alexis Hiniker.

Summary: Position paper with light empirical results arguing that the design of the OpenAI GPT Store encourages deceptive design patterns related to the presentation of expertise in customized versions of ChatGPT.

[w.2] **From Human to Model Overconfidence: Evaluating Confidence Dynamics in Large Language Models.**

NeurIPS 2024 Workshop on Behavioral Machine Learning.

Bingbing Wen, Chenjun Xu, Bin Han, Robert Wolfe, Lucy Lu Wang, and Bill Howe.

Summary: Proposes a method known as Answer-Free Confidence Estimation (AFCE) to improve estimation of language model confidence based on models' "verbalized" natural language descriptions.

2023 [w.1] **Regularizing Model Gradients with Concepts to Improve Robustness to Spurious Correlations.**

ICML 2023 Workshop on Spurious Correlations, Invariance, and Stability.

Yiwei Yang, Anthony Zhe Liu, Robert Wolfe, Aylin Caliskan, and Bill Howe.

Summary: Proposes a method known as CReg to penalize a machine learning model's sensitivity to a protected attribute, including the absence of group labels at the dataset level, outperforming the use of Empirical Risk Minimization (ERM) for regularization.

Under Submission

[u.4] **Understanding Privacy Norms Around LLM-Based Chatbots: A Contextual Integrity Perspective.**

Sarah Tran, Hongfan Lu, Isaac Slaughter, Bernease Herman, Aayushi Dangol, Yue Fu, Lufei Chen, Biniyam Gebreyohannes, Bill Howe, Alexis Hiniker, Nicholas Weber and Robert Wolfe.

Summary: Large-scale, preregistered survey study of how chatbot users perceive the privacy norms around the data they exchange through modern chat interfaces. Analysis of data collected via factorial vignettes finds that users largely defer to consistent baseline privacy expectations that they transfer from other contexts, rather than contextual factors relevant specifically to chatbots.

Note: I am the senior author on this paper.

[u.3] **Toward Needs-Conscious Design: Co-Designing a Human-Centered Framework for AI-Mediated Communication.**

Robert Wolfe, Aayushi Dangol, Jaewon Kim, and Alexis Hiniker.

Summary: Study introducing a human-centered framework for AI-mediated communication that builds on the principles of Nonviolent Communication (NVC).

[u.2] **Reading AI and Reading the World: Using an Interactive AI System to Promote Children's Understanding of AI Bias.**

Aayushi Dangol, Robert Wolfe, Akeiylah Dewitt, Ben Chickadel, Julie Kientz, and Sayamindu Dasgupta.

Summary: Introduces the interactive CLIP4KIDS system and studies how students understand AI biases in terms of "assumptions" and "stereotypes," drawing connections between historical injustices and present implicit biases in AI models.

[u.1] **Volitional AI: The Epistemic Risks of Generative AI Demand Visible Leadership.**
Robert Wolfe, Tanushree Mitra, and Bill Howe.
Summary: Position paper regarding the unique impact of generative AI on the public trust. Argues that generative and general-purpose models must be approached as epistemic infrastructure.

Grants, Awards, and Honors

2024 **UW iSchool Conference Travel Award**: $2,000.
Awarded $2,000 to cover conference travel for presentation of two first-author papers at ACM FAccT 2024.

UW iSchool Conference Travel Award: $658.
Awarded $658 to cover expenses related to travel for participation in the ACM CHI 2024 Dark Patterns workshop.

2023 **UW iSchool Strategic Research Initiative Award**: $15,000.
Awarded $15,000 grant co-written with PI Dr. Bill Howe, entitled Laboratory-Scale AI. Proposes empirical validation of domain-specific, instruction-tuned open models for their competitiveness with large, general, proprietary models like GPT-4. **Resulting research published at ACM FAccT 2024.**

Google Research: $60,000.
Awarded $60,000 grant co-written with PI Dr. Alexis Hiniker, entitled Encouraging Nonviolent Communication in Online Messaging Platforms. Proposes user-centered design of AI-driven technologies for promoting empathy and nonviolence. Research under submission.

2022 **UW iSchool Conference Travel Award**: $2,000.
Awarded $15,000 to cover conference travel for presentation of two first-author papers at ACM FAccT 2022.

2012-2014 **CNDLS Fellowship, Georgetown University**: Full tuition fellowship and stipend.
Full tuition fellowship and stipend covering two years in the English M.A. program. Connected to Fellow position at the Georgetown Center for New Designs in Learning and Scholarship (CNDLS), an organization to further technological approaches to higher education.

2008-2010 **Dean's Scholarship, University of Maryland**: $5,000 yearly.
Awarded $5,000 yearly tuition scholarship for the first two years of undergraduate studies.

Invited Research Talks

Oct-Dec 2024 **The Reliability Challenge of Generative AI**.
Rutgers University, SUNY Binghamton, SUNY Albany, Gonzaga University, American University, Oregon State University, College of the Holy Cross.
Talk on the challenges and opportunities of generative AI for building more reliable and democratic applications in the service of society.

August 2022 **Quantifying Biases and Societal Defaults in Word Embeddings and Language-Vision AI**.
National Institute of Standards and Technology. AI Metrology Colloquium Series.
Talk and discussion of recent research on biases in vision-language AI models, and comprehensive analyses of gender bias in word embeddings.

June 2022 **Manifestations of Implicit Biases in Language-and-Image AI**.
The Santa Fe Institute for Complex Systems. Language as a Window into Human Minds.
Overview of recent research for a small group of psychologists, decision scientists, and computer scientists at a meeting intended to foster interdisciplinary perspectives.

Research Mentorship

2024-2025 **Research Lead: User Empowerment Lab Directed Research Group**.
Research Topic: Epistemic Risk in General-Purpose AI.
Leading a DRG for undergraduates in the UW iSchool to study questions about epistemic risk in cutting-edge AI systems. Created syllabus and materials intended to facilitate development of students as capable researchers. Led to submission of a peer-reviewed conference paper first-authored by a DRG student.

Fall 2024 **Research Mentor: Responsible AI for Ukraine**.
Research Topic: Customizing Small LLMs for Scientific Reasoning Tasks.
Mentoring a project intended to create research opportunities for Ukrainian Catholic University students in coordination with New York University's Center for Responsible AI.

Press Coverage

May 2025 **UW HCDE Impact**. Helping kids think critically about AI. By Leah Pistorius.
Refers to AI Just Keeps Guessing (published at IDC 2025, first author Aayushi Dangol).

January 2025 **University of Washington News**. Study finds strong negative associations with teenagers in AI models. By Stefan Milne.
Refers to Representation Bias of Teenagers in AI (published at AIES 2024).

January 2025 **GeekWire**. UW study of how AI models portray teens finds strong negative associations. By Kurt Schlosser.
Refers to Representation Bias of Teenagers in AI (published at AIES 2024).

April 2023 **Business Insider**. Stable Diffusion and DALL-E display bias when prompted for artwork of 'African workers' versus 'European workers'. By Thomas Maxwell.
Refers to Markedness in Visual Semantic AI (published at ACM FAccT 2022).

April 2023 **The Intercept**. AI Art Sites Censor Prompts About Abortion. By Debbie Nathan.
Refers to Gender Bias in Word Embeddings (published at AIES 2022).

January 2023 **Insider**. ChatGPT could be used for good, but like many other AI models, it's rife with racist and discriminatory bias. By Hannah Getahun.
Refers to Markedness in Visual Semantic AI (presented at ACM FAccT 2022).

December 2022 **MIT Tech Review**. The viral AI avatar app Lensa undressed me – without my consent. By Melissa Heikkila.
Refers to Markedness in Visual Semantic AI (presented at ACM FAccT 2022).

Lab Memberships

2025 – Present Computational Social Science Lab (CSS). Rutgers University.

2023 – 2025 Volitional AI Lab (Howe Lab). UW iSchool.

2022 – 2025 User Empowerment Lab (Hiniker Lab). UW iSchool.

Academic Service

2025 AAAI 2025. **Program Committee**.
ACM FAccT 2025. **Program Committee**.
AIES 2025. **Program Committee**.
ACM CHI 2025. Reviewer.

2024 AIES 2024. **Program Committee**.
AAAI 2024. **Program Committee**.
ACL Rolling Review 2024. Reviewer.
NeurIPS 2024. Reviewer.

2023 AAAI 2023. **Program Committee**.
NeurIPS 2023. Reviewer.
Nature Humanities and Social Science. Reviewer.

2022 ACM FAccT 2022. **Program Committee**.
ICML 2022. Reviewer.
NeurIPS 2022. Reviewer.
AIES 2022. Reviewer.

Teaching Experience

Spring 2023 **Teaching Assistant, Informatics 371: Advanced Methods in Data Science (University of Washington)**.
*Average student rating: 4.9/5. (**Top 10-20% of TAs** at the University of Washington).*
Summary of Role: Managing lab sessions, updating programming and data analysis assignments, grading. Course content covers causality, machine learning, introductory Bayesian statistics, and introductory deep learning, including the fundamentals of computer vision and NLP. TA to Dr. Ott Toomet.

Winter 2023 **Teaching Assistant, Informatics 371: Core Methods in Data Science (University of Washington)**.
*Average student rating: 4.9/5. (**Top 10-20% of TAs** at the University of Washington).*
Summary of Role: Managing lab sessions, updating programming and data analysis assignments, grading. Course content covers data analysis, statistical inference, regression, and basic machine learning. TA to Dr. Ott Toomet.

Fall 2022 **Teaching Assistant, Informatics 270: Data Reasoning (University of Washington)**.
Average student rating: 4.6/5.
Developed and gave two class lectures on AI bias and epistemic risk; **created Jupyter notebook used as course material** for teaching statistical bias in AI to all discussion sections; participated in panel discussion on scientific uses of AI with course professors; taught discussion sessions, graded assignments. TA to Dr. Jevin West and Dr. Carl Bergstrom.

Invited Classroom Lectures and Discussions

Spring 2024 **Invited Class Discussion, Information Management 598: Epistemological Foundations of AI (University of Washington)**.
Gave invited discussion on recent approaches to the use of generative AI in professional fact-checking, contextualizing uses within class materials on AI epistemologies. With Dr. Bill Howe.

Winter 2023 **Invited Guest Lecture and Workshop, Informatics 466: Moral Reasoning and Interaction Design (University of Washington)**.
Developing lecture and class materials to support students in learning about the ethical consequences of technical design decisions, with particular attention to generative AI. Gave a lecture on pragmatism in design and ran a class workshop. With Dr. Alexis Hiniker.

Spring 2022 **Invited Guest Lecture, Information Management 575: Machine Learning 3: Applications, Scaling, and Ethics (University of Washington)**.
Developed and gave the invited guest lecture on Data Science for Social Good, covering historical trends on applications of AI for good, and recent research on bias in AI.

Community Involvement

2023-2024 **University Children's Development School, Seattle, WA.**
Developed a system for allowing children to interact with multimodal AI systems and reason about AI fairness in a controlled context. System used as part of the curriculum for upper-elementary school children.

2021-2022 **Responsible AI Systems and Experiences (RAISE), University of Washington. Student Organizing Committee.**
Organized RAISE speaker series in coordination with student volunteers and RAISE faculty.

2018-2021 **Trail Ranger, Montgomery County, MD.**
Assumed primary responsibility for condition of a system of local trails, in coordination with county officials.

2016 **Presidential Election Judge, Montgomery County, MD.**
Assisted in facilitating voting in 2016 presidential election.

Professional Memberships

2022 – Present	Association for Computing Machinery (ACM)
	Association for the Advancement of Artificial Intelligence (AAAI)
	Institute of Electrical and Electronics Engineers (IEEE)
2021 – Present	Association for Computational Linguistics (ACL)