

Amanda Coston
acoston@cs.cmu.edu 703-401-1212

Education

Carnegie Mellon University, Ph.D. student, Machine Learning and Public Policy

Research: algorithmic decision-making and decision-support systems; causal inference, counterfactual predictions, doubly-robust estimation, selection bias, missing data; algorithmic fairness, audits, and evaluation

Advisors: Alexandra Chouldechova and Edward H. Kennedy

Carnegie Mellon University, Masters of Science in Machine Learning, 2019

Princeton University, Bachelor of Science in Engineering, 2013

Magna cum laude in computer science

Certificate in the Princeton School of Public and International Affairs

Senior Thesis: Machine Learning Techniques for the Diagnosis of Pediatric Tuberculosis.

Advisor: Robert Schapire

Lead Author Research Papers

Characterizing Fairness over the Set of Good Models Under Selective Labels with co-lead author Ashesh Rambachan and with Alexandra Chouldechova. International Conference on Machine Learning 2021.

Leveraging Administrative Data for Bias Audits: Assessing Disparate Coverage with Mobility Data for COVID-19 Policy with Daniel E. Ho et al. ACM Conference on Fairness, Accountability, and Transparency 2021.

- Feat. in *Wall Street Journal* “Smartphone Location Data Can Leave Out Those Most Hit by Covid-19,” April 5, 2021.
VentureBeat “Stanford and Carnegie Mellon find race and age bias in mobility data that drives COVID-19 policy” Nov. 18, 2020.

Counterfactual Predictions under Runtime Confounding with Alexandra Chouldechova and Edward H. Kennedy. Neural Information Processing Systems 2020.

Counterfactual Risk Assessments, Evaluation, and Fairness with Alexandra Chouldechova, Edward H. Kennedy, Alan Mishler. ACM Conference on Fairness, Accountability, and Transparency 2020.

Fair Transfer Learning with Missing Protected Attributes with Kush Varshney et al. at IBM Research. AAAI / ACM Conference on Artificial Intelligence, Ethics, and Society 2019.

Other Research Papers and Presentations

Conditional Learning of Fair Representations. Han Zhao, Amanda Coston, Tameem Adel, Geoff J. Gordon. International Conference on Learning Representations 2020.

Neural Topic Models with Survival Supervision: Jointly Predicting Time-to-Event Outcomes and Learning How Clinical Features Relate. Linhong Li, Ren Zuo, Amanda Coston, Jeremy C. Weiss, George H. Chen. International Conference on Artificial Intelligence in Medicine 2020.

Offline Heterogeneous Policy Evaluation: A Causal Approach with Leqi Liu. Causal ML Workshop 2018.

Awards and Fellowships:

K&L Gates Presidential Fellow in Ethics and Computational Technologies

NSF Graduate Research Fellow

Tata Consultancy Services (TCS) Presidential Fellowship

Suresh Konda Best First Paper Award for 2019

Awarded for *Counterfactual Risk Assessments and Evaluation for Child Welfare Screening*

Carolyn Comer Graduate Student Involvement Award 2020
Awarded for co-organizing a Personal Protective Equipment drive in response to COVID-19

Phi Beta Kappa, Tau Beta Pi

Teaching

Teaching Assistant for 10-301/10601 Introduction to Machine Learning Spring 2021
Assisted Carnegie Mellon University course taught by Matt Gormley and Tom Mitchell.

Project Leader at AI4ALL Summer 2019
Organized and led a project on fairness in algorithmic risk assessments for high schoolers

Service

Area Chair for Responsible AI workshop at ICLR Spring 2021

Machine Learning for Development (ML4D) Workshop Co-organizer for NeurIPS 2018 and 2019
Workshop for methods and applications of machine learning for problems in the developing world

Program Committee Member/Reviewer

JRSS-B, Data Mining and Knowledge Discovery, AISTATS 2022, ICLR 2022, ICML 2020-2021, FAccT 2020-2021, NeurIPS 2020-2021, AIES 2020

Fairness, Ethics, Accountability, and Transparency Reading Group at Carnegie Mellon Fall 2019—Spring 2020
Co-organizer

Work

Facebook FAIAR Summer 2021
Research intern at Responsible AI. Conducted a creator-centric fairness assessment of Instagram Reels.

RegLab Summer Research Fellow Summer 2020
Researcher at the Regulation, Evaluation, and Governance Lab at Stanford Law School

IBM Research AI, Science for Social Good Fellow Summer 2018
Developed methods for fair risk assessments in domain adaptation when access to the protected attribute is limited (to either source or target)

Hivisasa.com, Technical Consultant (Kenya) Spring 2017
Built data analytics pipeline for the CEO and CTO to track most popular authors and content for the news site Hivisasa.com in Kenya

Teneo, Data Scientist Fall 2015—Jan 2017
Built predictive models and analytics dashboards to drive client insights in strategic communications and shareholder activism.

Microsoft Program Manager in Bing Local Fall 2013—Fall 2015
Shipped local search features for global markets that boosted revenue by \$5M annually, including the local recommendation engine.

Research talks

Merck Data Science All Hands Invited Talk October 2021
On algorithmic fairness and decision support systems

Johns Hopkins University Causal Inference Working Group Invited Talk May 2021
On counterfactual learning, evaluation, and fairness assessments

PlaceKey COVID-19 Data Consortium Invited Talk April 2021
On auditing mobility data for disparate coverage by age and race

University of Pennsylvania Epidemiology Invited Talk February 2021
On causal inference and fairness in algorithmic-assisted decision support systems

University of Chicago Crime Lab Invited Talk September 2020

On algorithmic-assisted decision support systems in public services

Civic Engagement

Court Appointed Special Advocate, Family Law CASA

Winter 2014—Fall 2015

Investigated the child's interest in family law cases and made recommendations to the court on the child's behalf

Committee on Discipline, Princeton University

Fall 2012—Spring 2013

Adjudicated cases of alleged academic and behavioral violations

Engineers Without Borders

Fall 2010—Fall 2012

Skills: R; Python including NumPy and Pandas; proficient in French