

Alexandra Chouldechova
Heinz College, Carnegie Mellon University

(412) 268 4414
achould@cmu.edu
<http://www.andrew.cmu.edu/~achoulde/>

- EMPLOYMENT **Carnegie Mellon University, H. John Heinz III College**
Estella Loomis McCandless Assistant Professor of
Statistics & Public Policy January 2020 —
Assistant Professor, Statistics and Public Policy October 2014 —
- Partnership on AI**
Research Fellow January 2020 — July 2020
- Microsoft Research, New England**
Visiting Researcher June 2018 — September 2018
- EDUCATION Stanford University, Stanford, CA
Ph.D., Statistics, 2009 — September 2014
- University of Toronto, Toronto, Canada
Hon. B.Sc. with High Distinction, Mathematical Statistics, 2009
- AWARDS **Research**
Best Thematic Paper Award, NAACL 2019
Honorable Mention, CHI 2019
Best Technical & Interdisciplinary Paper, FAT* 2018
NSERC Postgraduate Scholarship D, 2010-2013
NSERC Postgraduate Scholarship M, 2009-2010
- Teaching**
Stanford University Centennial Teaching Assistant Award, 2013
Stanford Statistics Department Teaching Assistant Award, 2010
- PUBLICATIONS Rhema Vaithianathan, Emily Putnam-Hornstein, Alexandra Chouldechova, Diana Benavides-Prado, Rachel Berger. (2020) Hospital Injury Encounters of Children Identified by a Predictive Risk Model for Screening Child Maltreatment Referrals: Evidence from the Allegheny Family Screening Tool. *To appear in JAMA Pediatrics*.
- Riccardo Fogliato, Max G'Sell and Alexandra Chouldechova. Fairness evaluation in the presence of biased noisy labels. *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS 2020)*

Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. A case for humans-in-the-loop: decisions in the presence of erroneous algorithmic scores. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (CHI 2020)

Alexandra Chouldechova. (2020) Transparency and Simplicity in Criminal Risk Assessment. Response to: “The Age of Secrecy and Unfairness in Recidivism Prediction” by Rudin, Wang and Coker. *Harvard Data Science Review*

Amanda Coston, Alan Mishler, Edward Kennedy and Alexandra Chouldechova. Counterfactual risk assessment, evaluation, and fairness. *In Proceedings of the ACM Conference on Fairness, Accountability, and Transparency* (FAT* 2020)

Alexandra Chouldechova and Aaron Roth. (2020) A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*

Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. (2019) Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-Making in Child Welfare Services. *In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (CHI 2019)

Best Paper Honorable mention

A. Romanov, M. De-Arteaga, H. Wallach, J. Chayes, C. Borgs, A. Chouldechova, S. Geyik, K. Kenthapadi, A. Rumshisky, A. Kalai, (2019) Whats in a name? Reducing bias in bios without access to protected attributes, *In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics* (NAACL 2019).

Best Thematic Paper Award

M. De-Arteaga, A. Romanov, H. Wallach, J. Chayes, C. Borgs, A. Chouldechova, S. Geyik, K. Kenthapadi, A. Kalai (2019) Bias in Bios: A Case Study of Semantic Representation Bias in a High-Stakes Setting, *In Proceedings of the ACM Conference on Fairness, Accountability, and Transparency* (FAT* 2019)

Zachary Lipton, Alexandra Chouldechova, and Julian McAuley. (2018) Does mitigating ML’s impact disparity require treatment disparity? *In Proceedings of the Thirty-second Annual Conference on Neural Information Processing Systems* (NeurIPS 2018)

De-Arteaga, M., Dubrawski, A., Chouldechova, A. (2018) Learning under selective labels in the presence of expert consistency. *5th Workshop on Fairness, Accountability, and Transparency in Machine Learning* (FAT/ML 2018)

Burckhardt, P., Chouldechova, A., Nugent, R., (2018) TeachIT: Turning the

classroom into a research laboratory via interactive e-learning tools. *10th International Conference on Teaching Statistics (ICOTS10)*

Chouldechova, A., Putnam-Hornstein, E., Benavides Prado, D., Fialko, O., Vaithianathan, R. (2018) A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions (with Diana Benavides Prado, Oleksandr Fialko and Rhema Vaithianathan) *Conference on Fairness, Accountability, and Transparency (FAT* 2018)*

Best Technical & Interdisciplinary Paper Award

Chouldechova, A. (2017) Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*.

Chouldechova, A. and G'Sell, M., (2017) Fairer and more accurate, but for whom? *4th Workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT/ML 2017)*

Diwakar, I.B., Chouldechova, A., Clements, M.A., Padman, R., (2017) Predictive Analytics for Pediatric Diabetic Keto-Acidosis (DKA) using Features Extracted from Asynchronous Multivariate Data Streams. *iHealth 2017 Clinical Informatics Conference*

Diwakar, I.B., Chouldechova, A., Clements, M.A., Padman, R., (2016) On Extracting Features from Asynchronous Multivariate Data Streams. *4th International Conference on Business Analytics and Intelligence (Best Paper Award)*

Chouldechova, A., (2016) Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *3rd Workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT/ML 2016)*

Kanji, H.D., Chouldechova, A., Harvey, C., O'dea E., Faulkner G., & Peek, G. (2016). Safety and Outcomes of Mobile ECMO Using a Bicaval Dual-Stage Venous Catheter. *ASAIO Journal*.

G'Sell, M. G., Wager, S., Chouldechova, A., & Tibshirani, R. (2015). False Discovery Rate Control for Sequential Selection Procedures, with Application to the Lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.

Chouldechova, A., & Mease, D. (2013). Differences in search engine evaluations between query owners and non-owners. *In Proceedings of the sixth ACM international conference on Web search and data mining (WSDM)*, 103-112.

Liu, J., Narsinh, K. H., Lan, F., Wang, L., Nguyen, P. K., Hu, S., Lee, A., Han, L., Gong, Y., Huang, M., Nag, D., Rosenberg, J., Chouldechova, A., Robbins,

R. C., Wu, J. C. (2012). Early Stem Cell Engraftment Predicts Late Cardiac Functional Recovery Clinical Perspective Preclinical Insights From Molecular Imaging. *Circulation: Cardiovascular Imaging*, 5(4), 481-490.

PRESENTATIONS **2020**

Algorithm-assisted decision-making in child welfare *Microsoft Research NYC*

Social license, legitimacy, and organizational justice *Berkeley Center for Law and Technology Symposium*

Human-centered predictive modeling in child welfare *Center for Social Data Analytics, Penn State*

2019

Fairness evaluation in the presence of biased noisy labels. *CMStatistics 2019*

Panel on Fairness and Ethics in Data Science *ACM/IMS Data Science Summit*

Toward accountability in public services *Simons Institute Workshop*

Panel on algorithms and the law *Columbia Law Symposium*

Evaluating criminal risk assessment tools *University of Pittsburgh Statistics Seminar*

2018

Predictive bias and disparate impact in criminal justice decision making *Microsoft Research, Cambridge*

Fairness and bias in predictive modelling, *Society of Epidemiology Research Colloquium*

Risk assessment in criminal sentencing: Target variable bias and disparate impact *University of Washington CSSS Seminar*

Fairness in classification: A look at predictive bias in recidivism prediction instruments *University of Washington Biostatistics Seminar*

Algorithm-assisted decision making by public service agencies *AI and the Law Conference, Seton Hall Law School, NJ*

Fairness in classification: A look at predictive bias in recidivism prediction instruments *Harvard Statistics Colloquium*

Predictive bias and disparate impact in recidivism risk prediction *UPenn Criminology Colloquium, Philadelphia, PA*

Predictive bias and disparate impact in criminal justice decision making *Microsoft Research, New York City*

“Science of Smart Cities” Panel *Metro21: Smart Cities Institute Launch, Pittsburgh, PA*

Data-driven decision making in criminal justice and human services *Workshop on Accountable Decision Systems, Cornell Tech*

“Algorithmic bias”: Practical and technical challenges. *BIRS Workshop on the Interface of Machine Learning and Statistical Inference, Banff, Canada*

2017

The Devil in the Details: Race and the use of algorithms to guide decision-making. *Beyond the Bench 2017, San Diego, CA*

“First, Do No Harm”: Applying predictive analytics to human services. Panel presentation at *2017 MetroLab Summit*

Predictive analytics for child welfare: Fundamental principles. Panel presentation at *2017 MetroLab Summit*

Fair prediction with disparate impact. *Wharton Statistics Seminar*

Fairness-aware predictive analytics in child protective services. *Presentation to visitors from Denmark’s TrygFonden’s Centre for Child Research*

Fairness in data-driven decision making. *Allegheny County Lunch & Learn Series*

A fair predictor of . . . who gets caught? *ACLU and NYU convening on: What Does Fairness Look Like? A Conversation on Race, Risk Assessment Tools, and Pretrial Justice*

Fairness-aware predictive analytics in child protective services. Panel presentation at *Bloomberg Data for Good Exchange 2017*

Fairness in data-driven decision making. *National Academies expert planning meeting on Predictive Analytics for Human Services and Education*

Predictive Analytics In Child Abuse and Neglect: Real world use-cases for fairness-aware, ethical and effective use of analytics in child welfare.

Panel presentation at *2017 Data for Policy Conference*

Fair prediction with disparate impact. *2017 Joint Statistical Meetings*

Fairer and more accurate, but for whom? *Max Planck Institute for Software Systems*

Fair prediction with disparate impact. *2017 International Chinese Statistical Association Applied Statistics Symposium*

Data science in society: Bias, discrimination and accountability *Barrett Lectures, University of Tennessee - Knoxville*

Bias and discrimination in data-driven decision making. *University of Toronto Department of Statistical Sciences 40th Anniversary Research Day*

Algorithms and explanations: Model comparison. *NYU Information Law Institute Workshop*

2016 and earlier

Fair prediction with disparate impact. *Cornell Statistics Department Seminar; Yale Statistics Department Seminar; Heinz College Seminar*

Small effect sizes with large disparate impact: Bias in recidivism prediction. *2016 International Indian Statistical Association Conference on Statistics*

False discovery rate control for spatial data. *ENAR 2015, JSM-NRC 2015, JSM 2015*

GAMSEL: A penalized likelihood approach to model selection for generalized additive models. *JSM 2014, SIAM Conference on Optimization 2014*

Clusterwise false discovery rate control in spatial data. *JSM 2013*

Error rate control in non-standard settings. *Asilomar statistics retreat 2012*

SOFTWARE **gamsel:** R package for fitting sparse generalized additive models for high-dimensional data

GRANTS PwC Risk and Regulatory Services Innovation Center (June 2020 - May 2021), \$302,790

Auditing for Quality and Fairness when Predictions become Policies

Joint with Zack Lipton (Tepper/MLD) and Matt Fredrikson (CSD)

MacArthur Foundation (March 2020 - March 2022), \$1,700,000

Pretrial Risk Management in the Safety and Justice Challenge

Joint with John Monahan (UVA), Kristian Lum (HRDAG), Logan Koepke (Upturn), David Robinson (Upturn/Cornell) and Sarah Desmarais (NCSU)

National Science Foundation in collaboration with Amazon (January 2020 – December 2022), \$1,037,000

FAI: Advancing Fairness in AI with Human-Algorithm Collaborations

Joint with Steven Wu (UMN), Min Kyung Lee (UT-Austin) and Haiyi Zhu (CMU)

Block Center (May 2019 – July 2020), \$75,000

Counterfactual Risk Assessment for Improved Decision Support in Child Welfare Services

Joint with Edward Kennedy (CMU)

Allegheny County Department of Human Services (May 2018 –), \$100,000

Predictive bias in child welfare risk assessment

Metro21 Institute (May 2018 - May 2019), \$51,024

Developing and Deploying Methodologies for Improving the Accuracy, Fairness, and Trustworthiness of Risk Assessment Models in Child Welfare

Berkman Development Fund Grant (PI, January 2017 – December 2017), \$4,500

FairCompare: An R package and dashboard for fairness assessment of data-driven prediction models

ProSEED/Simon Initiative (PI, May 2015 – May 2018), \$14,084

Improving Statistics Education with Interactive Computer-Enabled Lab Activities

SERVICE

Referee: Annals of Statistics; Journal of the American Statistical Association; Bernoulli; Electronic Journal of Statistics; Criminology & Public Policy; Journal of Statistical Planning and Inference; Journal of Quantitative Criminology; Statistical Science; Biometrika; Science; Proceedings of the National Academy of Sciences

Program committee: Track chair FAT* 2018; Senior Program Committee member, IJCAI 2019, 2020; Area Chair NeurIPS 2019, 2020; Area Chair ICML 2020

Organization:

Workshop co-organizer: CCC Fair Representations and Fair Interactive Learning Workshop (with Aaron Roth)

Program co-Chair: 2019 ACM FAT* Conference (with Fernando Diaz)

TEACHING

Heinz College, CMU

95-791: Data Mining I

94-942: Programming in R for Analytics

940-739: PPM Systems project (empirical reconstruction of the prior record score)

95-721: MISM Capstone project (crime recidivism prediction)