Michael Oberst

moberst@jhu.edu, moberst.com
Last updated: November 27, 2024

Summary

My research spans causal inference, machine learning, and statistical methods, with the goal of enabling reliable decision-making and prediction in high-risk domains like healthcare.

Appointments

Johns Hopkins University, Baltimore, MD — 2024–Onwards
Assistant Professor, Computer Science

Carnegie Mellon University, Pittsburgh, PA — 2023–2024
Postdoctoral Associate, Machine Learning Department

Education

Massachusetts Institute of Technology, Cambridge, MA — 2017–2023
PhD. Electrical Engineering & Computer Science
Dissertation: "Towards Rigorously Tested & Reliable Machine Learning for Health"
Advisor: David Sontag

Massachusetts Institute of Technology, Cambridge, MA — 2017–2019
M.S. Electrical Engineering & Computer Science
Thesis: "Counterfactual Policy Introspection with Structural Causal Models"
Advisor: David Sontag

Harvard University, Cambridge, MA — 2008–2012
B.A. Statistics (*Summa Cum Laude*)
Advisor: Edoardo Airoldi

Publications[1]

Conferences

EMNLP 2024 — **Medical Adaptation of Large Language and Vision-Language Models: Are We Making Progress?**. D. Jeong, S. Garg, Z. Lipton, M. Oberst. Conference on Empirical Methods in Natural Language Processing, 2024

AISTATS 2024 — **Auditing Fairness under Unobserved Confounding**. E. Byun, D. Sam, M. Oberst, Z. Lipton, B. Wilder. International Conference on Artificial Intelligence and Statistics, 2024

[1] ∗ denotes co-first author, † denotes a student (BS/MEng) that I supervised.

AISTATS 2024 **Benchmarking Observational Studies with Experimental Data under Right-Censoring**. I. Demirel, E. De Brouwer, Z. Hussain, M. Oberst, A. Philippakis, D. Sontag. International Conference on Artificial Intelligence and Statistics, 2024

AISTATS 2023 **Falsification of Internal and External Validity in Observational Studies via Conditional Moment Restrictions**. Z. Hussain, MC. Shih, M. Oberst, I. Demirel, D. Sontag. International Conference on Artificial Intelligence and Statistics, 2023

NeurIPS 2022 **Evaluating Robustness to Dataset Shift via Parametric Robustness Sets**. N. Thams*, M. Oberst*, D. Sontag. Neural Information Processing Systems, 2022

NeurIPS 2022 **Falsification before Extrapolation in Causal Effect Estimation**. Z. Hussain*, M. Oberst*, MC. Shih*, D. Sontag. Neural Information Processing Systems, 2022

NeurIPS 2021 **Finding Regions of Heterogeneity in Decision-Making via Expected Conditional Covariance**. J. Lim*†, C. Ji*, M. Oberst*, S. Blecker, L. Horwitz, D. Sontag. Neural Information Processing Systems, 2021

ICML 2021 **Regularizing towards Causal Invariance: Linear Models with Proxies**. M. Oberst, N. Thams, J. Peters, D. Sontag. International Conference on Machine Learning, 2021

AMIA 2021 **Trajectory Inspection: A Method for Iterative Clinician-Driven Design of Reinforcement Learning Studies**. C. Ji*, M. Oberst*, S. Kanjilal, D. Sontag. American Medical Informatics Association Annual Symposium, 2021

AISTATS 2020 **Characterization of Overlap in Observational Studies**. M. Oberst*, FD. Johansson*, D. Wei*, T. Gao, G. Brat, D. Sontag, KR. Varshney. International Conference on Artificial Intelligence and Statistics, 2020

KDD 2020 **Treatment Policy Learning in Multiobjective Settings with Fully Observed Outcomes**. S. Boominathan†, M. Oberst, H. Zhou, S. Kanjilal, D. Sontag. ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2020

ICML 2019 **Counterfactual Off-Policy Evaluation with Gumbel-Max Structural Causal Models**. M. Oberst, D. Sontag. International Conference on Machine Learning, 2019

Journal Articles

STM 2020 **A Decision Algorithm to Promote Outpatient Antimicrobial Stewardship for Uncomplicated Urinary Tract Infection**. S. Kanjilal, M. Oberst, S. Boominathan, H. Zhou, DC. Hooper, D. Sontag. Science Translational Medicine, 2020

SR 2020 **Predicting Human Health from Biofluid-Based Metabolomics using Machine Learning**. ED. Evans, C. Duvallet, ND. Chu, M. Oberst, MA. Murphy, I. Rockafellow, D. Sontag, EJ. Alm. Scientific Reports, 2020

Preprint / Working Paper

ACIC 2022 **Bias-robust Integration of Observational and Experimental Estimators**. M. Oberst, A. D'Amour, M. Chen, Y. Wang, D. Sontag, S. Yadlowsky. Presented at the American Causal Inference Conference, 2022.

Honors & Awards

Teaching	Fredrik C. Hennie III Award for Teaching Excellence	2021
Reviewing	Top Reviewer (Top 10%), Neural Information Processing Systems (NeurIPS)	2022
	Top Reviewer (Top 10%), International Conference on Machine Learning (ICML)	2022
	Top Reviewer (Top 5%), Uncertainty in Artificial Intelligence (UAI)	2021
Fellowships	Honorable Mention, NSF Graduate Research Fellowship Program (NSF-GRFP)	2018
	Analog Devices Graduate Fellowship	2017
Academic	John Harvard Scholar (Top 5% of students)	2012
	Phi Beta Kappa, Senior 48	2011

Teaching & Mentorship

Teaching	Teaching EN.601.788 (Machine Learning for Healthcare) at JHU	2024
	Head Teaching Assistant for 6.867 (Machine Learning) at MIT	2021
	Fredrik C. Hennie III award for teaching excellence	
	Overall Teaching Evaluation Rating: 6.9/7.0	
	Teaching Fellow for CS50 (Intro to Computer Science) at Harvard	2010
Mentorship	CMU Undergraduate — Andrew Wang	2024
	CMU Master's Thesis — Pratheek De Souza Rebello	2024
	MIT Master's Thesis — Justin Lim	2020
	First-author publication in Neural Information Processing Systems (NeurIPS)	
	MIT Master's Thesis — Sooraj Boominathan	2019
	First-author publication in Knowledge Discovery and Data Mining (KDD)	
	MIT Undergraduate Research Opportunities Program — Shreyas Balaji	2019
	MIT Undergraduate Research Opportunities Program — Justin Lim	2019
	MIT Undergraduate Research Opportunities Program — Elizabeth Han	2019
	MIT Master's Thesis — Helen Zhou	2018
	MIT Undergraduate Research Opportunities Program — Sooraj Boominathan	2018

Invited Talks & Presentations

SIAM Conference on Mathematics of Data Science, Atlanta, GA 2024
Invited Talk at Mini-Symposium: Mathematics of Trustworthy Machine Learning
Title: Auditing Fairness under Unobserved Confounding

INFORMS, Seattle, WA 2024
Invited Talk at Invited Session: Machine Learning Aided Causal Inference
Title: Auditing Fairness under Unobserved Confounding

LMU Munich, Munich, Germany 2024
AI Keynote Series, Institute of AI in Management
Title: Auditing Fairness under Unobserved Confounding

Artificial Intelligence in Medicine (AIME) Conference, Salt Lake City, UT 2024
Invited Talk at Workshop: AI for Reliable and Equitable RWE Generation in Medicine
Title: Auditing Fairness under Unobserved Confounding

Abridge AI, Pittsburgh, PA 2024

Abridge AI Seminar
Title: Towards Rigorously Tested and Reliable Machine Learning for Healthcare

University of California, Berkeley, Berkeley, CA 2024
Computational Precision Health Doctoral Seminar
Title: Towards Rigorously Tested and Reliable Machine Learning for Healthcare

Johns Hopkins University, Baltimore, MD 2023
Institute for Assured Autonomy Seminar
Title: Towards Rigorously Tested and Reliable Machine Learning for Healthcare

Carnegie Mellon University, Pittsburgh, PA 2022
Machine Learning / Duolingo Seminar
Title: What is the role of causality in reliable prediction?

Chalmers University, Gothenburg, Sweden 2022
Johansson Lab
Title: Evaluating Robustness to Dataset Shift via Parametric Robustness Sets

University of California, Berkeley, Berkeley, CA 2022
American Causal Inference Conference (ACIC)
Title: Bias-robust Integration of Observational and Experimental Estimators

Stanford University, Palo Alto, CA 2022
Online Causal Inference Seminar (OCIS)
Title: Regularizing towards Causal Invariance: Linear Models with Proxies

WHOOP (Wearables Company), Boston MA 2020
Presentation to Data Science team
Title: Learning Treatment Policies from Observational Data

Broad Institute of MIT and Harvard, Cambridge MA 2020
Models, Inference, and Algorithms Seminar
Title: Primer: Learning Treatment Policies from Observational Data

Service

Conferences	Reviewer, Neural Information Processing Systems (NeurIPS)	2024
	Reviewer, International Conference on Machine Learning (ICML)	2024
	Reviewer, Artificial Intelligence & Statistics (AISTATS)	2024
	Reviewer, Causal Learning & Reasoning (CLeaR)	2024
	Reviewer, Neural Information Processing Systems (NeurIPS)	2023
	Reviewer, Neural Information Processing Systems (NeurIPS)	2022
	Top Reviewer Award (Top 10%)	
	Reviewer, International Conference on Machine Learning (ICML)	2022
	Top Reviewer Award (Top 10%)	
	Reviewer, Machine Learning for Health Symposium (ML4H)	2022
	Reviewer, Neural Information Processing Systems (NeurIPS)	2021
	Reviewer, Uncertainty in Artificial Intelligence (UAI)	2021
	Top Reviewer Award (Top 5%)	
	Reviewer, Artificial Intelligence & Statistics (AISTATS)	2019
Journals	Reviewer, Journal of Causal Inference	2024
	Reviewer, European Journal of Epidemiology	2024

	Reviewer, Journal of Machine Learning Research (JMLR)	2024
	Reviewer, Journal of the Royal Statistical Society—Series A (JRSSA)	2023
	Reviewer, Journal of Causal Inference	2022
	Reviewer, Statistics & Computing	2021
	Reviewer, Bayesian Analysis	2021
Workshops	**AC**, Robustness of few/zero-Shot Learning in Foundation Models (NeurIPS)	2023
	Reviewer, Robustness of few/zero-Shot Learning in Foundation Models (NeurIPS)	2023
	Reviewer, Counterfactuals in Minds and Machines (ICML)	2023
	Reviewer, Principles of Distribution Shift (ICML)	2022
	Reviewer, Distribution Shifts (NeurIPS)	2022
	Reviewer, Casual Inference in Sequential Decision-Making (NeurIPS)	2021
	Reviewer, Distribution Shifts (NeurIPS)	2021
	Mentor, Machine Learning for Health (ML4H) Reviewer Mentorship Program	2021
	Organizer, Machine Learning for Health (ML4H)	2019
Admissions	Mentor, MIT EECS Graduate Application Assistance Program	2020
	Application Reviewer, MIT EECS PhD Admissions	2019

Work Experience

Abridge AI, Pittsburgh, PA 2024–Present
Visiting Scientist
Develop methods to improve the robustness and reliability of AI-based ambient clinical documentation in healthcare. Part-time (20%) since July 2024.

Google Brain, Cambridge, MA 2021
PhD Research Intern
Worked on learning reliable short-term surrogates for long-term outcomes in recommender systems, using recommendations as instrumental variables. Worked with Alexander D'Amour, Steve Yadlowsky, Minmin Chen, and Yuyan Wang.

Clarify Health Solutions, San Francisco, CA, USA 2016–2017
Manager, Data Science
Built and led the data science / data engineering team for a healthcare tech startup, creating software to help hospitals deliver on value-based care.

McKinsey & Company, Nairobi, Kenya 2014–2015
Associate, Africa Delivery Hub
Managed teams across Southern and Eastern Africa, helped launch Nairobi office and dedicated group focused on public sector work.

McKinsey & Company, New York City, NY 2012–2014
Business Analyst
Worked on projects in public health, education, renewable energy, and financial services.