

Miruna Oprescu

📍 New York, NY ✉ amo78@cornell.edu 🌐 mirunaoprescu.com in [antonia-miruna-oprescu](https://www.linkedin.com/in/antonia-miruna-oprescu)

Summary

I am a researcher in causal inference and machine learning. My work develops methods for treatment effect estimation, policy evaluation, and quasi-experimental data in real-world settings, with an emphasis on robustness to unobserved confounding, limited data, and structured domains such as spatiotemporal and networked systems. Across academia and industry, I have developed both new methodology and practical tooling for heterogeneous treatment effect estimation and policy evaluation, including foundational contributions to the open-source [EconML](#) library. I am motivated by building trustworthy and effective methods for complex real-world systems and socially consequential applications.

Areas: Causal inference, machine learning, heterogeneous treatment effect estimation, policy evaluation, quasi-experimental methods, robust causal inference, structured data, AI for science

Education

Cornell University (Cornell Tech)

Fall 2021 – Present

Ph.D. Candidate in Computer Science. Advisor: Nathan Kallus

M.S. in Computer Science (2024, awarded en route to Ph.D.)

Department of Energy Computational Science Graduate Fellow

Harvard University

2011 – 2015

Joint A.B. degree in Physics and Mathematics. Minor in Computer Science

Cum laude in field with High Honors in Physics and Mathematics

Experience

Machine Learning Research Intern

Summer 2025

Netflix, New York City, NY

- Developed and estimated a structural causal model of subscription choice under quasi-experimental plan migration, leveraging instrumental variation and purchasing power differences to identify plan value, price sensitivity, and switching costs. Mentors: David Hubbard, Yonatan Gur.

Research Intern

Summer 2024

Brookhaven National Laboratory, Brookhaven, NY

- Developed GST-UNet [2], a neural framework for spatiotemporal causal inference under time-varying confounding, integrating UNet architectures with iterative G-computation to estimate intervention effects in environmental and health applications. Mentor: Shinjae Yoo.

Machine Learning Research Intern

Summer 2022

Netflix, Los Gatos, CA

- Built a causal machine learning model to estimate the effect of watching novel content on user engagement, developing methods that combine weak instrumental variables with observational data to identify effects under low compliance [3]. Mentors: Sudeep Das, Aish Fenton.

Senior Data and Applied Scientist

2017 – 2021

Microsoft Research, Cambridge, MA

- Conducted research in causal machine learning, developing methods for heterogeneous treatment effect estimation and policy evaluation, with results published at top venues [8, 9, 10, 11].

- Founded and maintained to the open-source [EconML](#) library, implementing production-scale causal inference algorithms now widely used in industry and academia.
- Collaborated on projects in subseasonal weather forecasting and cancer modeling, applying ML techniques to real-world scientific data.

Software Development Engineer

2015 – 2017

Microsoft, Cambridge, MA

- Developed scalable, distributed machine learning algorithms as a core contributor to [MMLSpark](#), Microsoft’s open-source ML library for Apache Spark.

Selected Publications

† equal contribution, ‡ alphabetical authors. Full publication list available on [Google Scholar](#).

Preprints

- [1] Maresa Schröder, **Miruna Oprescu**, Stefan Feuerriegel, and Nathan Kallus. Causal inference on networks under misspecified exposure mappings: A partial identification framework. *arXiv preprint arXiv:2602.03459*, 2026.
- [2] Ayush Khot[†], **Miruna Oprescu**[†], Maresa Schröder, Ai Kagawa, and Xihai Luo. Spatial deconfounder: Interference-aware deconfounding for spatial causal inference. *arXiv preprint arXiv:2510.08762*, 2025.

Conference and Journal Publications

- [1] **Miruna Oprescu**, Brian M Cho, and Nathan Kallus. Efficient adaptive experimentation with non-compliance. *Advances in Neural Information Processing Systems*, 2025. To appear.
- [2] **Miruna Oprescu**, David K Park, Xihai Luo, Shinjae Yoo, and Nathan Kallus. Gst-unet: A neural framework for spatiotemporal causal inference with time-varying confounding. *Advances in Neural Information Processing Systems*, 2025. To appear.
- [3] **Miruna Oprescu** and Nathan Kallus. Estimating heterogeneous treatment effects by combining weak instruments and observational data. *Advances in Neural Information Processing Systems*, 37:118777–118806, 2025.
- [4] Andrew Bennett, Nathan Kallus, **Miruna Oprescu**[†], Wen Sun, and Kaiwen Wang[†]. Efficient and sharp off-policy evaluation in robust markov decision processes. *Advances in Neural Information Processing Systems*, 37:112962–113000, 2025.
- [5] Andrew Bennett[†], Nathan Kallus[†], and **Miruna Oprescu**[†]. Low-rank mdps with continuous action spaces. In *International Conference on Artificial Intelligence and Statistics*, pages 4069–4077. PMLR, 2024.
- [6] **Miruna Oprescu**, Jacob Dorn, Marah Ghoummaid, Andrew Jesson, Nathan Kallus, and Uri Shalit. B-learner: Quasi-oracle bounds on heterogeneous causal effects under hidden confounding. In *International Conference on Machine Learning*, pages 26599–26618. PMLR, 2023.
- [7] Nathan Kallus[†] and **Miruna Oprescu**[†]. Robust and agnostic learning of conditional distributional treatment effects. In *International Conference on Artificial Intelligence and Statistics*, pages 6037–6060. PMLR, 2023.

- [8] Keith Battocchi[‡], Eleanor Dillon[‡], Maggie Hei[‡], Greg Lewis[‡], **Miruna Oprescu[‡]**, and Vasilis Syrgkanis[‡]. Estimating the long-term effects of novel treatments. *Advances in Neural Information Processing Systems*, 34:2925–2935, 2021.
- [9] **Miruna Oprescu[†]**, Vasilis Syrgkanis[†], and Zhiwei Steven Wu[†]. Orthogonal random forest for causal inference. In *International Conference on Machine Learning*, pages 4932–4941. PMLR, 2019.
- [10] Vasilis Syrgkanis, Victor Lei, **Miruna Oprescu**, Maggie Hei, Keith Battocchi, and Greg Lewis. Machine learning estimation of heterogeneous treatment effects with instruments. *Advances in Neural Information Processing Systems*, 32:15193–15202, 2019. **Spotlight presentation.**
- [11] **Miruna Oprescu[†]**, Vasilis Syrgkanis[†], Keith Battocchi[†], Maggie Hei[†], and Greg Lewis[†]. EconML: A Machine Learning Library for Estimating Heterogeneous Treatment Effects. In *CausalML Workshop, NeurIPS*, 2019. **Spotlight presentation.**

Talks

- *Causal Inference for Spatiotemporal Interventions*, SIAM CSE, 2025.
- *Reliable Treatment Effect Estimation Using Weak Instruments and Observational Data*, Workshop in Operations Research and Data Science (WORDS), Duke University, 2024.
- *Uncertainty Quantification in Causal Inference: Sharp and Efficient Bounds on Heterogeneous Causal Effects Under Hidden Confounding*, Brookhaven National Laboratory, 2023.
- *Causal Inference and Machine Learning in Practice with EconML and CausalML*, SIGKDD, 2021.
- *EconML: A Machine Learning Library for Estimating Heterogeneous Treatment Effects*, ODSC East, invited talk, 2019.

Honors & Awards

- Department of Energy Computational Science Graduate Fellowship 2022–2026
- DOE Communicate Your Science & Engineering (CYSE) Award 2025
- Meta PhD Research Fellowship Finalist 2022
- *Cum laude*, Harvard University; High Honors in Physics and Mathematics 2015
- Derek C. Bok Award for Distinction in Teaching (Data Science), Harvard 2014

Professional Service

- Director, Summer Science Program (SSP), 2023–Present
- Reviewer: NeurIPS (Top Reviewer Award, 2025), ICML, AISTATS 2021–present

Teaching

- Teaching Assistant, Cornell University
 - Learning, Inference, and Decision Making from Data Spring 2022
 - Applied Machine Learning Fall 2021
- Teaching Fellow, Harvard University: Physics, Data Science, and Mathematics courses; recipient of the Derek C. Bok Award for Distinction in Teaching 2012–2014