

Elliot Creager

University of Waterloo
Department of Electrical and Computer Engineering
200 University Ave W
Waterloo, Ontario, N2L 3G3
Canada

email: creager@uwaterloo.ca
url: <https://ecreager.github.io/>
Github: [ecreager](#)
Google Scholar: [boebIUcAAAAJ](#)

Current position

2023 - *Assistant Professor*, University of Waterloo

Previous experience

2020 - 2021 *Graduate Fellow*, Schwartz Reisman Inst. for Technology and Society, Toronto, Ontario
2019 - 2020 *Student Researcher*, Google Brain, Toronto, Ontario
2019 *Research Intern*, Google Brain, Toronto, Ontario
2015 - 2017 *Research Scientist*, Analog Devices, Inc., Cambridge, Massachusetts
2014 *Research Intern*, Analog Devices, Inc.
2013 *Research Intern*, Analog Devices, Inc.

Education

2023 Ph.D. in Computer Science, University of Toronto
2015 M.A. in Music Technology, McGill University
2013 Sc.B. in Electrical Engineering (Honors) and A.B. in Music, Brown University

Publications

Conferences

- 2024b P. A. Alamdari, T. Q. Klassen, **E. Creager**, and S. McIlraith, “Remembering to Be Fair: On Non-Markovian Fairness in Sequential Decision Making”, *ICML 2024*
- 2024a B. Eyre, **E. Creager**, D. Madras, V. Pappayan, and R. Zemel, “Out of the Ordinary: Spectrally Adapting Regression for Covariate Shift”, *ICML 2024*
- 2023a A. Mani, I. P. Chandratreya, **E. Creager**, C. Vondrick, and R. Zemel, “SurfsUp: Learning Fluid Simulation for Novel Surfaces”, *ICCV 2023*
- 2022a S. Pitís, **E. Creager**, A. Mandlekar, and A. Garg, “MoCoDA: Model-based Counterfactual Data Augmentation”, *NeurIPS 2022*
- 2021b F. Trauble, **E. Creager**, N. Kilbertus, F. Locatello, A. Dittadi, A. Goyal, B. Schölkopf, and S. Bauer, “On Disentangled Representations Learned from Correlated Data”, *ICML 2021* (Oral)
- 2021a **E. Creager**, J.-H. Jacobsen, and R. Zemel, “Environment Inference for Invariant Learning”, *ICML 2021*
- 2020c S. Pitís, **E. Creager**, and A. Garg, “Counterfactual Data Augmentation for Locally Factored Dynamics”, *NeurIPS 2020* (also “outstanding paper” at *ICML 2020 Object-oriented Learning Workshop*)
- 2020b M. Mladenov, **E. Creager**, O. Ben-Porat, K. Swersky, R. Zemel, and C. Boutilier, “Optimizing Long-term Social Welfare in Recommender Systems: A Constrained Matching Approach”, *ICML 2020*
- 2020a **E. Creager**, D. Madras, T. Pitassi, and R. Zemel, “Causal Modeling for Fairness in Dynamical Systems”, *ICML 2020*
- 2019a D. Madras, **E. Creager**, T. Pitassi, and R. Zemel, “Fairness Through Causal Awareness: Learning Latent-Variable Models for Biased Data”, *ACM FAT* 2019*
- 2019c **E. Creager**, D. Madras, J.-H. Jacobsen, M.A. Weis, K. Swersky, T. Pitassi, and R. Zemel, “Flexibly Fair Representation Learning by Disentanglement”, *ICML 2019*
- 2019b C.-H. Chang, **E. Creager**, A. Goldenberg, and D. Duvenaud, “Explaining Image Classifiers by Counterfactual Generation”, *ICLR 2019*
- 2018a D. Madras*, **E. Creager***, T. Pitassi, and R. Zemel, “Learning Adversarially Fair and Transferable Representations”, *ICML 2018*
- 2016 **E. Creager**, N.D. Stein, R. Badeau, and P. Depalle, “Nonnegative Tensor Factorization with Frequency Modulation Cues for Blind Audio Source Separation”, *ISMIR 2016*,

Workshops

- 2024d R. Solanki and **E. Creager**, “Promoting User Data Autonomy During the Dissolution of a Monopolistic Firm”, *NeurIPS 2024 Workshop on Regulatable ML*
- 2024c B. Jaipersaud, Z. Zhu, F. Rudzicz, and **E. Creager**, “Show, Don’t Tell: Uncovering Implicit Character Portrayal using LLMs”, *NeurIPS 2024 Workshop on Creativity and Generative AI*
- 2022b B. Eyre, R. Zemel and **E. Creager**, “Towards Environment-Invariant Representation Learning for Robust Task Transfer”, *ICML 2022 Workshop on Spurious Correlations, Invariance*,

* denotes equal contribution

and Stability

- 2021d D. Dickson and **E. Creager**, “**Measuring User Recourse in a Dynamic Recommender System**”, *ICML 2021 Workshop on Algorithmic Recourse*
- 2021c **E. Creager** and R. Zemel, “**Online Algorithmic Recourse by Collective Action**”, *ICML 2021 Workshop on Algorithmic Recourse*
- 2020d R. Adragna, **E. Creager**, D. Madras, and R. Zemel, “**Fairness and Robustness in Invariant Learning: A Case Study in Toxicity Classification**”, *NeurIPS 2020 Workshop on Algorithmic Fairness Through the Lens of Causality* (Oral)
- 2018b W. Grathwohl*, **E. Creager***, S.K.S. Ghasemipour*, R. Zemel, “**Gradient-Based Optimization of Neural Network Architecture**”, *ICLR 2018 Workshop*

Teaching

Course instructor

- 2024-2025 *Algorithm Design and Analysis*, University of Waterloo
- 2022 *Introduction to Artificial Intelligence*, University of Toronto

Conference tutorials

- 2022 *Algorithmic Fairness: at the Intersections*, NeurIPS

Teaching assistant

- 2022 *Introduction to Machine Learning*, University of Toronto
- 2021 *Introduction to Machine Learning*, University of Toronto
- 2021 *Probabilistic Learning and Reasoning*, University of Toronto
- 2019 *AI and Ethics: Mathematical Foundations and Algorithms*, University of Toronto
- 2019 *Fairness and Privacy in Machine Learning*, African Institute for Mathematical Sciences (Rwanda)
- 2018 *Machine Learning and Data Mining*, University of Toronto
- 2018 *Probabilistic Learning and Reasoning*, University of Toronto
- 2017 *Introduction to Artificial Intelligence*, University of Toronto
- 2014 *Digital Audio Signal Processing*, McGill University
- 2013 *Communications Systems*, Brown University
- 2012 *Communications Systems*, Brown University

Invited talks

- 2024 *To Build AI That Works For Everyone, Adapt Models and Coordinate Data*, Waterloo.AI Seminar Series, Waterloo, Canada
- 2023 *Out of the Ordinary: Spectrally Adapted Regression for Covariate Shift*, McGill Equity and Equality Using AI and Learning algorithms (EQUAL) lab meeting, Montreal, Canada
- 2023 *Methods for Counterfactual Data Augmentation in Reinforcement Learning*, Forging a Path:

	Causal Inference for Improved Policy Workshop, Toronto, Canada
2023	<i>Can “Adversaries” Play a Positive Role in Ethical AI?</i> , Vector Machine Learning Security and Privacy Workshop, Toronto, Canada
2023	<i>Counterfactual Reasoning in Reinforcement Learning and Algorithmic Fairness</i> , DEFIRST reading group at Mila (Quebec AI Institute), Montreal, Canada
2023	<i>Society and Ethics Concerns in Machine Learning</i> , Pursue STEM Outreach Program for High Schoolers, Toronto, Canada
2022	<i>Bias in AI: Mitigation Strategies</i> , Vector Institute Bias in AI Program for Industry Sponsors, Toronto, Canada
2021	<i>Fair Representation Learning with Disentanglement</i> , Vector Institute Endless Summer School, Toronto, Canada
2021	<i>Bias in AI: Mitigation Strategies</i> , Vector Institute Bias in AI Program for Industry Sponsors, Toronto, Canada
2021	<i>An Algorithmic Fairness Perspective on Robust Representation Learning</i> (Keynote), Domain Adaptation and Representation Transfer Workshop at MICCAI
2019	<i>Causal Modeling for Fairness in Dynamical Systems</i> , Microsoft Research Guest Lecture Series, Montreal, Canada
2018	<i>Learning Adversarial and Transferable Representations</i> , CIFAR Deep Learning and Reinforcement Learning Summer School, Toronto, Canada

Academic service

2025	<i>Program Committee</i> , ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)
2025	<i>Program Committee (Area Chair, Ethics Reviewer)</i> , ICML
2025	<i>Program Committee</i> , International Conference on Learning Representations (ICLR)
2025	<i>Program Committee</i> , AAAI Conference on Artificial Intelligence (AAAI)
2024	<i>Program Committee</i> , Workshop on Algorithmic Fairness Through the Lens of Metrics and Evaluation (NeurIPS)
2024	<i>Program Committee</i> , Workshop on Generative AI and Creativity (NeurIPS)
2024	<i>Program Committee</i> , Machine Learning For Health Symposium (ML4H)
2024	<i>Program Committee (Area Chair)</i> , NeurIPS
2024	<i>Program Committee</i> , Workshop on the Next Generation of AI Safety (ICML)
2024	<i>Program Committee</i> , Canadian Artificial Intelligence Conference
2024	<i>Program Committee</i> , Workshop on Recommendation Ecosystems: Modeling, Optimization and Incentive Design (AAAI)
2023	<i>Program Committee</i> , Workshop on Robustness of Few-shot and Zero-shot Learning in Foundation Models (NeurIPS)
2023-2024	<i>Program Committee</i> , Workshop on Regulating Machine Learning (NeurIPS)
2023	<i>Program Committee</i> , Workshop on Distribution Shifts: New Frontiers with Foundation Models (NeurIPS)
2023	<i>Program Committee</i> , Workshop on Causal Representation Learning (NeurIPS)
2023-2024	<i>Program Committee</i> , Conference on Health, Inference, and Learning (CHIL)
2022	<i>Program Committee</i> , Workshop on Distribution Shifts (NeurIPS)

2022 *Program Committee*, Workshop on Robustness in Sequence Modeling (NeurIPS)

2022 *Program Committee*, Workshop on A Causal View on Dynamical Systems (NeurIPS)

2022 *Program Committee*, Workshop on Algorithmic Fairness Through the Lens of Causality and Privacy (NeurIPS)

2022 *Program Committee*, Workshop on Continuous-time Methods for ML (ICML)

2022 *Program Committee*, Workshop on Principles of Distribution Shifts (ICML)

2022-2023 *Program Committee*, Workshop on Spurious Correlations, Invariance, and Stability (ICML)

2021 *Program Committee*, Workshop on Distribution Shifts: Connecting Methods and Applications (NeurIPS)

2021 *Program Committee*, Workshop on Algorithmic Fairness Through the Lens of Causality and Robustness (NeurIPS)

2021-2024 *Ethics Reviewer*, NeurIPS

2020-2024 *Program Committee*, ICML

2020 *Co-organizer*, Resistance AI Workshop (NeurIPS)

2020-2023 *Program Committee*, ACM FAccT Conference

2020 *Program Committee*, Workshop on Algorithmic Fairness Through the Lens of Causality and Interpretability (NeurIPS)

2019-2023 *Program Committee*, NeurIPS

Program Committee, Fair Machine Learning for Health Workshop (NeurIPS)