

EDUCATION

Carnegie Mellon University

Ph.D. candidate in Machine Learning. Advised by Pradeep Ravikumar and Andrej Risteski.

Pittsburgh, PA

2018–Present

Carnegie Mellon University

B.S. in Computer Science, additional major in Statistics & Machine Learning

Pittsburgh, PA

2012–2016

– Senior Thesis: “Human-Usable Password Schemas: Beyond Information-Theoretic Security”

Advised by Manuel Blum and Santosh Vempala. Awarded “Exemplary Senior Honors Thesis”.

PUBLICATIONS

Conference Publications

- Y. Chen, E. Rosenfeld, M. Sellke, T. Ma, and A. Risteski. [Iterative Feature Matching: Toward Provable Domain Generalization with Logarithmic Environments](#). NeurIPS 2022 (to appear)
- B. Liu, E. Rosenfeld, P. Ravikumar, and A. Risteski. [Analyzing and Improving the Optimization Landscape of Noise-Contrastive Estimation](#). In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2022.
- I. Apostolopoulou, I. Char, E. Rosenfeld, and A. Dubrawski. [Deep Attentive Variational Inference](#). In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2022.
- E. Rosenfeld, P. Ravikumar, and A. Risteski. [An Online Learning Approach to Interpolation and Extrapolation in Domain Generalization](#). In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.
- E. Rosenfeld, P. Ravikumar, and A. Risteski. [The Risks of Invariant Risk Minimization](#). In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2021.
- E. Rosenfeld, E. Winston, P. Ravikumar, and J. Z. Kolter. [Certified Robustness to Label-Flipping Attacks via Randomized Smoothing](#). In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.
- J. Cohen, E. Rosenfeld, and J. Z. Kolter. [Certified Adversarial Robustness via Randomized Smoothing](#). In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.

Preprints and Manuscripts

- E. Rosenfeld, P. Ravikumar, and A. Risteski. [Domain-Adjusted Regression or: ERM May Already Learn Features Sufficient for Out-of-Distribution Generalization](#). *arXiv preprint arXiv:2202.06856*, 2022.
- E. Rosenfeld, S. Vempala, and M. Blum. [Human-Usable Password Schemas: Beyond Information-Theoretic Security](#). *arXiv preprint arXiv:1906.00029*, 2016.

Workshop Papers

- I. Apostolopoulou, E. Rosenfeld, and A. Dubrawski. [Self-Reflective Variational Autoencoder](#). In *ICLR Workshop on Hardware Aware Efficient Training*, 2021.

INVITED TALKS

Generalizing to New Distributions via Invariance: Challenges and Opportunities

Computational Data Science Seminar, Technion – Israel Institute of Technology. 3 April 2022.

The Risks of Invariant Risk Minimization

Failures of Learning Seminar, Technion – Israel Institute of Technology. 31 March 2022.

International Conference on Learning Representations (Virtual). 3 May 2021.

[Artificial Intelligence Seminar, CMU](#) (Virtual). 6 April 2021.

Causal Discovery & Causality-Inspired Machine Learning Workshop, NeurIPS (Virtual). 11 December 2020.

Certified Robustness to Label-Flipping Attacks via Randomized Smoothing

[International Conference on Machine Learning](#) (Virtual). 14 July 2020.

Human-Usable Password Schemas: Beyond Information-Theoretic Security

Meeting of the Minds, CMU. 4 May 2016.

AWARDS AND HONORS

- **Exemplary Senior Honors Thesis**, CMU May 2016
- **Carnegie Mellon School of Computer Science College Honors** May 2016
- **Carnegie Mellon University Honors** May 2016
- **Senior Leadership Recognition Award**, CMU May 2016

PROFESSIONAL EXPERIENCE

J.P. Morgan

New York, NY

Research Intern

Summer 2019

- Investigated reinforcement learning techniques for detecting and responding to shifts in environment dynamics.

Google

New York, NY

Software Engineer

2016-2018

- Designed, engineered, and launched features to enhance Google Search by supplementing Knowledge Graph with User-Generated Content (UGC).
- Designed and constructed data processing pipelines to manipulate UGC and share the resulting aggregate data with Google users.
- Performed statistical analysis of key metrics on data generation and feature usage to inform design decisions and drive feature improvements.

Pinterest

San Francisco, CA

Application Security Intern

Summer 2015

- Implemented role-based access control on top of OAuth for Pinterest's internal development tools and data stores.

Microsoft

Bellevue, WA

Software Development Engineering Intern

Summer 2014

- Designed and engineered a hybrid content- and collaboration-based recommendation engine for Bing Local Search.

Institute for Physical Sciences, Inc.

McLean, VA

Software Development Intern

Summer 2013

- Developed, implemented, and tested a statistically-derived contextual thesaurus using syntactic ngrams parsed by Google Research.

TEACHING

- **Head Teaching Assistant** at Carnegie Mellon University
 - Probabilistic Graphical Models (10-708)* Spring 2020
 - Advanced Deep Learning (10-707)* Spring 2019
 - Andrew's Leap* Summer 2010, 2012
- **Teaching Assistant** at Carnegie Mellon University
 - Mathematical Foundations for Computer Science (15-151)* Spring 2013
 - Summer Academy of Math and Sciences (SAMS)* Summer 2011

SERVICE

Reviewing

ICML: 2020, 2021 (Expert Reviewer)
NeurIPS: 2019, 2020 (Top 10%), 2021 (Top 10%), 2022
ICLR: 2022
AISTATS: 2022 (Top 10%)

Outreach

CMU AI Undergraduate Mentor 2019-Present
– *Mentoring undergraduates from underrepresented groups who want to get involved in AI research.*

CMU Graduate Application Support Program (GASP) Mentor 2020-Present
– *Providing support and detailed feedback on graduate school applications to students from traditionally underrepresented backgrounds. cs.cmu.edu/~gasp/*