Elan Rosenfeld

Website: cs.cmu.edu/~elan
Email: elan.rosenfeld@gmail.com

Education

Carnegie Mellon University — Pittsburgh, PA
Ph.D. in Machine Learning. — 2018–2024

- Thesis Title: "Understanding, Formally Characterizing, and Robustly Handling Real-World Distribution Shift"
  Advised by Andrej Risteski and Pradeep Ravikumar

Carnegie Mellon University — Pittsburgh, PA
B.S. in Computer Science, additional major in Statistics & Machine Learning — 2012–2016

- Senior Thesis Title: "Human-Usable Password Schemas: Beyond Information-Theoretic Security"
  Advised by Manuel Blum and Santosh Vempala
  Awarded "Exemplary Senior Honors Thesis"

Selected Works

- E. Rosenfeld and A. Risteski. Outliers with Opposing Signals Have an Outsized Effect on Neural Network Optimization. *ICLR 2024*
- E. Rosenfeld and S. Garg. (Almost) Provable Error Bounds Under Distribution Shift via Disagreement Discrepancy. *NeurIPS 2023*
- E. Rosenfeld, P. Ravikumar, and A. Risteski. Domain-Adjusted Regression or: ERM May Already Learn Features Sufficient for Out-of-Distribution Generalization. *NeurIPS DistShift 2022*
- E. Rosenfeld, P. Ravikumar, and A. Risteski. The Risks of Invariant Risk Minimization. *ICLR 2020*
- J. Cohen, E. Rosenfeld, and J. Z. Kolter. Certified Adversarial Robustness via Randomized Smoothing. *ICML 2019*

More Publications

- E. Rosenfeld and N. Rosenfeld. One-Shot Strategic Classification Under Unknown Costs. *ICML 2024*
- S. Saengkyongam, E. Rosenfeld, P. Ravikumar, N. Pfister, and J. Peters. Identifying Representations for Intervention Extrapolation. *ICLR 2024*
- S. Buchholz, G. Rajendran, E. Rosenfeld, B. Aragam, B. Schölkopf, and P. Ravikumar. Learning Linear Causal Representations from Interventions under General Nonlinear Mixing. *NeurIPS 2023*
- Y. Chen, E. Rosenfeld, M. Sellke, T. Ma, and A. Risteski Iterative Feature Matching: Toward Provable Domain Generalization with Logarithmic Environments. *NeurIPS 2021*
- B. Liu, E. Rosenfeld, P. Ravikumar, and A. Risteski. Analyzing and Improving the Optimization Landscape of Noise-Contrastive Estimation. *ICLR 2021*
- I. Apostolopoulou, I. Char, E. Rosenfeld, and A. Dubrawski. Deep Attentive Variational Inference. *ICLR 2021*
- E. Rosenfeld, P. Ravikumar, and A. Risteski. An Online Learning Approach to Interpolation and Extrapolation in Domain Generalization. *AISTATS 2021*
- E. Rosenfeld, E. Winston, P. Ravikumar, and J. Z. Kolter. Certified Robustness to Label-Flipping Attacks via Randomized Smoothing. *ICML 2020*

Workshop Papers and Manuscripts

- S. Kangaslahti, E. Rosenfeld, and N. Saphra. Loss in the Crowd. In *ICML Workshop on Mechanistic Interpretability*, 2024.
- E. Rosenfeld, P. Nakkiran, H. Pouransari, O. Tuzel, and F. Faghri. APE: Aligning Pretrained Encoders to Quickly Learn Aligned Multimodal Representations. In *Has It Trained Yet? NeurIPS Workshop*, 2022.
- I. Apostolopoulou, E. Rosenfeld, and A. Dubrawski. Self-Reflective Variational Autoencoder. In *ICLR Workshop on Hardware Aware Efficient Training*, 2021.
- E. Rosenfeld, S. Vempala, and M. Blum. Human-Usable Password Schemas: Beyond Information-Theoretic Security. *CMU Senior Honors Thesis*, 2016.

Invited Talks

Building a Foundation for Trustworthy Machine Learning
AWS Fundamental Research, May 2024
FAIR Labs, May 2024
Vector Institute for Artificial Intelligence, April 2024
Microsoft Research, April 2024
Google Research, April 2024
New York University, April 2024
University of Pennsylvania, April 2024.
Carnegie Mellon University, March 2024
Boston University, March 2024.
University of Texas at Austin, March 2024.

Outliers with Opposing Signals Have an Outsized Effect on Neural Network Optimization
Deep Learning: Classics and Trends, ML Collective (Virtual), January 2024.
Carnegie Mellon University, December 2023.

Generalizing to New Distributions via Invariance: Challenges and Opportunities
Technion – Israel Institute of Technology, April 2022.

The Risks of Invariant Risk Minimization
Technion – Israel Institute of Technology, March 2022.
Princeton University (Virtual), July 2021.
Carnegie Mellon University (Virtual), April 2021.

Human-Usable Password Schemas: Beyond Information-Theoretic Security
Carnegie Mellon University, May 2016.

Awards and Honors

Exemplary Senior Honors Thesis, Carnegie Mellon University 2016
School of Computer Science College Honors, Carnegie Mellon University 2016
University Honors, Carnegie Mellon University 2016
Senior Leadership Recognition Award, Carnegie Mellon University 2016
Semi-Finalist, National Merit Scholarship Program 2011
1st Place, League of Women Voters Mock Election Design 2010
Finalist, North American Computational Linguistics Olympiad 2010

Teaching

- **Head Teaching Assistant** at Carnegie Mellon University

  Probabilistic Graphical Models (10-708) Spring 2020

  – Designed math and coding assignments on various topics in PGMs, wrote solutions, and coordinated with other TAs to ensure productive office hours.

  – Advised multiple groups of students on their semester-long research project.

  Advanced Deep Learning (10-707) Spring 2019

  – Worked one-on-one with instructor to design all assignments and write solutions.

  – Created and presented several course lectures.

  Summer Academy of Math and Sciences (SAMS) Summer 2011

  Andrew's Leap Summer 2010, 2012

- **Teaching Assistant** at Carnegie Mellon University

  Mathematical Foundations for Computer Science (15-151) Spring 2013

Service

Organizing

Sole lead organizer of the ICML 2022 Workshop on Principles of Distribution Shift (PODS)

Reviewing

Reviewing Awards

ICML Expert Reviewer x2
NeurIPS Top 10% x3
AISTATS Top 10%

ICML: 2020-present
NeurIPS: 2019-present
ICLR: 2022-present
AISTATS: 2022
TMLR: 2022-present
JMLR: 2023-present
CLeaR: 2024

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/*

Teaching Assistant at the CMU Summer Academy for Math and Science (SAMS) Summer 2011

– *Taught C programming and LEGO robotics to students from underrepresented communities.*

Instructor at Google Pittsburgh Spring/Summer 2010

– *Taught Android programming to students from underserved neighborhoods as part of a computer science enrichment program.*

Professional Experience

Apple ML Research — Seattle, WA
Research Intern — Summer 2022

- Studied approaches to learning large aligned multimodal encoders more efficiently, reaching better zero-shot image classification performance with substantially less paired data and compute.

J.P. Morgan Research — 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 & constructed processing pipelines to manipulate UGC and share resulting aggregate data with 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 internal development tools and data stores.

Microsoft — Bellevue, WA
Software Development Engineering Intern — Summer 2014

- Designed & 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.
- Contributed to a larger project for automatic abstraction-based document summarization.