Harrison Rosenberg

rosenberg.harrison.j@gmail.com hrosenberg@ece.wisc.edu 3107406232
harrisonrosenberg.com wiscprivacy.com Citizen of United States of America

Research Interests

Generative AI, Human Factors in AI, Machine Learning Biases, Face Recognition, Robust Machine Learning, Privacy Implications of Machine Learning, Feature Extraction and Identification

Education

2017–2024 **Ph.D.** in Electrical and Computer Engineering — University of Wisconsin–Madison

- Advisor: Professor Kassem Fawaz
- Dissertation: Empirically and Theoretically Understanding Machine Learning Fairness through Face Recognition

2017–2019 **M.S.** in Electrical Engineering — University of Wisconsin–Madison

2013–2017 **B.S.** in Electrical Engineering and Computer Sciences — University of California, Berkeley

Work Experience

2017–2024	Research Assistant	University of Wisconsin–Madison
2021–2022	Teaching Assistant	University of Wisconsin–Madison
2016–2017	(Head) Undergraduate Student Instructor	University of California, Berkeley
2016	Software Engineering Intern	StubHub Inc
2014, 2015	Programming Intern	Lieberman Software Corporation

Projects

2022–Current Faces and Generative Modeling — University of Wisconsin–Madison

Text-to-image diffusion models have achieved widespread popularity due to their unprecedented image generation capability. Their ability to synthesize and modify human faces has spurred research into using **generated images in face recognition pipelines**. **Limitations** of generative model utility in face generation are identified with a combination of **qualitative and quantitative** measures. Innovations apply beyond face generation: similar analysis is applicable to large language models and audio generators.

Tools and Concepts: Generative AI, Diffusion Networks, PyTorch, GPT-3.5/GPT-4, CUDA, Pandas, Jupyter, NumPy, User Surveys, Dataset Synthesis, Mathematical Modeling

Datasets: Synthetic Datasets, Labeled Faces in the Wild, CelebA, LAION

2021–2023 Face Obfuscation: Limitations and Fairness — University of Wisconsin–Madison

Face obfuscation systems promise to mitigate privacy incursions associated with widespread automated face recognition. In this project, face obfuscation systems themselves are **shown to yield significant privacy incursions**. Furthermore, privacy incursions have strong **dependence on demographics**.

Tools and Concepts: TensorFlow, PyTorch, CUDA, Pandas, NumPy, Adversarial Learning, Mathematical Modeling

Datasets: Labeled Faces in the Wild, VGGFace2

2020–2022 DARPA Guaranteeing AI Robustness Against Deception (GARD) — University of Wisconsin–Madison

Modern reconnaissance techniques yield an abundance of data far too large for human analysts to carefully sort through entirely unaided. Both computerized image and audio recognition systems have potential to **reduce burden** on human intelligence analysts, yet they are **vulnerable** to acts of deception. This DARPA-funded project explores techniques to **harden computational systems against military deception**.

Tools and Concepts: PyTorch, CUDA, Adversarial Robustness Toolbox, Armory, NumPy, Adversarial Learning

Datasets: APRICOT, LibreSpeech, So2Sat, CIFAR-10

2019-2022 Convergence Bounds for Fairness — University of Wisconsin–Madison

Among the most common characterizations of machine learning fairness are sample complexities for multicalibration error convergence. The main result demonstrates **sample complexities** for multicalibration error convergence can be obtained by **reparametrizing pre-existing bounds** for empirical risk minimization learning.

Tools and Concepts: PyTorch, CUDA, Scikit-learn, Pandas, Jupyter, NumPy, Statistical Learning Theory, Concentration Inequalities, Confidence Intervals, Probabilistic Modeling

Datasets: UCI Adult, COMPAS, CelebA

Publications

[1] **Harrison Rosenberg***, Shimaa Ahmed*, Guruprasad Viswanathan Ramesh*, Ramya Korlakai Vinayak, Kassem Fawaz. Limitations of Face Generation *Thirty-Eighth AAAI Conference on Artificial Intelligence, 2024*

[2] **Harrison Rosenberg**, Brian Tang, Kassem Fawaz, Somesh Jha. Fairness Properties of Face Recognition and Obfuscation Systems. *The 32nd USENIX Security Symposium, 2023*

[3] **Harrison Rosenberg**, Robi Bhattacharjee, Kassem Fawaz, Somesh Jha. An Exploration of Multicalibration Uniform Convergence Bounds. *Arxiv*

[4] Zachary Charles, **Harrison Rosenberg**, Dimitris Papailiopoulos. A Geometric Perspective on the Transferability of Adversarial Directions. In *The 22nd International Conference on Artificial Intelligence and Statistics. PMLR, 2019.*

Teaching Experience

2021–2022 Teaching Assistant — University of Wisconsin–Madison

- Hold office hours, grade assignments, and mentor students
- Set homework and exam rubrics
- Both traditional and flipped format courses
- **ECE 203** Signals, Information, and Computation; **ECE 431** Digital Signal Proc.; and **ECE 533** Image Proc.

2016–2017 (Head) Undergraduate Student Instructor — University of California, Berkeley

- Responsible for $250k+ staff budget: Hired and led 50+ total student instructors, graders, and lab assistants
- Manage logistics and administrivia in a course of 500+ students per semester
- Lead discussion sections and exam review sessions
- **EE16A** Designing Information Devices and Systems I

Professional Activities

2023	Talk: Publication [2]	USENIX Security 2023
2022	Talk: Publication [2]	UW–Madison Computer Engineering Seminar Series
2020, 2022	Reviewer	International Conference on Artificial Intelligence and Statistics (AIStats)
2021	Sub-Reviewer	IEEE Symposium on Security and Privacy (Oakland)
2019–2020	Student Member	Department Ph.D. Graduate Committee

Notable Achievements

2023	Dissertator Travel Award	University of Wisconsin–Madison
2022–2023	Awardee (C/O Kassem Fawaz)	American Family Funding Initiative
2018–2020	NSF Trainee	LUCID, University of Wisconsin–Madison
2018	QIF Fellowship Finalist	Qualcomm Inc.
2017–2018	Wisconsin Distinguished Graduate Fellow	University of Wisconsin–Madison

Skills

Strong Knowledge of AI, Machine Learning, Signal Processing, Probability, Linear Algebra, and LaTeX
Programming Languages (in order of proficiency): Python, MATLAB, SQL, Java, C#
Other: Red Cross First-Aid/CPR/AED Certification, FCC Technician Class Licensee