

# Yahav Bechavod

Postdoctoral Researcher

University of Pennsylvania  
Department of Computer and Information Science  
3333 Chestnut Street  
Philadelphia, PA 19104

Phone: 267-506-3795  
Email: [yahav@seas.upenn.edu](mailto:yahav@seas.upenn.edu)  
Webpage: [yahavbe.github.io](https://yahavbe.github.io)

---

## Research Interests

I study the foundations of **reliable ML for consequential decision-making**. My work draws on and extends ideas from **machine learning, optimization, and algorithmic game theory**. Through this interdisciplinary lens, I approach reliability balancing theoretical rigor with practical applicability, aiming to build a unified understanding of how ML systems can make sound, equitable, and responsible decisions in complex real-world settings.

## Education

---

Ph.D.	Computer Science, The Hebrew University Dissertation: “Socially Aware Algorithms: Learning, Fairness, and Incentive Awareness” Dissertation Committee: Profs. Aaron Roth, Guy Rothblum, Amit Daniely	2022
M.S.	Computer Science, The Hebrew University Graduated Summa Cum Laude Thesis: “Algorithms for Fairness in Binary Classification” Thesis Advisor: Prof. Katrina Ligett	2018
B.S.	Mathematics and Computer Science (Double Major) The Hebrew University	2015

## Grants and Awards

---

Penn Arts & Sciences AI x Science Fellowship	2024–2026
Apple Seed Grant	2024
Israeli Council for Higher Education Postdoctoral Fellowship	2023–2024
KLA Outstanding Doctoral Research Award	2021
Apple Scholars in AI/ML PhD Fellowship	2020–2022
Charles Clore Foundation PhD Fellowship (declined prize money)	2020–2023

## Research Experience

---

Postdoctoral Researcher, University of Pennsylvania Department of Computer and Information Science Host: Prof. Aaron Roth	2023–present
<ul style="list-style-type: none"><li>Worked on 5 research projects: published 2 papers in ICML; 2 under review; 1 working paper</li><li>Secured Apple Seed Grant funding for a new line of research on fairness in online settings</li><li>Mentored PhD students from Roth’s group in various research projects.</li></ul>	

- Developed new frameworks integrating constrained learning and omniprediction with multiple strategic decision makers

**Visiting Scholar**, Simons Institute (UC Berkeley) May–July 2019

- Invited participant of the “Algorithmic Fairness” summer cluster.
- Worked with Prof. Steven Wu on a project on individual fairness in online learning.
- Resulting paper was selected as an oral presentation at NeurIPS.

**Visiting Student Researcher**, University of Pennsylvania Jul–Oct 2017

- Worked with Prof. Aaron Roth and Prof. Michael Kearns on a project on learning in sequential settings.
- Resulting paper was accepted and presented at NeurIPS.

## Service and Leadership

---

### Journal Referee

Journal of Machine Learning Research (JMLR)

### Reviewer

International Conference on Machine Learning (ICML)	2022–2024
Conference on Neural Information Processing Systems (NeurIPS)	2021–2024
ACM Conference on Fairness, Accountability, and Transparency (FAccT)	2021

### Ethical Reviewer

International Conference on Machine Learning (ICML)	2025
Conference on Neural Information Processing Systems (NeurIPS)	2021–2024

### Co-Organizer

December 2021

NeurIPS 2021 Workshop on Learning and Decision-Making with Strategic Feedback ([StratML'21](#))

- Co-initiated and co-organized the workshop along with Hoda Heidari, Eric Mazumdar, Celestine Mendler-Dünner, Tijana Zrnic, and Aaron Roth.
- Co-coordinated the participation of 9 invited speakers and panelists: Francesca Parise, Lillian Ratliff, Nir Rosenfeld, Moritz Hardt, Navin Kartik, Yang Liu, Cristobal Cheyre, Jon Kleinberg, and Steven Wu.
- Oversaw the reviewer recruitment and reviewing process of paper submissions and selection of contributed talks, with 42 peer-reviewed contributed papers presented.
- Co-moderated panel on “Strategic classification and modeling of agent behavior”, with Moritz Hardt, Navin Kartik, and Yang Liu.

## Teaching and Mentoring

---

### Teaching Assistant

Fall 2015, Fall 2016

Hebrew University School of Engineering and Computer Science  
Class: Introduction to Algorithms (CS67504)

- Graded assignments for approximately 300 students.
- Interacted with students on weekly office hours.
- Held midterm oral examinations.

### **Teaching Assistant**

Spring 2016, 2017

Hebrew University School of Engineering and Computer Science  
Class: Introduction to Machine Learning (CS67577)

- Graded assignments for approximately 200 students.
- Interacted with students on weekly office hours.
- Mentored students as part of a class Hackathon.

### **Graduate Mentor**

2018

Advised a Master's student on a project of implementing a practical framework for fair classification. Designed the project and Held weekly mentoring meetings. The work resulted in an open-source Python code available for public use [here](#).

## **Publications**

---

### **Working Papers**

#### **Oracle-Efficient Learning with Long-Term Constraints**

Yahav Bechavod, Aaron Roth  
Working Paper, 2025

### **Work In Submission**

#### **Dynamic Regret Bounds for Online Omniprediction with Long-Term Constraints**

Yahav Bechavod, Jiuyao Lu, Aaron Roth  
Submitted, 2025

#### **Online Omniprediction with Long-Term Constraints**

Yahav Bechavod, Jiuyao Lu, Aaron Roth  
Submitted, 2025

Preliminary version accepted at NeurIPS 2025 Workshop on Constrained Optimization for Machine Learning

### **Peer-Reviewed Published Papers**

#### **Monotone Individual Fairness**

Yahav Bechavod  
International Conference on Machine Learning (ICML), 2024

#### **Individually Fair Learning with One-Sided Feedback**

Yahav Bechavod, Aaron Roth  
International Conference on Machine Learning (ICML), 2023

#### **Information Discrepancy in Strategic Learning**

Yahav Bechavod, Chara Podimata, Steven Wu, Juba Ziani  
International Conference on Machine Learning (ICML), 2022

## **Gaming Helps! Learning from Strategic Interactions in Natural Dynamics**

Yahav Bechavod, Katrina Ligett, Steven Wu, Juba Ziani

International Conference on Artificial Intelligence and Statistics (AISTATS) 2021

## **Metric-Free Individual Fairness in Online Learning**

Yahav Bechavod, Christopher Jung, Steven Wu

Advances in Neural Information Processing Systems (NeurIPS) 2020

Selected for Oral Presentation (long talk, 1% of submissions)

## **Equal Opportunity in Online Classification with Partial Feedback**

Yahav Bechavod, Katrina Ligett, Aaron Roth, Bo Waggoner, Steven Wu

Advances in Neural Information Processing Systems (NeurIPS) 2019

## **Penalizing Unfairness in Binary Classification**

Yahav Bechavod, Katrina Ligett

Fairness, Accountability, and Transparency 2017

---

## **Selected Invited and Contributed Talks**

### Individual Fairness in Online Classification

- University of Pennsylvania CS Theory Seminar. Philadelphia, PA. February 2023.
- Computation for Fairness (TOC4Fairness) Seminar Series. Virtual. February 2023.

### Information Discrepancy in Strategic Learning

- ICML Spotlight Presentation. Baltimore, Maryland. July 2022.

### Metric-Free Individual Fairness in Online Learning

- Apple Machine Learning Speaker Series. Virtual. June 2021.
- Symposium on the Foundations of Responsible Computing (FORC). Virtual. June 2021.
- Neural Information Processing Systems (NeurIPS) 2020. Virtual. December 2020.
- Hebrew University Machine Learning Seminar Series. Virtual. February 2020.

### An Introduction to Algorithmic Fairness

- Hebrew University Federmann Center for the Study of Rationality. Eilat, Israel. February 2020.

### Equal Opportunity in Online Classification with Partial Feedback

- Simons Institute Workshop on Developments in Research on Fairness. Berkeley, California. July 2019.

---

## **Poster Presentations**

Bechavod, Y. Monotone individual fairness. Poster Presentation. The 41st International Conference on Machine Learning. Vienna, Austria. 2024.

Bechavod, Y. & Roth, A. Individually Fair Learning with One-Sided Feedback. Poster Presentation. The 40th International Conference on Machine Learning. Honolulu, Hawaii. 2023.

Bechavod, Y., Podimata, C., Wu, Z. S., & Ziani, J. Information Discrepancy in Strategic Learning. Poster Presentation. The 39th International Conference on Machine Learning. Baltimore, Maryland. 2022.

Bechavod, Y., Ligett, K., Wu, S., & Ziani, J. Gaming Helps! Learning from Strategic Interactions in Natural Dynamics. Poster Presentation. The 24th International Conference on Artificial Intelligence and Statistics. Virtual. 2021.

Bechavod, Y., Jung, C., & Wu, Z. S. Metric-Free Individual Fairness in Online Learning. Poster Presentation. Advances in Neural Information Processing Systems 33. Virtual. 2020.

Bechavod, Y., Ligett, K., Roth, A., Waggoner, B., & Wu, Z. S. Equal opportunity in online classification with partial feedback. Poster Presentation. Advances in Neural Information Processing Systems 32. Vancouver, Canada. 2019.

Bechavod, Y. & Ligett, K. (2017). Penalizing Unfairness in Binary Classification. Poster Presentation. Fairness, Accountability, and Transparency. Halifax, Canada. 2017.

## References

---

### **Aaron Roth, PhD**

Henry Salvatori Professor of Computer and Cognitive Science  
The University of Pennsylvania  
3333 Chestnut Street, Philadelphia, PA 19104  
Phone: 917-270-9542  
Email: aaroth@cis.upenn.edu

### **Katrina Ligett, PhD**

Professor of Computer Science  
Director, Federmann Center for the Study of Rationality  
The Hebrew University  
Givat Ram, Jerusalem, Israel 9190401  
Phone: +972-58-796-1416  
Email: katrina@cs.huji.ac.il

### **Guy Rothblum, PhD**

Research Scientist, Apple,  
Professor, Department of Computer Science and Applied Mathematics,  
Weizmann Institute of Science.  
Rehovot, Israel, 7610001  
Phone: 857-998-9820  
Email: rothblum@alum.mit.edu

### **Steven Wu, PhD**

Associate Professor, School of Computer Science,  
Carnegie Mellon University  
5000 Forbes Ave. Pittsburgh, PA 15213  
Phone: 917-562-5830  
Email: zstevenwu@cmu.edu

### **Juba Ziani, PhD**

Assistant Professor, School of Industrial and Systems Engineering  
Georgia Tech,  
755 Ferst Drive NW Atlanta, GA 30332  
Phone: 626-905-0523  
Email: jziani3@gatech.edu