Bias in AI Reading Group Fall 2026

Thursdays 3-4pm | Sherrerd 008

Sign up to get mailing list updates here

This reading group focuses on justice, bias and fairness concerns in AI from a wide range of disciplinary perspectives—including CS, law, philosophy, and STS. The group is open to anyone interested, regardless of whether bias/fairness is your primary research focus. For an example of the kinds of topics we have discussed in the past, check out our reading list from the previous three semesters.

We meet every week. Sessions take a variety of formats, including book club discussions, presentations of work in progress/recent work by members, presentations of papers in the research area or series of papers on a theme, and guest speakers.

If you have any questions, feel free to reach out to Patty Liu (patty.liu@princeton.edu) or Kara Schechtman (schechtk@princeton.edu).

Expectations for Participants

Participants are encouraged to review any assigned readings in advance. But don't worry if you haven't read. You're still very welcome to attend and participate!

Discussion will be in person. For sessions with a presentation component, Zoom will be available in case you want to listen in, but we will not be monitoring/addressing the Zoom audience.

Expectations for Presenters

If you're interested in presenting, please sign up below with your name and a link to anything you'd like people to read in advance. You can present your own work, a paper you find particularly interesting, or a series of papers on a common theme.

Presenters should provide a summary of the paper's key contributions, including any necessary background for understanding it, and prepare a few questions to guide the discussion. We welcome different presentation formats: Especially if you're seeking feedback on your own work and would prefer a different structure, feel free to adapt as needed.

Presentation Schedule & Sign-Ups

Contact Kara or Patty if you need edit access to the doc to sign up!

Date	Presenter	Topic/Paper
September 10, 2026	N/A	Welcome & Book Club: Empire of AI Prologue
September 17, 2026	N/A	Book Club: Empire of AI Part I
September 24, 2026	N/A	Book Club: Empire of AI Part II
October 1, 2026	N/A	Book Club: Empire of AI Part III
October 8, 2026	N/A	Book Club: Empire of AI Part IV and Epilogue
October 15, 2026	Vikram	Straightened or Caricatured: How Text-to-Image Models Portray Queer Gender Identities
October 22, 2026	No meeting - fall break	
October 29, 2026	Kara	Topic: Data center impacts and activism Reading: What Happens When Data Centers Come to Town?
November 5, 2026	N/A	Book Club: Data Driven, Chapter 1, & Chapter 2
November 12, 2026	N/A	Book Club: Data Driven, Chapter 3 & Chapter 4
November 19, 2026	N/A	Book Club: Data Driven, Chapter 5 & Chapter 6
November 26, 2026	No meeting - Thanksgiving	
December 3, 2026	N/A	Book Club: Data Driven, Chapter 7 & Chapter 8

Paper Pool 📄🛟

If you're interested in presenting but not sure what to present, please feel free to claim any paper from the pool. If you're interested in discussing a paper but not sure if you'd like to present it, feel free to add it to the pool.

1. Tipping the Balance: Predictive Algorithms and Institutional Decision-Making in Context
2. Amazing Things Come from Having Many Good Models
3. Auditing large language models for race & gender disparities: Implications for artificial intelligence-based hiring
4. Identity-related Speech Suppression in Generative AI Content Moderation
5. Prediction, Performativity, and Potential Outcomes: Communicative Rationality in Prediction-Allocation Problems
6. Risk scores, label bias, and everything but the kitchen sink
7. One Label, One Billion Faces: Usage and Consistency of Racial Categories in Computer Vision
8. Tech Ethics: Speaking Ethics to Power, or Power Speaking Ethics?
9. What's Sex Got to Do With Fair Machine Learning?
10. Disparate Treatment and Discriminatory Harm
11. The Contestation of Tech Ethics: A Sociotechnical Approach to Technology Ethics in Practice