

Introduction to the Social Consequences of Computing

EECS 298

Winter 2024

4 Credits

Instructor: [Ben Fish](#)

Meeting times:

Lecture MW 1:30-3PM

Lab F 1:30-2:30PM

Overview: Computing is now used in every facet of life affecting countless people, including in social media, advertising, policy decisions, admissions, and hiring. In doing so, the process of computing and algorithm design now involves understanding the role of computing in society.

This class will introduce you to the ways in which applications of computing affect social institutions and how these social consequences produce questions about **how to conceptualize, critique, and ensure our all-too-human values in computing**. To accomplish this, we will explore how the design of computing systems, including machine learning systems, bureaucratic decision making systems, social media, and so forth intersect with issues of fairness, discrimination, privacy, and other values. We will critically examine the philosophical and sociological underpinnings of these values and the strategies commonly used to promote them, and seek to connect these conceptualizations to the emerging algorithmic tools proposed for promoting those values. This class will heavily feature **in-class discussion** to do this work. In addition, in order to practice reasoning through these problems, this class will feature programming in **Python**. No previous programming experience in Python is needed.

Prerequisites: EECS 280 or permission of instructor

Requirements satisfied: Satisfies the **EECS 496** requirement for CS & DS majors

Learning objectives:

This is a dive into an area that's still largely unsettled. Rather than providing an off-the-shelf methodology, the focus is on providing the tools to examine, analyze and construct computational systems with their social consequences in mind. In particular, the goals of this class are to be able to do the following:

- To explain how examples of computing systems from the (recent) history of computer science teach us about the harms of those systems
- To identify the potential harms created or maintained by computing systems
- To analyze existing algorithms or computing systems to understand their consequences in society in terms of complex and contested values
- To evaluate and critique potential solutions to overcome the harms of computing systems
- To integrate these findings into your own practice of building computational systems
- To become familiarized with Python

Topics:

- Sociotechnical systems and their properties: measurement, constructs, and proxies, stakeholders and affordances
- Harm in computing: discrimination & oppression, race & gender, feedback loops, personalization, and phrenology
- Surveillance and privacy: societies of control, worker surveillance, anonymization and PII, targeted advertising, contextual integrity
- Fairness and justice: ethical theory, machine learning and fairness, praxis and interventions
- Information and platforms: content moderation and recommendation, polarization, misinformation/disinformation, platform power and governance, activism

Homework: This class will heavily feature **readings**, due each class. There will also be required responses to the readings, including a presentation. In addition, there will be programming assignments in Python. The readings will be a chance to engage more heavily with the conceptual underpinnings of the topics discussed in class, while the programming assignments will be a chance to see how harm can be enacted or perpetuated by systems students could create.

Labs: Labs will consist of programming assignments, and are an opportunity to focus on the Python component of this class. Assignments will be an opportunity to simultaneously practice Python and engage in the process of critical reflection of what you code.

Final Project: There will be a final group project. Each group will tackle a particular computing system of interest. They will use the tools learned in this class to explore the social consequences of computing in that particular domain, producing a final report and/or code to demonstrate their findings and report out on their design process and methodology. At the end of the semester, they will present their work to the class.

Textbooks:

- Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*, 2019.
- Virginia Eubanks. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, 2018.
- Kirsten Martin. *Ethics of Data and Analytics: Concepts and Cases*, 2022.