AI, Decision-Making, & Society
FALL SEMESTER 2024

INSTRUCTORS: Dylan Hadfield-Menell, Ashia Wilson
aidms-staff@mit.edu

Class Description

This class focuses on how values interact with AI decisions & decision-making systems. Students will design extensive evaluations and use exploratory data analysis in order to understand how to develop AI systems in anticipation of downstream impacts. The course also centers on participatory AI design, and emphasizes the importance of equity, transparency, and pluralistic alignment, as well as the challenges of data collection and representativeness and bias.

Logistics

Class Meetings

Lectures: Tuesday, Thursday 1-2:30 (Location: TBA)

Recitations: Friday

Websites

- We will post materials and collect homework through Canvas
- We will answer questions and moderate a discussion forum on Piazza

Pre-requisites

6.100A(?)

Grading Scheme

Final grade will be determined according to the following breakdown: homework assignments (50%), participation in discussion sessions (20%) and final project (30%)

Homework assignments: We will have six problem sets. Each problem set will contain additional problems intended for students taking the graduate version of the class. Typically, each problem set will be due between one to two weeks of its release. You are expected to submit your homework on time. However, because unexpected situations might occur, you have a budget of two late submissions. Each late submission allows you to have four extra days to complete the assignment. (You do not need to inform us about your use of your budget, just make sure you do not exceed it.) Beyond this mechanism, late homework will not be accepted, except in extreme (and, when requested, documented) circumstances.

Discussion sessions: Recitations will be in the form of discussion sessions. Each of these sessions will have a reading material as well as question(s)—meant to ground the discussion—assigned to it. The material needs to be read ahead of the discussion session, and you will also be required to supply your (short, 10–20 sentences) answers to the questions in advance. You are expected to attend (and take active part in) all the discussion sessions. Excused absences can be requested with S^3 support.

Final project: The final project will ask each student (or a 2–4 student team) to explore—in the form of an essay and an in-class presentation—one of the topics discussed in the class in more depth. This might involve further exploration of one of the phenomena, approaches and/or scenarios presented in class; application of one of the discussed methodology to a new scenario; or a cross-cutting synthesis of (a subset of) the class material. (For the graduate version of the class, the expectations of the quality and scope of

the project will be commensurably higher.) The preliminary, 2–3 page project proposal will be due in the middle of the semester