Pre-read: privacy and ethics (part 1)

Throughout the course we've made brief mention of normative limits to prediction: privacy, fairness, and ethics. We will discuss them this week.

The "inputs" and "outputs" of predictive models may both raise privacy concerns. That is, the data that is used to make predictions may be sensitive (think Fragile Families Challenge, or browsing histories that are used for ad targeting) and it would be a breach of privacy if those data were to be collected without consent or to fall into the wrong hands. Separately, the predictions themselves may be considered private: think of an embarrassing targeted ad when someone is watching over your shoulder. More gravely, inferring or predicting private attributes may allow powerful entities to manipulate people (political actors inferring personality traits for election interference) or make adverse decisions about them (health insurance companies raising premiums based on pre-existing conditions due to predicted increases in future healthcare costs).

The first reading is a very broad overview of commercial data collection today: what data is collected and how it is used. It is ok to skim this or pick some sections to read. Note that this report is written by privacy advocates and is different in tone and substance from most of the papers you've read. As you encounter various types of inferences and predictions based on personal data, here are some questions to think about:

- Where do these applications fall on the spectrum between readymade (repurposing existing data) and custommade (collecting data for a specific purpose)?
- Do you have an intuition for how accurate these predictions are? And how important is accuracy to commercial success?

The second reading is an excerpt from a book called The Power of Habit. It is well known for an anecdote about how Target figured out that a teenager was pregnant before her father did, but it has deeper lessons of relevance to us. In particular, it is a case study of how the commercial value of personal data arises in part from the ability to use it to persuade — and perhaps manipulate — people, and not just make inferences or predictions about them.

The third reading is the paper that led to the birth of Cambridge Analytica. Like all our readings, we will read this paper critically. Incidentally, the lead author, Michal Kosinski, is a coauthor of several other papers with similar themes, including "Psychological targeting as an effective approach to digital mass persuasion" (PNAS 2017; arguing for the effectiveness

of the kind of targeting that Cambridge Analytica is alleged to have engaged in) and “Deep neural networks are more accurate than humans at detecting sexual orientation from facial images” (JPSP 2018; a paper that drew vociferous criticism). In class, we will discuss the ethics of conducting this kind of research.