Course Syllabus

Instructor: Jens Ludwig

Time and Location: Mondays 1:30-4:30pm, Keller Center 0021

Contact: jludwig@uchicago.edu

Office Hours:

Teaching Assistant: Rafael Batista [rafael.batista@chicagobooth.edu] and Megan Kang [megankang@uchicago.edu]

Course Description

It is hard to name a sector that will not be dramatically affected by artificial intelligence (or machine learning), from the private sector to government and nonprofits. There are many excellent courses that teach you the *mechanics* behind these innovations -- helping you develop an engineering skill set. This course takes a different approach. It is aimed at people who want to *deploy* these tools, either in business or policy, whether through start-ups or within a large organization. While this requires some knowledge of how these tools work, that is only a small part of the equation, just as knowing how an engine works is a small part of understanding how to drive. What is really needed is an understanding of what these tools do well, and what they do badly. This course focuses on giving you a *functional*, rather than mechanistic, understanding. By the end, you should be an expert at identifying ideal (and problematic) use-cases and thereby well-placed to create new products, businesses and policies that use artificial intelligence.

Prerequisites

There are no formal requirements for this class although we assume students have familiarity with basic concepts in statistics and regression analysis.

Format

The course is intended to be highly interactive, where students are expected to contribute their ideas to the discussion. There will be a series of activities and experiential material which will require active participation. Throughout the term the modules are designed to craft a simpler mental model of AI, starting with the fundamentals of what exactly makes AI so powerful and going through to what are good (and bad) uses of AI and why. Students should expect a mix of lectures, hands-on activities, discussions, and guest speakers.

Requirements and Grading

Considering the aim of the course is to develop a stronger, intuitive understanding of AI, there will be no in-class exams. Instead, students will be expected to complete 7 take-home assignments--memos--that will help digest the material we are teaching. Students will also be required to complete a final project which they will be expected to complete outside of class.

Participation (15%):

Students are expected to participate in class. This involves coming to class prepared and playing an active role in discussions.

Take-Home Assignments (25%):

Each of weeks 1-7. Submit memos via Canvas.

- For weeks 1 - 6 memos are due **Friday** (so Jan 11, 18, 25, Feb 1, 8, 15) your 1-2 page memo (500 words suggested) will describe one of the following:
  - *Critical. Find an article in the popular press, a startup or a policy proposal that illustrates one of the pitfalls we discussed in class - or is somewhat related to something discussed in class. You are looking for something incomplete, sensationalistic or just plain wrong/misleading. Describe what it is and the error.*
  - Generative. Find an interesting new AI application e.g. “I could apply AI to solve...” Again ideally you’d build on the material you are learning in class.
  - Feel free to use bullet points and lists. We reward clarity. We suggest writing 1000 words and then editing out verbage to reach 500 words: brevity is meant to force clarity, not to reduce content. Each week we will pick the best of the memos to feature in class. These chosen memos will be given extra credit.

- For week 7, memos are due by **Thursday** (February 21): These memos describe potential candidates for your final project and are written jointly, one memo per team. The memo should include a long list (> 3) of candidates; as well starring 1-3 candidates for a shortlist. We will read these before office hours so each idea should be described in enough detail that office hours can be used to sift through the candidates rather than describe them. The better your memos, the more useful will be class time.

Checklists (20%):

Submitted via Canvas **Wednesday March 20.** You will be given a sequence of blank checklists on the first day of class. Throughout the lectures we will ‘populate’ these checklists. It is your job to take the suggestions in class (and whatever additional ideas you have that are inspired by the class discussion) and populate these checklists. We suggest making them initially very long, and then to winnow down. On March 2 you will submit your final checklists. Written in your own words checklists should ideally make practical the lessons of the class; and allow you to use them in your career going forward.

Final Project (40%):

Working in teams of 5-6 you will develop a project of your own - from conception to pitch. The project will be either a startup, a policy project or an internal project for a corporate. The outputs will be:

1. A poster to be displayed on **Monday March 4** in class;
2. A pitch deck due **Monday March 11** (all teams will submit and some will be chosen, based on their poster, to present to class on March 11); and
3. A longer write up (a business plan, a policy memo or a strategy memo) that is due on **Wednesday March 21.**

Objectives and Goals

This course aims to equip students with a functional, rather than a mechanical, understanding of AI. Through a series of interactive lessons students will develop a better intuition for AI applications, helping students become expert at identifying ideal use-cases and thereby well-placed to create new products, businesses and policies that use artificial intelligence.

Our goal is to make students smarter consumers of AI. The class is intended as a complement to, not substitute for, standard machine learning classes that focus on the nuts and bolts of how to be an algorithm *producer*. But there is a distinct set of intuitions that are under-developed in standard machine learning classes, and perhaps under-appreciated even by people who engage in industrial-strength machine learning as their profession. We hope that by the end of the quarter each of you will be positioned to tell Netflix, with their giant team of data scientists, something that they currently only dimly understand -- or if they understand it, they and their C-suite executives do not fully appreciate how critical it is to the company's entire future as a business (or similarly for, say, the office of a big-city chief information officer).

More generally by the end of the course, students will have:

- A framework for thinking about and finding new opportunities to apply AI
- A series of "checklists" to help think through key questions like is this a feasible project involving AI, or should I worry about bias in the algorithm in this application
- An understanding of signals, or useful data which exists in the environment, that people don't typically know to look for
- Both an "entrepreneurial" as well as a "systems-wide" way of thinking about AI

Students should also know that this course will NOT teach:

- How to *build* AI systems
- How to program code
- How to estimate data models

- The mathematical formalisms behind AI or machine learning algorithms

The class, therefore, does not require a background in programming. It does, however, require enough understanding of key statistical concepts such as mean, variance, sampling and correlation. The course is set up to avoid heavy mathematical notation and formalism; thus imposing a heavy demand for critical thinking.