ORIE 6170: Engineering Societal Systems

Lecture 1: Introduction

Nikhil Garg

Course webpage: https://orie6170.github.io/Spring_2022/

Plan for today

- Content overview
- Syllabus/Course structure
- Questions
- Hopefully end a few minutes early for specific questions

Please interrupt with questions at anytime

(but raise your hand via zoom)

Who am I?

Instructor: Nikhil Garg

Asst Professor, Cornell Tech, ORIE

Research on the application of algorithms, data science, and mechanism design to the study of democracy, markets, & societal systems

Past experiences/collabs: Uber, Upwork, other marketplaces, campaign data science, NYC Parks Department

Content overview

What is this class about?

- Every societal system is “designed” to some extent, either implicitly or explicitly
  - What are the rules of the game?
  - What can people do?
  - Who gets what, at what cost?
  - How do people find each other?
- There is a large toolkit in the intersection of computer science, economics, and operations to *understand* and *engineer* such systems

  Both theoretical and empirical

Who is this course for?

PhD students across ORIE/CS/IS/Econ/CAM/Business

The objectives of the class are: tldr – learn to do research in this area

- Introduce and discuss modeling and methodological tools used in the related literatures that are helpful in studying societal systems
- Expose students to the recent developments and state-of-the-art research in the application domains
- Develop students' abilities to understand and critique research papers and presentations and to conduct original research.

Organization: likely topics

Introduction (~1 weeks)

Transportation systems as a representative case study (~2-3 weeks)

Online marketplaces more generally (~2 weeks)

Crowdsourcing, social choice, information design (~2 weeks)

Education systems (~2 weeks)

Limits of technical approaches (~1 week)

Miscellaneous methodologies and applications

What is “Engineering societal systems?”

What is market design and the contribution of engineers in designing societal systems? What are the main ideas?

- What are the common methodologies?
- Some history and success stories
- What tools do people use?

Transportation systems as a case study

Transportation marketplaces (Uber, Lyft, etc)

- Pricing, matching, wages

Public transportation

- School bus routing and stop placement
- Shared vehicles and transportation

Congestion pricing

- Pricing usage of roadways to raise money, reduce congestion/pollution

Online markets more broadly

- Pricing and wages (centralized and decentralized)
- Reputation systems and opinion dynamics
- Recommendations and matching, assortment optimization
- Discrimination

Education systems

- School choice (matching + recommendations)
- School zone design/optimization
- Designing admissions processes

Crowdsourcing, social choice, information design

How do we {make decisions with, learn from} groups of people?

- Wisdom of crowds, herding, information design
- Voting in complicated spaces (rankings, participatory budgeting)
- Optimization + voting (gerrymandering, sortition)
- Social choice + machine learning/participatory design of ML models

Limits of technical approaches

- What are the limits to engineering methodologies?
- What (and who) is missed when we try to mathematize/optimize societal problems?
- How do we incorporate qualitative methods?
- What are the major criticisms made by others of market designers?

Cross-cutting methodology

Questions you need to answer

- What is your [the system's] *lever?*
- What is your *objective function?*
- How do people *react* to your lever?
  - What are people's *preferences?*
  - What are people's *strategy spaces?*
  - How to model people's behavior? Rational…?
- How do people affect *each other?*
- What is the information space?
  - What do you know? What data do you have?
  - What do people know?
  - How do you acquire more information?

Common tasks

- Understand your domain
- Write a model for the ?s on the left
- Calculate "equilibria"
- Estimate preferences from historical data
- Simulate counter-factual worlds
- Experiment/Psuedo-experiment
- Deploy a system

Methods used

- Applied modeling/stochastics
- Game theory/mechanism design
- Optimization, Algorithms
- Machine learning/statistics/data science
- Online learning/decision-making

Course themes

Be able to articulate what matters in a system

- “All models are wrong”, “The map is not the territory”
- Why did the authors include/exclude certain things in their model? What would change if they made different choices?

Different questions require different methods

- Sometimes theory, sometimes empirical
- Often a mix: how do we do research at the intersection?
- Why did the authors choose the methods they did? What would the paper look like if it was a theory/empirical paper instead?

What is this class not?

This is not an algorithmic game theory class, or even a mechanism design class

Tim Roughgarden: Algorithmic Game Theory (Lecture 1: Introduction and Examples) – YouTube

We won't cover details of auctions, Gale Shapley matching

It is also not a machine learning or optimization or a "methods" course

- We're not going to go deep on any particular method
- The papers of course use (advanced) methods; we will discuss them; I will provide further resources; and in your paper reviews you will go deep on understanding a paper's methods

Syllabus

https://docs.google.com/document/d/1jb3OyBg9Iv5YsSgp0E1dRVmNpvfK_1-Z5e1fazlm7Uk/edit

Assignments + Grading

Final project: 35%

Project proposal/presentation, report, class presentation, peer review of a classmate's project

Paper review + presentation: 25%

Read a paper and write a journal-level review for it [suggested list posted soon]

Give a 10-15 minute presentation to the class on your chosen paper

Presentation feedback: Watch two presentations and give feedback

A paper review presentation by a classmate

A presentation available online by an established researcher

Paper reading and discussion: 15%

Choose 2 papers that we'll discuss in class and be a discussion leader [list posted on rolling basis]

Participation: 10%

(Remote) attendance mandatory. Each class that you miss, please email me a 300-word reflection on the material for that day. Missing more than 2 classes without a (good) excuse may affect your grade.

I don't expect everyone to read each discussed paper in detail, but I expect you to read abstract + intro for almost all of them, read the model/setup for most, read main results and methods for many, and go into proof/method details for a few.

Class structure

~10-15 days discussing papers

~3-5 lectures by me

~5 guest lectures

~3 days paper review presentations

~1 day project proposal presentations

~2-3 days project presentations

Course communication

Course Slack channel: First place for any question/comment

Office hours: Happy to chat about anything – sign up on link in syllabus

Email: Try to avoid; but preferred over private message on Slack.

Classroom norms

- Take space, make space: allow others to join the conversation, but please contribute as you feel comfortable.
- Embrace a growth mindset. Not understanding something in a paper is the default.
- Ask questions!
- Be willing to give and receive feedback respectfully.
- Zoom norms

  Feel free to take video-off breaks as necessary, and a couple lectures of video off the entire time. But I expect you to mostly keep video on and participate.

Announcements

- Watch out for the course pre-survey, posted on Slack soon

Next time: “Great ideas and papers” relevant to market design

- Please read (some of): The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics - Roth - 2002 - Econometrica

Questions?