ORIE 6170: Engineering Societal Systems

Lecture 1: Introduction

Nikhil Garg

Course webpage: https://orie6170.github.io/Spring2025/

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)

Please complete the following poll, logged in with your cornell Netid:

PollEv.com/nikhilgarg713

Who am I?

Instructor: Nikhil Garg

Asst Professor, Cornell Tech, ORIE

Research computational understanding of and data-driven decision-making within societal systems, methodologically spanning OR, CS, policy, economics

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?
  - How do we allocate scarce resources?
- There is a large toolkit in the union of computer science, social sciences, operations, humanities to *understand* and *engineer* such systems: mathematically, computationally, empirically, qualitatively

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.

What is “Engineering societal systems?”

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

- What are the common methodologies?
- Some history and success stories, as well as cautionary tales
- What tools do people use?

(Partial) pipeline of a computational system

Most (even fairness/social good) work is on predictive modeling

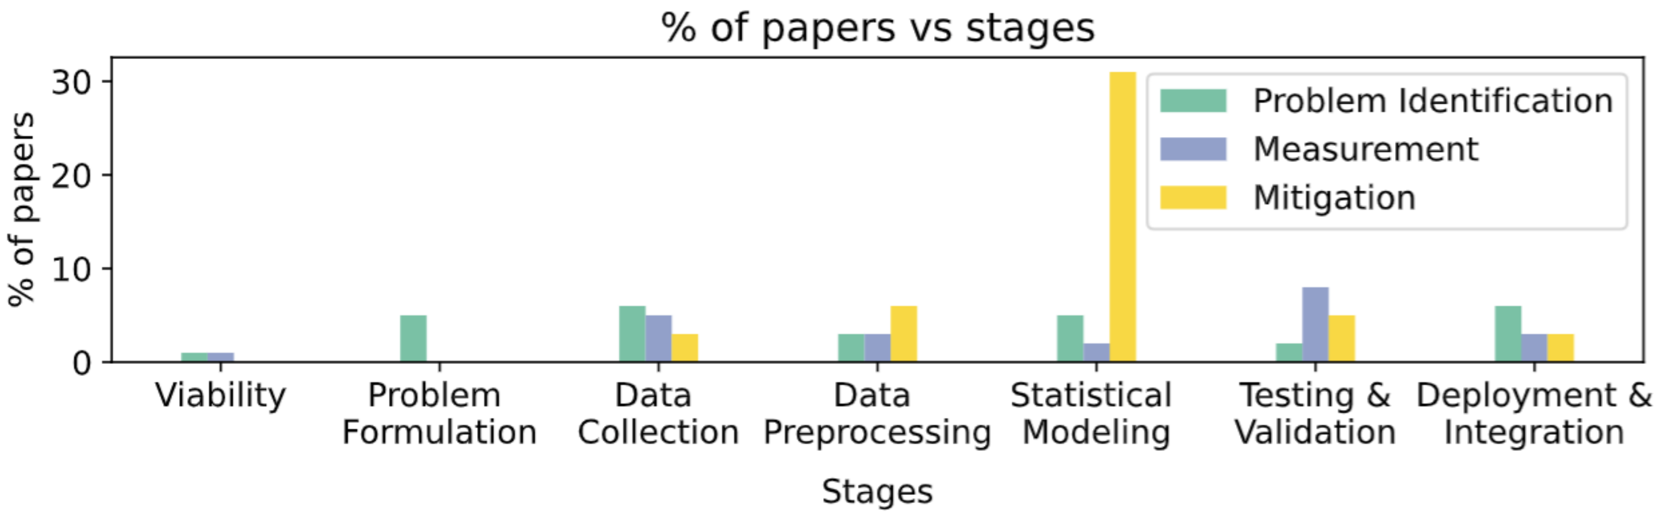

Black, Naidu, Ghani, Rodolfa, Ho, Heidari (EAAMO `23)

Interlocking questions at every stage

What are we predicting? Is that the right target?
What data do we have for prediction?
How do human generated data biases affect predictions?
Will our predictions hold up over time?

How to use noisy/accurate predictions to make decisions?
How do prediction errors lead to decision errors?
How should we allocate scarce resources?
How do we make fair decisions under constraints and uncertainty?

What are the *population* effects of individual models?
What happens when the world changes?
Can people respond to your system to change what they're doing?

How do we design and evaluate real-world systems and policies?

Can we measure the thing we care about?

Is the computational system better than the human alternative? Under what metric?

Organization: likely topics

- Introduction
- Prediction
- Optimization + resource allocation
- Matching & recommendations; market design broadly
- Pricing
- Miscellaneous methodologies and applications
  - Generative AI, human-AI collaboration, experimentation + evaluation, etc

Special focus this year

- algorithms in government
- Limitations and pitfalls, but also cautious “success” stories

Prediction

Machine learning/predictive systems that aim to predict future outcomes to help make decisions, broadly in use in health care, criminal justice, education.

- Aspects of “successful” deployed predictive systems
- “Against predictive optimization”; pitfalls in deploying predictive systems
- Comparing human and machine predictions

Case studies/papers

- Risk assessments for bail decisions in criminal justice
- Prediction models in child welfare
- Student educational “tracking” and predicting educational attainment

Optimization + resource allocation

Optimization and scheduling approaches used for large scale system design such as placement of public transportation and bike share systems, school bus routing, emergency management, and others. Resource allocation systems broadly such as in social services, informed by prediction models.

Case studies/papers

- Vehicle/school bus routing, for equity and efficiency constraints; Emergency vehicle routing and allocation
- Resource allocation in social services, e.g., for housing and other services
  - Resource allocation over space and time
  - Inspections and maintenance resources
- Gerrymandering/optimization for democracy

Ranking/Matching/recommendation/market design

Stable matching systems or algorithmic recommendations to help match people to each other, such as for residencies, schools, jobs, dating, etc. Market design more broadly. Ranking design for social network feeds

Case studies/papers

- School choice (matching + recommendations)
- Ranking design for social network feeds (Bridging/community notes, personalization)
- Refugee matching

Pricing

Pricing/monetary decisions by government or regulated by government, including congestion pricing, taxation auditing, regulation of algorithmic pricing collusion, etc.

Case studies/papers

- Congestion pricing
- Wireless spectrum auctions
- Tax auditing
- Algorithmic pricing collusion
- Personalized Pricing and regulation/discrimination broadly

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?
- Can we meaningfully address "wicked problems" through technology?
- How do we bring "systems level thinking" to these systems

Miscellaneous

- Algorithms in NYC government as an extended case study
- Experimentation + evaluation
- Human-AI collaboration
- Generative AI in high-stakes settings
- Limits of technical approaches
- Voting and social choice
- Other applications, such as online marketplaces

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 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?
  - Will this change over time?
- *What happens when your system is wrong?*

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
- Evaluate as close to "real world" as possible
- Deploy a system

Methods used

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

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?
  What data do people have? What data is needed to answer this question?

Different questions require different methods

- Sometimes theory, sometimes empirical, sometimes qualitative
- 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/qualitative 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 in lecture
- 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
  - You are welcome to focus your assignments on papers whose methods you find most appropriate for your goals

Syllabus

ORIE6170 syllabus Spring 2025 - Google Docs

Assignments + Grading

Final project: 40%

- Project proposal/presentation, report, class presentation, peer review of a classmate's project
- Option of participating in a joint review research paper with a larger class group

Paper review + presentation: 20%

- 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: 10%

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: 10%

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

Attendance + Participation: 10% **Each**

- Attendance mandatory. **We will take attendance in many of the classes.**
- **Participation is also mandatory.** 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. **Especially important for those of you in Ithaca.**

Class structure

~10-15 days discussing papers

~3-5 lectures by me

~3-5 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~~

Ed Discussion (linked from Canvas)

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

Paper discussion signup + instructions out soon

Next time: "Great ideas and papers" relevant to the area

Please read (some of):

- The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics - Roth - 2002 – Econometrica
- Rittel and Webber. "Dilemmas in a General Theory of Planning" sympoetic.net/Managing_Complexity/complexity_files/1973 Rittel and Webber Wicked Problems.pdf

Questions?