Online Decision-Making and Market Design

ORIE 6180: Syllabus

Spring 2026

Essential Course information:

Lectures and Recitations
Class time/location: Tue/Thur 11:40am-12:55pm, Tang Hall 205

Instructor
Sid Banerjee <`sbanerjee@cornell.edu`>, 229 Rhodes Hall

Zoom Link
`https://cornell.zoom.us/j/93996106771?pwd=vXG8vjDMt63aTaRh5KevYN3AsV6OH4.1`
(Meeting ID: 939 9610 6771, Passcode: samuelcahn)

Course description:

TLDR: *How does one solve optimization problems without knowing all the inputs?*

In this class, we will study models and algorithms for solving two types of problems:

- **Online decision-making**: Settings involving multiple decisions over time, with uncertainty about future events, and where actions affect current and future outcomes.
- **Market design**: Settings involving multiple decision-making agents, with uncertainty about agents' types, and where agents' actions affect each other as well as overall outcomes.

The above problems have a long history across multiple fields (Operations Research, Control, Computer Science, Economics). The aim of this course is to survey the main paradigms for these problems – how to formulate them; what are common algorithmic tools for solving them; how to analyze the performance and limits of these algorithms; how to compare different models of knowledge in terms of their effect on decision-making.

Why study these problems together? From an *optimization viewpoint*, these two problems turn out to be very closely related. We will start by discussing how to formulate these problems as optimization programs (insulating variables, revelation principle, minimax techniques), and common approaches towards solving them (LP/convex duality, iterative methods, potential functions). Our approach will be theoretical, focusing on mathematical techniques for designing and analyzing algorithms with formal guarantees. However the problems we consider have significant practical motivation, and we will see lots of examples throughout the course.

Course Goals

1. Models and Tools: Introduce important models and algorithmic tools from online algorithms, control, microeconomics and game theory which are useful for reasoning about online markets.
2. Research Problems: Identify open problems in these topics.
3. Applications: Exposure to interesting platforms, and the use of these techniques in practice.

Tentative Course Structure

The course is geared towards giving an overview of the landscape of online decision-making and markets, and prepare students for research on these topics. The following is a superset of topics; in the spirit of the course topics, we will dynamically adjust the exact mix.

1. **From Decision-making Models to Optimization Programs**
   - Markov Decision Processes and stochastic control
     – State-action frequencies, value functions, HJB equations
     – Examples: value and policy iteration
   - Mechanism design
     – Incentive compatibility and rationality constraints; the revelation principle
     – Implementability and Border's criterion
     – Examples: Walrasian prices, rationalizability (Afriat's theorem), Bayesian persuasion
   - Non-Bayesian models for online decision-making
     – Zero-sum games and the minimax theorem, Yao's lemma
2. **Markov Decision Processes: Exact Solutions**
   - Structural characterizations and exact solutions
     – Bang-bang control, LQR and linear policies
     – Threshold policies (Value function convexity, optimal stopping)
     – Index policies (the Gittins' index, prevailing-cost arguments, polymatroids)
3. **Mechanism Design: Exact Solutions**
   - Basics of mechanism design
     – Welfare maximization and VCG
     – Single parameter settings: Myerson's lemma and optimal mechanisms
     – Impossibility theorems (bilateral trade and Myerson-Satterthwaite, public goods)
   - Mechanism design in complex settings
     – Correlated valuations (Cremer-McLean, Milgrom-Weber)
     – Multi-item valuation classes (unit-demand, submodular, XOS)
   - Beyond welfare maximization
     – Fairness in allocation (Varian's characterization)
     – Non-monetary approaches (stable matching, cost sharing)
4. **Approximation in MDPs and Mechanism design**
   - Additive approximation in MDPs – LP relaxations and weakly-coupled MDPs
     – The compensated coupling
     – Finite-horizon bandits: UCB and Thompson sampling, 'good events' analysis, information theoretic lower bounds
     – Reinforcement learning via 'optimistic' algorithms
   - Approximate mechanism design
     – Bayesian settings (prophet inequalities, black-box reductions)
     – Worst-case settings (sample-based pricing, Bulow-Klemperer)
     – Flow dualities for approximate mechanism design
     – Multi-objective settings and the theory of majorization

4. **Non-Bayesian Paradigms for Online Decision-Making**
    - Prediction from expert advice
      – Multiplicative weights and applications, Follow-the-perturbed-leader
      – Blackbox reductions via importance sampling: partial feedback, bandits
    - Competitive analysis for online algorithms
      – Canonical problems: Rent-or-buy, paging and k-server, online matching
      – Competitive analysis via dual fitting and primal-dual approaches
      – Online contention resolution schemes
    - Algorithms with predictions

Prerequisites:

I will assume knowledge of basic probability and algorithms/optimization (ideally at the level of ORIE 6500/CS 6820 and ORIE 6300 or equivalent) – in particular, you should be comfortable with (or willing to read up) LP duality, basic convex optimization, Markov chains, coupling, concentration inequalities. Prior exposure to game theory would be helpful, but is not necessary. Send me a mail if you are concerned about having the appropriate prerequisites.

Course Logistics:

Your grade will be based on a project (40%), assignments (30%) and scribing/class participation (30%). Scribing guidelines will be given in the first class.

The main component of the grade is based on the final project. For this, students need to submit a 1-2 page proposal on **Wednesday, November 3rd, 2021**. Subsequently, we will have student presentations in the last week of classes, and a **final report due during the finals period**. The proposal is worth 5%, the final presentation is worth 15%, and the remaining 20% is for the report.

References:

There is no single textbook for the course; I will post my notes, as well as the scribed notes, for all the topics we cover. A lot of what we discuss will be drawn from papers/tutorials, which will be linked on the website. However, you may find some of the following references helpful:

- The videos from the Simons Institute semester on Data-Driven Decision Processes
- Bayesian Decision-Making settings:
  – Richard Weber's course notes: Optimization and Control
  – Kamesh Munagala's notes on Optimization Under Uncertainty
- Non-Bayesian Decision-Making:
  – Sahil Singla's course on Advanced Algorithms and Uncertainty
  – Thomas Kesselheim's course on Algorithms and Uncertainty
  – Bobby Kleinberg's course notes for Learning, Games and Electronic Markets
- For Mechanism and Market Design:
  – Mechanism Design: A Linear Programming Approach by Rakesh Vohra
  – Mechanism Design and Approximation by Jason Hartline
  – Tim Roughgarden's lecture notes; also available as a book
  – Online and Matching-Based Market Design - a collection of essays summarizing the contents of the Simons Institute semester on the same topic.