Online Decision-Making and Market Design

ORIE 6180: Syllabus

Fall 2021

Essential Course information:

Lectures and Recitations
Class time/location: TR 9:40-10:55am, Phillips 307
Website: `http://people.orie.cornell.edu/sbanerjee/ORIE6180/orie6180f21.html`

Instructor
Sid Banerjee <`sbanerjee@cornell.edu`>, 229 Rhodes Hall (Office hours by appointment)

Course description:

TLDR: *How does one make 'optimal' decisions without knowing all the inputs?*

Online platforms and marketplaces are ubiquitous in most aspects of our lives. At their core, these platforms are engines 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.

To do so, we will adopt an *optimization viewpoint*, discussing common principles for formulating these settings as large-scale optimization programs (insulating variables, revelation principle, minimax techniques), and common approaches towards solving them (LP/convex duality, iterative methods, potential/Lyapunov functions). Our approach will be primarily 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 problem settings and platforms, and empirical and anecdotal evidence for 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/Potryagin maximum principle
      – Examples: value and policy iterations
    - Non-Bayesian models for online decision-making
      – Zero-sum games and the minimax theorem, Yao's lemma
      – vector-valued games and Blackwell approachability
    - Mechanism design
      – Incentive compatibility and rationality constraints; the revelation principle
      – Examples: Walrasian prices, rationalizability (Afriat's theorem), Bayesian persuasion)
2. **Markov Decision Processes**
    - Structural characterizations and exact solutions
      – Threshold policies (Value function convexity, Bruss' Odds Ratio for optimal stopping)
      – Index policies (mthe Gittins' index, prevailing-cost arguments, polymatroids)
    - Approximation techniques for MDPs
      – LP relaxations and weakly-coupled MDPs
      – Hardness bounds via Martingale relaxations
      – The compensated coupling
      – Finite-horizon bandits: UCB and Thompson sampling, the 'good events' analysis, information theoretic lower bounds
      – Reinforcement learning via 'optimistic' algorithms
3. **Mechanism Design and Markets**
    - 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)
      – Implementability and Border's criterion
      – Flow dualities for approximate mechanism design
    - Approximate mechanism design
      – Bayesian settings (prophet inequalities, black-box reductions)
      – Worst-case settings (sample-based pricing, Bulow-Klemperer)
    - Other strategic models
      – Two-sided market models (Rochet-Tirole, Armstrong, Weyl)
      – Fairness in allocation (Varian's characterization)

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

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:

- Bayesian Decision-Making settings:
  - Richard Weber's course notes: Optimization and Control
  - Kamesh Munagala's notes on Optimization Under Uncertainty
- Non-Bayesian Decision-Making:
  - Bobby Kleinberg's course notes for Learning, Games and Electronic Markets
  - Alex Slivkins' book on Introduction to Multi-Armed Bandits
- 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