Introduction to Algorithmic Economics

Spring 2023

Instructors: Hongning Wang, Denis Nekipelov

Prerequisites: Computer Science majors should register for the CS 4501 course, which has CS 2150 as a prerequisite. Economics majors should register for the ECON 4559 course, which has ECON 3010 or 3110 and ECON 3720 or 4720 as a prerequisite.

Communication: Course materials will be posted on the designated Microsoft Teams group. We will use the Teams for most communications with the class, and encourage you to use it for asking questions and discussions with the course staff and your classmates.

Course description: Algorithmic tools are becoming commonplace tools in large marketplace to help market agents make their decisions or even make decisions on their behalf. At the core of those tools are learning algorithms that take the data from the market environment and update the recommendation or decision with the arrival of each new datum. The automated nature of such algorithms allows the tools to be more agile and make decisions faster and more precisely than humans could. Even more importantly, they can make those decisions in very complex multi-dimensional market environments.

For instance, online marketplaces like Amazon or Walmart offer a vast variety of items that are sold by thousands of vendors (including Amazon and Walmart themselves) with the product selection that far exceeds any brick and mortar grocery, convenience or furniture store. Individual vendors compete with each other over the price and variety of products they offer so that the platform can display their products to consumers who search for them. Given the number of competitors and the selection of products that can serve as substitute to the product offered by a particular vendor, the majority of vendors rely on AI-based software that adjusts prices for their products automatically given their inventory and behavior of competitors. As a result, prices in those marketplaces change dynamically, responding to the changes in the overall consumer demand, stock-outs of individual vendors and the dynamics in the markets of complementary or substitute products.

There are two fundamental differences between the framework that is discussed in standard Economics courses to analyze markets and the reality of these modern marketplaces. First, the agents in such marketplaces rely on fully automated tools to make their decisions that use relatively simple directions from human operators. In contrast, classic Economic analysis relies on the idea that actions of market agents are a direct expression of their preferences over market outcomes (such as prices and quantities). With automated tools agents' actions no longer express their preferences directly. Instead they reflect how a learning algorithm updates those actions to adjust to changing market environment in order to optimize for the goal specified by a human operator.

Second, these marketplaces are inherently unstable. Learning algorithms constantly update the actions and the updates of one algorithmic tool triggers the updates of the rest of the tools. At the same time, Economic theory instructs the application of the concept of market (Nash) equilibrium. Market equilibrium has stability in its core: if the equilibrium has been achieved, no agent in the market has an incentive to deviate from their current action.

The course will cover recent work in Computer Science and Economics the enables the appropriate analysis of these marketplaces and allows us to understand how the agents in these marketplaces can successfully compete with each other.

Course contents: We will discuss several fundamental concepts that enable development of learning algorithms and the analysis of non-stable environments. The list of topics in order is provided below.

- Information and competition; Economic behavior with private information; Bayesian updating and Bayesian learning
- Introduction in learning and machine learning
- PAC learnability
- Complexity and bias in learning problems: Computational complexity of learning
- Convex learning problems and computational stability; Regularization and regularized learning; Gradient and Stochastic gradient descent algorithms
- Online learning; Distinction between full information problems and bandit problems; The concept of regret; Online classification and online convex optimization
- Online learning in multi-arm bandit environments;
- Stability of environments with online learning agents; Nash equilibrium and Coarse Correlated equilibrium; Learning algorithms and implementation of an equilibrium;
- Analysis of outcomes in non-stable environments; Price of anarchy bounds; Smoothness and other approaches for price of anarchy evaluation; From price of Anarchy to empirical price of anarchy;
- System with long-run feedback; Dynamic optimization and dynamic Markov games with incomplete information;
- Reinforcement learning; Model-based and model-free algorithms for reinforcement learning; Reinforcement learning algorithms as strategies in dynamic Markov games.

Assignments: The main assignments for the class will be a series of projects where students will work in small teams to analyze data from real electronic platforms (e.g.,Yahoo!'s Webscope data and Amazon AWS open data). For most assignments, students will be grouped into teams of two or four students, including students in both the CS and ECON sections of the course in each team. The project teams and assignments will emulate the work of real interdisciplinary teams at leading companies. It will be the team's responsibility to distribute the tasks, communicate with each other and explain in the project report what part of the project was completed by what team member. For the final project, students will have an opportunity to define their own open-ended project to answer an interesting question about the world using methods and techniques from the class. In addition to the project assignments, there will be reading assignments and other preparation assignments for classes, and there may be some individual problem sets or quizzes

(we hope that these will not be necessary, but will use them if it seems like they will be helpful for improving engagement in the class, or if there are students who are not contributing fully to team assignments).

Projects: There will be a total of 5 projects that involve algorithm development, simulations and data analysis related to the topics of the class.

Evaluation: Grading will be based primarily on your performance on project assignments, with adjustments for outstanding contributions to the class. For team projects, everyone in a team will normally receive the same grade, except in situations where there are problems or where not all team members contribute fully to the project. All students are expected to be present at the final project proposal and final project presentation.

Grades will be calculated with several different weightings, where your grade is based primarily on whichever weighting results in the highest score. The range of possible weightings is:
– Project 1: 0-5%
– Project 2: 5-15%
– Project 3: 10-15%
– Project 4: 10-15%
– Project 5: 10-15%
– Final Project: 20-65%
– Class Contribution: 0-25%

Spend your energy focusing on what you are learning, instead of worrying about your grade.
Honor: We believe strongly in the value of a community of trust, and expect all of the students in this class to contribute to strengthening and enhancing that community.

As a student at the University of Virginia, you are trusted to be honorable and expected to behave in ways that merit that trust. We take advantage of this trust to provide a better learning environment for everyone. The course will be better for everyone if everyone can assume everyone else is trustworthy, and we start from the assumption that all students at the university deserve to be trusted. For most assignments in this course, you will be encouraged to discuss ideas and work with others to develop your ideas, and often expected to work in a team. We expect you to be honest, fair, and respectful with your teammates, and to contribute fully to your team to the best of your ability. We also expect you to notify the course staff if there are any problems with your team, or if you have teammates that are not behaving honorably.

You will always be expected to credit any collaborators and properly cite any resources you use. The honor expectations for each assignment should be clearly stated and make it unambiguous what is and is not permitted. If it is ever unclear what is considered acceptable on an assignment, please check with the course staff.

Expectations and Accommodations: If you anticipate any issues related to the format, materials, or requirements of this course, please meet with one of the instructors outside of class so we can explore potential options. Students with disabilities may also wish to work with the

Student Disability Access Center to discuss a range of options to removing barriers in this course, including official accommodations. Please visit their website for information on this process and to apply for services online: sdac.studenthealth.virginia.edu.