Fall 2017 – EE 290-O / IEOR 290
Societal-Scale Cyber-Physical Systems: Machine Learning and the Internet of Things
Course Syllabus

Course description

This course will cover theoretical tools for the analysis of data and human agents in cyberphysical systems. Concepts will include optimization, game theory, differential privacy, behavioral methods, statistical estimation, information theory, and utility function learning, with a focus on applications in resilience, security, privacy, and markets of data exchange. Throughout, we will emphasize the underlying mathematical tools required to understand the current research in each of these fields.

Prerequisites

Students should have some familiarity with linear algebra, optimization, and statistics.

Administrivia

- **Term:** Fall 2017.
- **Units:** 3.
- **Instructor:** S. Shankar Sastry and Roy Dong.
- **Course time:** TuTh 1100-1230.
- **Course room:** 521 Cory.
- **Website:** The course website is the best place to get the most up-to-date information about the course's schedule, current homeworks, &c. It can be found at:
  https://people.eecs.berkeley.edu/~roydong/2017fa_ee290o_ieor290.html
  Or, in brief: http://bit.ly/2wyYQ48
- **Office hours:** See website.

Course structure

The structure of this course is lectures presenting the mathematical framework and tools used for analysis in a variety of fields. Afterwards, the students will present recent papers in these fields showing how these tools are employed in cutting-edge research in cyber-physical systems and the Internet of Things. The focus of the presentations will be on: 1) what models and assumptions the authors use, 2) what are the consequences and conclusions of these assumptions, 3) if applicable, what real-world data is used, 4) what question are the researchers trying to answer. The goal is to both understand the mathematics behind different fields, but also the philosophy: what part of this problem does this community care about most?

There will be a participation grade, in part based on these student presentations, as well as a few homeworks throughout the semester to ensure everyone understands the mathematics. There will be a final project, which can either be research related to this course's theme or an overview of a general problem in a field, e.g. privacy. In either case, there will be presentations at the end of the semester.

Please refer to the course website regularly for both announcements, updates to the schedule, and deliverables requested throughout the semester.

Grading

Homework	10%
Participation	50%
Final report	20%
Final presentation	20%

Extended course description

The Internet of Things (IoT) is a term that represents a huge technological trend that is taking place: almost every device is being imbued with the intelligence of a microprocessor and an Internet connection. The interconnection in IoT promises an infrastructure that can drastically change how consumers live their day-to-day lives, with huge gains in efficiency, value, and possibility due to the shared knowledge and autonomy allowed. In profound ways, as the technology develops, the modalities of existence people experience will grow and shift.

However, the scale and scope of IoT raises new problems for engineers to consider. These problems are significantly different from ones previously explored in the design of comparatively isolated systems, and require a new theoretical underpinning to analyze IoT with models that capture all salient facets of these new technologies. This textbook contains a handful of theoretical frameworks, and their applications, as a first step into this new research frontier.

The goal of this course is to cover the theoretical foundations for the analysis of such systems.

This course will cover some of the new problems that arise due to: a) the scale, complexity, and homogeneity of multiple devices, b) the interactions between new service models and human agents, c) the new vulnerabilities that arise due to new interconnections, and d) the structure of these new disruptive markets. In the process, we'll introduce the theoretical background underpinning many of the different formulations of these novel problems; researchers from many different fields have approached different facets of the new problems in cyberphysical systems, and the goal of this course is to provide the student with enough resources and context to be able to understand the cutting-edge research in several of these fields.

One of the main focuses of this course will be the study of the role of information in these cyberphysical systems. First, we must consider what data is transmitted. Then, we ask ourselves the legitimate value of the data, i.e. how can data be leveraged for effective system operation? Once this data is being used in closed-loop system operation, be it through classical control methods, incentive schemes, or some other new service model, we have to analyze the effect of data in closed-loop. Particularly, once the data is used in system operation, do data sources have incentives to modify or obfuscate their data? What statistical information is contained in the data and does it raise privacy concerns?

Another main focus will be the interaction between humans and the system. We will discuss different methods for modeling humans and how we can learn parameters of these models from data.

This course will **not** cover how to design physical devices or the protocols for interfacing different devices. This course will not deeply treat machine learning qua machine learning, but will touch on various aspects of machine learning as it relates to these new service models. For example, we will consider how machine learning must change when data is not drawn from a distribution but, rather, is the reported values of strategic agents.