ESE 680-004: Learning and Control

Instructor: Nikolai Matni, Assistant Professor, ESE Department
Email: nmatni@seas.upenn.edu
Course website: https://nikolaimatni.github.io/courses/ese680-fall2019/
Office hours: Levine 374, please email instructor to schedule a mutually convenient time

Course Outline

This advanced topics course will provide students with an introduction to current areas of research at the intersection of machine learning and control. We will investigate machine learning and data-driven algorithms that interact with the physical world. Topics of study will include learning models of dynamical systems, using these models to robustly meet performance objectives, optimally refining models to improve performance, and verifying the safety of machine learning enabled control systems. The course will be a combination of lectures and student led presentations of papers drawn from a list of both classical and modern texts. Students will be evaluated based on their paper presentation, as well as a class project. Suitable choices for projects include implementing, evaluating, and comparing tools introduced in class, extending existing theoretical results, and applying tools to a domain specific problem of their choosing.

Tentative List of Topics:

- **System identification**
  - asymptotic vs finite-time guarantees
  - full state vs. partially observed systems
  - system identification for control

- **Learning theory**
  - Risk and empirical risk minimization
  - Concentration inequalities
  - Uniform convergence
  - Stability and generalization

- **Control of uncertain systems**
  - Modeling assumptions, system norms, robust stability
  - Optimal, robust, and model predictive control
  - Model validation and robust control

- **Model based control of learned systems**
  - Metrics for learning and control problems: regret, PAC, and beyond
  - End-to-end sample complexity guarantees
  - Comparison of online approaches

- **Model free methods**
  - Approximate dynamic programming and reinforcement learning
  - Sample complexity guarantees for continuous control
  - Complexity gaps between model free and model based methods

- **Safe learning and control**
  - Robust invariant sets and model predictive control
  - Lyapunov functions and regions of attraction
  - Control Lyapunov and control barrier functions

Grading

Grading will be based on course participation, paper presentation, lecture note scribing, and the course project.

Semester project: Projects may be done individually or in groups of two. Students are encouraged (but not required) to propose a topic that connects class material to aspects of their research. The first component of the project will consist of a written proposal (2 page limit, excluding references). This proposal should include: motivation, prior work, problem statement, and ideas/methods/approach to solving the problem, and will be due approximately 1 month into the semester. The second component will be a midterm update (3 page limit, excluding references), and will be due around the end of October. This update should include a description of current progress, which may include preliminary results (either computational or theoretical), a description of attempted methods/approaches, current roadblocks, and an outline of your plans for the rest of the semester. The final component of the project will be an end-of-semester report (8 page limit, excluding references) and accompanying presentation. The final report is meant to mimic a conference setting, and therefore the report should be a self-contained document with an introduction, literature review, problem formulation, main results (and experiments if applicable), and discussion/conclusion sections. Similarly, each group will be expected to give a 20min presentation (during the last two classes), where they describe the problem their project addresses, their progress and potential innovations, and concluding remarks.

Deciding on a topic: Examples of open problems will be given during the first weeks class. Each team is **strongly encouraged** to meet with me before the proposal submission. This will give them the opportunity to explore potential projects, as well as receive feedback and suggestions. If a team decides *post proposal* to switch to another problem, they may do so long as they resubmit a new proposal.

Project grading: The goal of the project is to stimulate research and discussions between you, your peers, and the instructor. The project grade will not be based on the amount of progress towards solving an open problem. Rather, it will be a reflection of the amount of effort put into the project, and on the clarity and quality of the submitted proposal/update/report and final presentation.

Paper presentation and scribing: Depending on enrollment numbers, each student will be assigned a paper to present and/or asked to scribe a lecture.
Presentations: Presentations should address the following questions:

- Problem statement: what problem is the paper solving, and why is it meaningful?
- Prior work: how does this paper fit into the current research landscape.
- Key idea/main result: what is the main takeaway from the paper? What is the main result and what are its implications?˜
- Key technical tools: what are the main technical contributions of the paper? Be prepared to present and explain these in detail.
- Points of confusion: are there technical or conceptual arguments that are unclear or incomplete? Bring these up so that we can discuss and try to work through them in class.
- Shortcomings/areas for improvement: what are some of the flaws of the paper? What are possible directions for future work to address these?

Presenters should plan for about 45min of presentation time and 45min of discussion time. To ensure that student presentations are of high quality, students will be required to sign up to present a paper at least 2 weeks in advance, and meet with the instructor 1 week in advance to go over their presentation.
Scribing: Scribes will be expected to both summarize the assigned reading and the lecture itself, and should address all of the same points that a presentation would. It is expected that you will put several hours into these notes, as they should go beyond just summarizing what was presented in class.

Prerequisites

ESE 500: Linear Systems, and one of *CIS 520: Machine Learning* or *ESE 605: Convex Optimization*, or permission from the instructor.

This course is ideal for advanced graduate students, who are interested in applying novel research concepts to their own work. Students are expected to be mathematically mature with a solid background in linear algebra, calculus, and probability, and be familiar with basic concepts in optimization and control.