



ESE 618-001: Learning for Dynamics and Control

Instructor: Nikolai Matni (nmatni@seas.upenn.edu), Assistant Professor, Dept. of ESE,

Course website: <https://nikolaimatni.github.io/courses/ese618-fall2021/index.html>

Lectures: Tu/Th 3:30-4:45pm ET, Moore 212 and on Zoom (check Piazza for Link/Passcode)

Office hours: Tu/Th 5:00-6:00pm ET, Levine 374 and on Zoom (check Piazza for Link/Passcode)

Course description

This course will provide students an introduction to the emerging area at the intersection of machine learning, dynamics, and control. We will investigate machine learning and data-driven algorithms that interact with the physical world, with an emphasis on a holistic understanding of the interplay between concepts from control theory (e.g., feedback, stability, robustness) and machine learning (e.g., generalization, sample-complexity). 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 also expose students to the ethical considerations that need to be considered when designing learning algorithms that interact with and are placed in feedback with the world.

Prerequisites

This is an advanced theory-intensive course. A solid foundation in linear systems (at the level of ESE 500), probability theory (at the level of ESE 530), and optimization (at the level of ESE 605), as well as mathematical maturity (comfort with reading and writing proofs) is required. Familiarity with Python is helpful, but not required. Undergraduates need permission.

Intended audience

This course is ideal for advanced graduate students who are interested in applying novel research concepts to their own work. By the end of this course, students will be ready to start doing research in the Learning for Dynamics and Control (L4DC) space.

Tentative List of topics

Part 0: Introduction & Fundamentals

- **Week 1: Introduction, history, course goals, fundamentals (Markov Decision Processes, Probably Approximately Correct, Regret)**

Part 1: Linear Systems

- **Week 2: Learning to control the linear quadratic regulator**

- **Week 2-3: A fresh look at system identification**
- **Week 4-5: Using learned models for control**
- **Week 6-7: Model-free methods**
- **Week 8: Non-stochastic control and online convex optimization**
- **Week 9: Fundamental limits**

Part 2: Nonlinear Systems

- **Week 10: Learning for nonlinear systems: nonlinear stability certificates, empirical risk minimization**
- **Week 11: Generalization without distribution shift: learning stability certificates from data, learning nonlinear predictors, stability constrained policy optimization**
- **Week 12: Dealing with distribution shift in learning and control: stability constrained imitation learning, regret bounds for adaptive nonlinear control**

Part 3: Feedback, learning, and fair ML

- **Week 13-14: recommender systems and statistical discrimination**

Grading

- **Homework (60%):** there will be five homework assignments. An initial homework assignment, Homework 0, will be handed out on the first day of class, and will be worth 10%. **Homework 0 is mandatory, and must be passed to a satisfactory level:** it is used to check your knowledge of prerequisites. The remaining four homework assignments will each be worth 12.5%. The homework assignments **must be LaTeXed and submitted on Canvas via Gradescope**. Please use the latex template provided on the course website – we ask that you write out detailed and rigorous solutions. You will be given **6 free late days** which you may use as you please throughout the semester, after which no late assignments will be accepted. You are allowed, even encouraged, to work on homework in small groups, but you must write up your own homework solutions and code to hand in -- please indicate who you collaborated with on your assignments. Each homework problem will be graded on a scale of 0-4.
- **Final project (40%):** students will be expected to work on a theory-focused project (in groups of up to 2 students). See course website for more details.

Note that these weights are approximate, and we reserve the right to change them later.

Code of Academic Integrity: All students are expected to adhere to the University's [Code of Academic Integrity](#).