

Tuesdays 3:00–4:30 pm and Thursdays, 4:05–5:35 pm

Prof. Sean Meyn (use canvas for email)

Cloud

Reinforcement learning is a collection of tools for the design of decision and control algorithms. What makes RL different from traditional control is that the modelling step is avoided, and instead the control design is based on observations of the system to be controlled.

Its origins go all the way back to Claude Shannon in the 1950s ([click here for the amazing survey of Sutton and Barto](#)), and the field made headlines in the public press more recently following the success of AlphaGo and other RL algorithms that beat grand masters at complex games like Go and Chess.¹ Today it is hoped that RL will be an engine behind autonomous cars, as well as better decision making in fields ranging from medicine to finance. This course provides an introduction to RL through the lens of control theory. We will find that the DQN algorithm behind AlphaGo is related to classical control concepts going back to the 1960s. Given this intuition we will discover techniques to create new and potentially more reliable algorithms for decision and control.

Intended for graduate students and senior undergraduates without the usual background in stochastic processes (though this is desirable). Experience with *Matlab* or *Python* is essential, and the always-essential signals & systems mathematical toolbox.

Office hours (via zoom)

Joel Mathias: Fridays, 10-11am

Sean Meyn: Tuesdays after lecture, and also Thursdays if desperately needed

We can be reached for questions any time via Canvas, and will extend our office hours when needed. Meyn misses 80% of email, so please stick to Canvas for communication.

Exams, homework, and grading Homework problems will be assigned on a weekly basis, to be handed in on the date due (usually Fridays at 5pm). *Late homework will be ignored*

There will be two evening midterm exams, March 10 and April 21, from 6:00–7:30 p.m. You will be allowed *one* sheet of notes ($8\frac{1}{2} \times 11$; both sides) in the first exam, and *two* in the second. Otherwise, the exams are closed-book and closed-notes.

Tentative grading scheme: Homework problems will count 15%, the two midterm exams 60%, and 25% for a final project.

Resources: The course will follow the monograph found on Canvas. See lecture 1 for a big list of resources. The videos by Richard Murray on the foundations of control design are especially valuable as an introduction to the course: <https://simons.berkeley.edu/talks/murray-control-1>

Please see `EEL6935_Ctr1_Sys_Reinforce_Learn_Meyn_Approved.pdf` (on Canvas or click link at top of this file) for more information about logistics and course topics.

¹Bell Labs hosted Richard Bellman and Claude Shannon in the 1950s, as they were inventing the U.S. brand of control theory and information theory!