

EE/ISE 556: Stochastic Systems and Reinforcement Learning Fall 2018

Lecture Time & Venue: 11:00-12:20pm on Tue and Thu, VHE 210

Discussion Section: Friday 4-4:50, GFS 207

Office Hours: Tue., 12:30-1:30pm

Course page: <http://www-bcf.usc.edu/~rahuljai/Teaching.html>

Instructor: Prof. Rahul Jain, EEB 328, rahul.jain@usc.edu

TA: Hiteshi Sharma, hiteshis@usc.edu

Course Description:

Engineering Systems are increasingly expected to be *autonomous*. Intelligent control of autonomous systems requires accounting for uncertainty and dynamics. Thus, this course will focus on automatic control of stochastic systems. We will start by considering basic stochastic system models such as Markov decision processes and linear stochastic systems. We will then look at principles in finding optimal control laws for such system models based on dynamic programming. We will also cover estimation and control for linear systems, Kalman filtering. In the second part of the course, we will shift our focus to Reinforcement Learning, where we will cover Multi-armed bandit learning, Monte Carlo methods, Temporal Difference Learning, Policy Gradient Methods and On-policy RL methods. Course material is relevant for wide variety of applications such as Robotics, Finance, and Autonomous Vehicles. Expected audience for the course is MS and PhD students in EE, Computer Science and ISE.

Pre-requisites:

Students are expected to be familiar with probability theory at the level of EE 503. EE 512 is recommended preparation. Please contact the instructor for D clearance. The course evaluation will be through class participation, home works, an exam and/or a project (depending on class size). Some programming skill in a high level language are also going to be needed.

Grading: Exam (50%), Home-works (20%), Project (25%), Class Participation (5%)

Recommended Texts:

P. R. Kumar and P. Varaiya, *Stochastic Systems*, 1986

R. S. Sutton and A. G. Barto, *Reinforcement Learning, second ed. draft, July 2018*

Optional References:

D. Bertsekas, *Dynamic Programming and Optimal Control (Vol. I & II)*, 3rd ed., 2007

D. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*, Athena Scientific, 1996

Tentative Topics:

- (1) Markov Decision Processes
- (2) Dynamic Programming for finite horizon
- (3) DP for Infinite horizon
- (4) Linear Stochastic Systems
- (5) Estimation & Control of Linear Systems
- (6) System Identification
- (7) Adaptive Control
- (8) Introduction to Reinforcement learning
- (9) Multi-armed Bandit Learning
- (10) Monte-Carlo Methods
- (11) Temporal Difference Learning
- (12) Policy Gradient Methods
- (13) On-policy RL
- (14) Deep Reinforcement Learning

Projects: There is going to be a research project. You will be expected to read a few papers and do something creative. Work done will have to be summarized in a project report.

Statement for Students with Disabilities:

Any student requesting academic accommodations based on a disability is required to register with Disability Services and Programs (DSP) each semester. A letter of verification for approved accommodations can be obtained from DSP. Please be sure the letter is delivered to me (or to TA) as early in the semester as possible. DSP is located in STU 301 and is open 8:30 a.m.–5:00 p.m., Monday through Friday. The phone number for DSP is (213) 740-0776.

Statement on Academic Integrity:

USC seeks to maintain an optimal learning environment. General principles of academic honesty include the concept of respect for the intellectual property of others, the expectation that individual work will be submitted unless otherwise allowed by an instructor, and the obligations both to protect one's own academic work from misuse by others as well as to avoid using another's work as one's own. All students are expected to understand and abide by these principles. Scampus, the Student Guidebook, contains the Student Conduct Code in Section 11.00, while the recommended sanctions are located in Appendix A: <http://www.usc.edu/dept/publications/SCAMPUS/gov/>. Students will be referred to the Office of Student Judicial Affairs and Community Standards for further review, should there be any suspicion of academic dishonesty. The Review process can be found at: <http://www.usc.edu/student-affairs/SJACS/>.