




NYU

TANDON SCHOOL
OF ENGINEERING



ECE-GY 9223 Reinforcement Learning

Prof. Quanyan Zhu

Meeting Time: 3:00PM - 5:30PM Fridays

Course description: This course is a graduate level course focusing on the theory and practice of reinforcement learning. Reinforcement learning is a paradigm that focuses on the question: How to interact with an environment when the decision maker's current action affects future consequences. This course provides an accessible in-depth treatment of reinforcement learning and dynamic programming methods using function approximators. The course starts with a concise introduction to Markov Decision Processes and optimal control problems, in order to build the foundation. We present an extensive review of state-of-the-art approaches to dynamic programming and reinforcement learning with approximations. Theoretical guarantees are discussed on the solutions obtained, and numerical examples and applications are used to illustrate the properties of the individual methods.

Pre-requisites: The course is offered as an advanced topic graduate course. The pre-requisites or co-requisites for this course are EL-GY 6233 System Optimization Methods, EL-GY 6253 Linear Systems, and EL-GY 6303 Probability and Stochastic Processes, or their equivalent.

Grading and Exams:

Homework: 40%

Midterm: 30%

Projects: 30%

Main Reference:

[BT] D.P. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*, Athena Scientific, 1996.

[FV] J. Filar, K. Vrieze, *Competitive Markov Decision Processes*, Springer 1997.

[CBL] N. Cesa-Bianchi, G. Lugosi, *Prediction, Learning, and Games*, Cambridge University Press, 2006.

Supplementary Text:

[BBSE] K. Busoniu, R. Babuska, B. Schytter and D. Ernts, *Reinforcement Learning and Dynamic Programming*, CRC Press, 2010.

[CS] C. Szepesvari, *Algorithms for Reinforcement Learning*, Morgan and Claypool Publishers, 2010.

[SB] R.S. Sutton and A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 1998.

[SM] S. Meyn, *Control Techniques for Complex Systems*, Cambridge University Press, 2007.

Topics:

- Stochastic Control
- Dynamic Programming
- Approximate Dynamic Programming
- Q- and TD- Learning
- Stochastic Approximation
- Convergence Analysis
- Markov Decision Problem
- Multi-Armed Bandit Problem
- Dynamic programming
- Stochastic Games
- Learning in Games
- Distributed Learning

