

EECS 598: Reinforcement Learning Theory (Tentative)

Instructor: Lei Ying (leiying@umich.edu)

Credits:3

Prerequisite: EECS 502 (**required**)

Meets: TTH 10:30-12PM; 1005 DOW

Topics: This course covers fundamental theories and principles of reinforcement learning. Topics to be covered include:

1. Dynamic programming and the principle of optimality
2. Multi-armed bandit: epsilon-greedy, Upper Confidence Bound (UCB) algorithm, Thompson Sampling
3. Markov chains and Markov Decision Process (MDP)
4. Value iteration, policy iteration, and LP formulation
5. Q-Learning: Model-based and model-free
6. Linear function approximation and deep reinforcement learning
7. Temporal-difference learning
8. SARSA
9. Policy gradient algorithm and variance reduction
10. The ODE methods and convergence analysis