An Introduction to Stochastic Control and Reinforcement Learning

structure of the course

Monday, July 7

08:30-09:00: Introduction to the course

09:00-10:30: PART 1 (Simone Garatti)
Finite state Markov chain (discrete-time) and Markov decision processes (MDP) (controlled Markov chain) and their applications. Discrete-time stochastic control. Finite horizon stochastic control problem, principle of optimality (Bellman equation)

11:00-12:30: PART 2 (Subhrakanti Dey)
Dynamic Programming and its solutions, Closed form solution for the Linear Quadratic Gaussian (LQG) control problem.

14:30-16:00: PART 3
Infinite horizon stochastic control problems (discounted and average cost with finite state and action space), Bellman optimality equation, existence of stationary control policy.

16:30-18:00: PART 4 (Simone Garatti)
Solution methodologies – value iteration and policy iteration and related algorithms.

Tuesday, July 8

09:00-10:30 + 11:00-12:30: PART 5 (S. Garatti)
Curse of dimensionality in solving Dynamic Programming algorithms, Approximate Dynamic Programming algorithms – approximation in policy space and value space, contraction properties and error bounds, simulation-based implementation.

14:30-16:00 + 16:30-18:00: PART 6 (S. Garatti)
Intro to reinforcement learning in the setting of MDP. Temporal difference methods (TD(0), TD(λ)), convergence properties. On-policy TD control (SARSA), Off-policy TD control such as Q-learning and its convergence properties, Applications (3 hrs with a break)

Wednesday, July 9

09:00-10:30: PART 7 (S. Garatti)
Advanced reinforcement learning. Value function approximation with Linear methods and function approximation. Deep reinforcement learning.

11:00-12:30: PART 8
policy gradient methods, actor-critic based reinforcement learning and their applications to continuous control (such as LQG) problems (3 hrs with a break)