| | | |
|---|---|---|
| 1. | Title of the course | Reinforcement Learning |
| 2. | Course number | CS5224 |
| 3. | Status of the course | Elective |
| 4. | Structure of credits | 3-0-0-3 |
| 5. | Offered to | PG |
| 6. | New course/modification to | New course |
| 7. | To be offered by | Department of Computer Science and Engineering |
| 8. | To take effect from | July 2020 |
| 9. | Prerequisite | CoT |
| 10. | Whether approved by the Department | Yes |
| 11. | **Course Objective(s):** To provide a theoretical and algorithmic perspective of the sequential decision-making paradigm under uncertainty with an emphasis on reinforcement learning. | |
| 12. | **Course Content:** Basics: sequences, probability theory, stochastic processes, discrete-time Markov chain, stochastic approximation theory; Reinforcement Learning (RL) framework: Markov decision process, value function, optimal policy/value function, Bellman equations; Exact solution methods: dynamic programming; Prediction methods: Monte-Carlo methods, temporal difference methods; Control methods: Q-learning, sarsa, expected sarsa, dyna; Eligibility traces; Approximation methods: fitted value and policy iteration, on/off-policy prediction using linear and nonlinear function approximations, policy gradient method, actor-critic method; Advanced topics: deep RL, generalized Bellman equation, stochastic recursive inclusions, proof of convergence of RL methods. | |
| 13. | **Textbook(s):**<br>1. Bertsekas D P, *Reinforcement Learning and Optimal Control*, 1st Edition, Athena Scientific (2019).<br>2. Sutton R S and Barto A G, *Reinforcement Learning: An Introduction*, 2nd Edition, MIT Press (2018). | |
| 14. | **Reference(s):**<br>1. Ash R B, *Real Analysis and Probability*, 1st Edition, Academic Press (1972).<br>2. Borkar V S, *Stochastic Approximation: A Dynamical Systems Viewpoint*, 1st Edition, Hindustan Book Agency (2008).<br>3. Szepesvari C, *Algorithms for Reinforcement Learning*, 1st Edition, Morgan & Claypool (2010). | |