Name of the Programme: M.Sc. in Artificial Intelligence
Course Code: CSI-512
Title of the Course: Reinforcement Learning
Number of Credits: 2(2L-0T-0P)
Effective from AY: 2023-24

Prerequisites for the course	Linear algebra, multivariable calculus Basic machine learning knowledge	
Objectives	To enable the student to understand the reinforcement learning paradigm, to be able to identify when an RL formulation is appropriate, to understand the basic solution approaches in RL, to implement and evaluate various RL algorithms.	
Content	Review of ML fundamentals – Classification, Regression. Review of probability theory and optimization concepts. RL Framework; Supervised learning vs. RL; Explore-Exploit Dilemma; Examples. MAB: Definition, Uses, Algorithms, Contextual Bandits, Transition to full RL, Intro to full RL problem Intro to MDPs: Definitions , Returns, Value function, Q-function. Bellman Equation, DP, Value Iteration, Policy Iteration, Generalized Policy Iteration. Evaluation and Control: TD learning, SARSA, Q-learning, Monte Carlo, TD Lambda, Eligibility Traces. Maximization-Bias & Representations: Double Q learning, Tabular learning vs. Parameterized, Q-learning with NNs Function approximation: Semi-gradient methods, SGD, DQNs, Replay Buffer. Policy Gradients: Introduction, Motivation, REINFORCE, PG theorem, Introduction to AC methods Actor-Critic Methods, Baselines, Advantage AC, A3C Advanced Value-Based Methods: Double DQN, Prioritized Experience Replay, Dueling Architectures, Expected SARSA. Advanced PG/A-C methods: Deterministic PG and DDPG, Soft Actor-Critic (SAC) HRL: Introduction to hierarchies, types of optimality, SMDPs, Options, HRL algorithms POMDPS: Intro, Definitions, Belief states, Solution Methods; History-based methods, LSTMS, Q-MDPs, Direct Solutions, PSR. Model-Based RL: Introduction, Motivation, Connections to Planning, Types of MBRL, Benefits, RL with a Learnt Model, Dyna-style models, Latent variable models, Examples, Implicit MBRL. Case study on design of RL solution for real-world problems.	2 hrs 2 hrs 2 hrs 2 hrs 2 hrs 2 hrs 2 hrs 2 hrs 3 hrs 3 hrs 4 hrs 4 hrs
Pedagogy	Hands-on assignments / tutorials / peer-teaching / flip classroom/ presentations.	
References/ Readings	1. Reinforcement learning -Introduction by Richard sutton and Andrew barto, 2nd edition, MIT press. 2. Algorithms for reinforcement learning by Csaba Szepesvari, Ronald Brachman, et al,2010.	
Course Outcomes	1.Solid understanding of reinforcement learning concepts, theories, and algorithms. 2.Ability to implement and apply reinforcement learning algorithms to real-world problems. 3.Evaluation and analysis of reinforcement learning systems. 4.Critical thinking skills, staying updated with current research and trends.	

(Back to top)