III B.Tech I Semester

		L	T	P	C
23A39502	**INTRODUCTION TO REINFORCEMENT LEARNING (Professional Elective-I)**	**3**	**0**	**0**	**3**

Course Objectives:

- To understand the foundational concepts of Reinforcement Learning (RL) and its mathematical formulations.
- To explore dynamic programming, Monte Carlo methods, and temporal-difference learning.
- To study the advanced function approximation methods using neural networks.
- To apply RL algorithms to real-world decision-making problems.
- To introduce policy gradient methods and deep reinforcement learning techniques.

Course Outcomes:

Upon completion of the course, students will be able to:

- Explain the core principles of reinforcement learning and its interaction model.
- Apply tabular and approximate solution methods for prediction and control.
- Evaluate and compare Monte Carlo, TD, and policy gradient methods.
- Design reinforcement learning models for real-world environments.
- Integrate neural networks with reinforcement learning techniques.

UNIT I – Introduction to Reinforcement Learning

Introduction to Machine Learning and RL, Agent-environment interface, Goals and rewards, Returns: episodic and continuing tasks, Markov Decision Processes (MDP), Value functions: state-value and action-value functions.

UNIT II – Dynamic Programming and Monte Carlo Methods

Policy evaluation and improvement, Policy iteration and value iteration, Generalized policy iteration, Monte Carlo prediction and control, On-policy and off-policy MC methods.

UNIT III – Temporal-Difference Learning and Eligibility Traces

TD Prediction (TD(0)), SARSA and Q-Learning, Expected SARSA, n-step returns, Eligibility traces, TD(λ) methods.

UNIT IV – Function Approximation and Deep RL

Linear and non-linear function approximation, Feature construction, Deep Q Networks (DQN), Experience replay and fixed Q-targets, Double DQN and Dueling DQN, Challenges in deep RL

UNIT V – Policy Gradient and Actor-Critic Method

Policy gradient theorem, REINFORCE algorithm, Variance reduction techniques, Actor-Critic architecture, Proximal Policy Optimization (PPO), Applications in Robotics and Games

Textbooks:

1. Richard S. Sutton and Andrew G. Barto, "Reinforcement Learning: An Introduction", 2nd Edition, MIT Press, 2018. (Free online at http://incompleteideas.net/book/the-book-2nd.html)

Reference Books:

1. Csaba Szepesvári, "Algorithms for Reinforcement Learning", Morgan & Claypool, 2010.
2. Marco Wiering and Martijn van Otterlo, "Reinforcement Learning: State-of-the-Art", Springer, 2012.
3. David Silver, Reinforcement Learning Lecture Series, University College London (UCL).
4. François-Lavet et al., "An Introduction to Deep Reinforcement Learning", Foundations and Trends® in Machine Learning, 2018.

Online Learning Resources:

1. NPTEL Online Course:
   https://nptel.ac.in/courses/106106143 – Reinforcement Learning by Prof. Balaraman Ravindran, IIT Madras
2. DeepMind & UCL Lectures (David Silver):
   https://www.davidsilver.uk/teaching/
3. Coursera – Reinforcement Learning Specialization:
   https://www.coursera.org/specializations/reinforcement-learning