21AID12	REINFORCEMENT LEARNING	L	T	P	C
		3	0	0	3

Course Objectives

- To knowledge on the model the trial-and-error learning process that is needed in many problem situations
- To understand operations research, behavioral psychology, and AI.
- To apply the basic mathematical foundations of reinforcement learning.

UNIT I	INTRODUCTION TO REINFORCEMENT LEARNING	9 HOURS

Definition of reinforcement learning-History of reinforcement learning -Key concepts: agent, environment, reward, policy, value function, Q-learning-Examples of reinforcement learning applications-Markov Decision Processes (MDPs)- Properties of MDPs: state space, action space, transition function, reward function, discount factor -Value iteration and policy iteration algorithms-Applications of MDPs

UNIT II	DYNAMIC PROGRAMMING	9 HOURS

Bellman equations and optimality conditions-Policy evaluation, improvement, and iteration – algorithms-Finite and infinite horizon problems-Examples of dynamic programming algorithms in reinforcement learning

UNIT III	REINFORCEMENT LEARNING METHODS	9 HOURS

Monte Carlo estimation of state-value and action-value functions-Incremental Monte Carlo updates and exploration techniques-First-visit and every-visit Monte Carlo algorithms-Policy evaluation with Monte Carlo methods- Temporal-Difference Learning- TD prediction and control algorithms: Sarsa, Q-learning, expected Sarsa-Off-policy TD learning with importance sampling-Comparison with Monte Carlo methods

UNIT IV	EXPLORATION AND EXPLOITATION	9 HOURS

Exploration-exploitation trade-off and its importance in reinforcement learning-Epsilon-greedy, optimistic initialization, and Boltzmann exploration methods-Upper confidence bound (UCB) and Thompson sampling algorithms-Bayesian reinforcement learning- Policy Search- Proximal policy optimization (PPO) and trust region policy optimization (TRPO) algorithms- Multi-Agent Reinforcement Learning- Multi-agent Q-learning and policy gradient methods

UNIT V	REINFORCEMENT LEARNING APPLICATIONS	9 HOURS

Reinforcement learning applications in robotics, gaming, finance, and healthcare-Current research topics in reinforcement learning- Discussion of potential research projects and applications

UNIT VI	RECENT TRENDS	

Recent trends on applications of Reinforcement Learning

TOTAL PERIODS: 45

Course Outcomes:

- Define the key features of reinforcement learning that distinguishes it from AI and non-interactive machine learning
- Given an application
- Describe (list and define) multiple criteria for analyzing RL algorithms and evaluate algorithms on these metrics: e.g. regret, sample complexity, computational complexity, empirical performance, convergence, etc
- Describe the exploration vs exploitation challenge and compare and contrast at least two approaches for addressing this

Textbooks:

1. Reinforcement Learning: An Introduction, Sutton and Barto, 2nd Edition.

Reference Books:

1. Deep Reinforcement Learning" by Pieter Abbeel and John Schulman
2. Hands-On Reinforcement Learning with Python" by Sudharsan Ravichandiran
3. Reinforcement Learning: State-of-the-Art" by Marco Wiering and Martijn van Otterlo