INSTITUTE OF AERONAUTICAL ENGINEERING
(Autonomous)
Dundigal - 500 043, Hyderabad, Telangana

COURSE CONTENT

REINFORCEMENT LEARNING								
VII Semester: CSE(DS)								
Course Code	**Category**	**Hours / Week**			**Credits**	**Maximum Marks**		
		L	**T**	**P**	**C**	**CIA**	**SEE**	**Total**
ACDD23	Elective	3	0	0	3	40	60	100
Contact Classes: 48	**Tutorial Classes: Nil**	**Practical Classes: Nil**				**Total Classes: 48**		
Prerequisites: Basic knowledge of linear algebra, probability theory, calculus, machine learning fundamentals, and programming (preferably in Python).								

I. COURSE OVERVIEW:

This course introduces students to the foundational concepts and algorithms of Reinforcement Learning (RL), a key area in machine learning where agents learn optimal behaviors through interactions with an environment. The course covers Markov Decision Processes (MDPs), dynamic programming, Monte Carlo methods, Temporal Difference learning, policy and value iteration, Q-learning, Deep Q Networks (DQN), and policy gradient methods. Students will explore applications in robotics, game AI, and autonomous systems while gaining hands-on experience in implementing RL algorithms using Python-based frameworks like OpenAI Gym and TensorFlow/PyTorch.

II. COURSE OBJECTIVES:

The students will try to learn:

I. The mathematical foundations of reinforcement learning including agent-environment interaction, rewards, states, and policies.
II. The formulation and solution of decision-making problems using Markov Decision Processes (MDPs).
III. The Core RL algorithms including value-based, policy-based, and actor-critic methods.
IV. The Practical applications of RL in real-world systems such as game playing, robotics, and self-learning agents using simulation tools.

III. COURSE OUTCOMES:

CO1 **Explain** the fundamental concepts of reinforcement learning, including agents, environments, states, actions, and rewards.
CO2 **Formulate** decision problems as Markov Decision Processes and apply Bellman equations to model optimal behavior.
CO3 **Analyze** and **implement** Dynamic Programming, Monte Carlo, and Temporal Difference learning techniques.
CO4 **Compare** and **evaluate** different RL algorithms such as Q-learning, SARSA, and DQN in simulated environments.
CO5 **Design** and **train** reinforcement learning agents using deep learning frameworks and policy gradient methods.
CO6 **Assess** the performance and convergence behavior of RL algorithms in solving real-world problems like autonomous control and game AI.

IV. COURSE CONTENT:

MODULE I: INTRODUCTION TO REINFORCEMENT LEARNING (9)
What is Reinforcement Learning? Agent-Environment Interface, Goals and Rewards, Exploration vs Exploitation, Applications of RL, Types of RL Algorithms, Elements of RL Problems.

MODULE II: MARKOV DECISION PROCESSES (MDPs) (10)
Definition of MDPs, Transition Probabilities, Reward Functions, Policies and Returns, Value Functions, Bellman Expectation Equations, Optimal Value Functions, Solving MDPs using Dynamic Programming.

MODULE III: MONTE CARLO AND TEMPORAL DIFFERENCE LEARNING (10)
Monte Carlo Methods for Prediction and Control, First Visit vs Every Visit MC, TD(0), SARSA.

Q-Learning, n-step Bootstrapping, TD(λ), Convergence Properties

MODULE IV: FUNCTION APPROXIMATION AND DEEP Q-LEARNING (9)
Need for Function Approximation, Linear and Non-linear Function Approximators, Deep Q Networks (DQN), Experience Replay, Target Networks, DQN Variants (Double DQN, Dueling DQN)

MODULE V: POLICY GRADIENT METHODS AND ADVANCED TOPICS (10)
Policy Gradient Theorem, REINFORCE Algorithm, Actor-Critic Methods, Proximal Policy Optimization (PPO), Trust Region Policy Optimization (TRPO), Case Studies: AlphaGo, Robotics, Autonomous Driving

V. TEXTBOOKS:

1. Richard S. Sutton and Andrew G. Barto, Reinforcement Learning: An Introduction, 2nd Edition, MIT Press, 2018.
2. Phil Winder, Reinforcement Learning: Industrial Applications of Intelligent Agents, O'Reilly, 2020.
3. Laura Graesser & Wah Loon Keng, Foundations of Deep Reinforcement Learning, Addison-Wesley, 2019.

VI. REFERENCE BOOKS:

1. Mimura Masashi, Deep Reinforcement Learning, Packt Publishing, 2020
2. Lucian Busoniu et al., Reinforcement Learning and Dynamic Programming Using Function Approximators, CRC Press, 2010

VII. WEB REFERENCES

1. NPTEL – Reinforcement Learning by IIT Madras
2. OpenAI Gym Documentation
3. DeepMind Research on Reinforcement Learning
4. Spinning Up in Deep RL – OpenAI

VIII. MATERIALS ONLINE

1. Course template
2. Tutorial question bank
3. Tech talk topics
4. Open-ended experiments
5. Definitions and terminology
6. Assignments
7. Model question paper – I
8. Model question paper – II

9. Lecture notes
10. PowerPoint presentation
11. E-Learning Readiness Videos (ELRV)