

Branch Name:	Artificial Intelligence & Data Science
Program Code:	ET108
Semester/Year:	Semester VII/Fourth Year
Subject Title:	Professional Elective – III Reinforcement Learning and its application
Subject Code:	1ET1080704
Pre-requisite Course:	Probability, Machine Learning

Course Objective:

1. In this course, we will investigate the process of learning through an agent's mistakes. This is very dissimilar to supervised machine learning and more closely resembles how people learn via experience.
2. The focus of Reinforcement Learning (RL) is on issues that call for a series of decisions to be made in order. The fundamentals of reinforcement learning and its application are the focus of this course.

Teaching and Examination Scheme:

Teaching Scheme(Hours per week)				Evaluation Scheme (Marks)				
Lecture	Tutorial	Practical	Credit	Theory		Practical		Total
				University Assessment	Continuous Assessment	University Assessment	Continuous Assessment	
04	-	02	05	70	30	30	20	150

Subject Contents

Sr. No	Topic	Total Hours	Weight (%)
1	Introduction: Course logistics and overview. Origin and history of Reinforcement Learning research. Its connections with other related fields and with different branches of machine learning.	6	15
2	Probability Primer: Brush up of Probability concepts - Axioms of probability, concepts of random variables, PMF, PDFs, CDFs, Expectation. Concepts of joint and multiple random variables, joint, conditional and marginal distributions. Correlation and independence.	6	15
3	Markov Decision Process: Introduction to RL terminology, Markov property, Markov chains, Markov reward process (MRP). Introduction to and proof of Bellman equations for MRPs along with proof of existence of solution to Bellman equations in MRP. Introduction to Markov decision process (MDP), state and action value functions, Bellman expectation equations, optimality of value functions and policies, Bellman optimality equations.	8	20
4	Prediction and Control by Dynamic Programming; Overview of dynamic programming for MDP, definition and formulation of planning in MDPs, principle of optimality, iterative policy evaluation, policy iteration, value iteration, Banach fixed point theorem, proof of contraction mapping property of Bellman expectation and optimality operators, proof of convergence of policy evaluation and value iteration algorithms, DP extensions.	8	25
5	Monte Carlo Methods for Model Free Prediction and Control: Overview of Monte Carlo methods for model free RL, First visit and every visit Monte Carlo, Monte Carlo control, On policy and off policy learning, importance sampling.	6	15

6	TD Methods: Incremental Monte Carlo Methods for Model Free Prediction, Overview TD(0), TD(1) and TD(λ), k-step estimators, unified view of DP, MC and TD evaluation methods, TD Control methods - SARSA, Q-Learning and their variants. Applications of RL for real-world problems	6	15
----------	--	----------	-----------

Course Outcome:

Upon successful completion of this course the student would be able to:

1. Knowledge related to Reinforcement Learning and its application.
2. Learn how to define RL tasks and the core principals.
3. Apply RL for real-world problems.

List of References:

1. "Reinforcement Learning", Sutton and Barto, MIT Press, 2018
2. "An Introduction to Internet of Things: Connecting Devices, Edge Gateway, and Cloud with Applications", Rahul Dubey, Cengage India Publication
3. "Reinforcement Learning: Theory and Algorithms" Alekh Agarwal, Nan Jiang, Sham Kakade, and John Langford
4. Machine Learning: A Reinforcement Approach" by Satinder Singh and Sanjay Lal
5. "Artificial Intelligence: Reinforcement Learning in Python: Complete guide to artificial intelligence and machine learning", Lazy Programmer, Kindle Edition

E-Resources:

1. <https://www.synopsys.com/ai/what-is-reinforcement-learning.html>
2. <https://www.simplilearn.com/tutorials/machine-learning-tutorial/reinforcement-learning>
3. <https://www.udemy.com/course/beginner-master-rl-1/>
4. <https://www.guru99.com/reinforcement-learning-tutorial.html>
5. <https://towardsdatascience.com/reinforcement-learning-101-e24b50e1d292>

List of Experiments:

1. Understand Reinforcement Learning.
2. Study the concepts of random variables.
3. Learn conditional and marginal distributions.
4. Study about Markov property and Markov chains.
5. Learn about proof of Bellman equations for MRPs.
6. Learn dynamic programming for MDP.
7. Study proof of convergence of policy evaluation and value iteration algorithms.
8. Learn about Monte Carlo methods for model free RL.
9. Study Incremental Monte Carlo Methods for Model Free Prediction.
10. Study the Reinforcement Learning application.

Course Learning Outcomes(CLO): On completion of this course, the students will be able to:

CLO	Description	Bloom's Taxonomy Level
CLO1	Learn how to define RL tasks and the core principals behind the RL	2 Understanding
CLO2	Understand the basics of probability and distribution functions	2 Understanding
CLO3	Understand Bellman equations for MRPs	2 Understanding 3 Applying
CLO4	Applying RL for real-world problems	3 Applying

Mapping of CLOs with POs & PSOs

Course Learning Outcomes	Program Outcomes(POs)												Program Specific Outcomes(PSOs)			
	PO 1	PO 2	PO 3	PO 4	PO 5	PO 6	PO 7	PO 8	PO 9	PO 10	PO 11	PO 12	PS O 1	PS O 2	PS O 3	PS O 4
CLO1	M	M	H	M	H	M			L		L	M	H	M		
CLO2	M	M	H	M	L	L				L	L	L	M	L		
CLO3	M	M	H	L	M	M				L	L		M	L		
CLO4	H	M	M	L	M	H			L	M	M	M	H	M		

H:High, M:Medium, L:Low