	May-June 2025 B. Tech. Program: Artificial Intelligence and Data Science, Scheme: II Regular Examination: Last Year, Semester: VIII Course Code: AIC801 and Course Name: Reinforcement Learning		
Date of Exam: 19/05/2025	Duration: 02.5 Hours		Max. Marks: 60

Instructions:
(1)All questions are compulsory.
(2)Draw neat diagrams wherever applicable.
(3)Assume suitable data, if necessary.

Q. No.	Question	Max. Marks	CO	BT level
Q 1	Solve any **two** questions out of three: (05 marks each)	10		
a)	Explain the basic concept of an agent and environment in the context of Reinforcement Learning..		CO1	U
b)	What is the fundamental building block of deep neural networks, and how does it function?		CO3	U
c)	Define linear policies and explain their advantages and limitations.		CO4	U
Q 2	Solve any two questions out of three: (05 marks each)	10		
a)	What is a Multi-arm Bandit problem, and how does it relate to Reinforcement learning?		CO2	U
b)	Define deep reinforcement learning and highlight its differences from traditional reinforcement learning.		CO3	U
c)	Outline the key stages in the life cycle of a reinforcement learning project.		CO4	U
Q.3	Solve any two questions out of three. (10 marks each)	20		
a)	Define a Markov Decision Process (MDP). Explain the components of an MDP—states, actions, rewards, transition probabilities, and the discount factor.		CO2	U
b)	What is experience replay, and how does it address certain challenges in training deep Q-learning models? Explain the role of a replay buffer in experience replay.		CO3	U
c)	Define multi-agent reinforcement learning (MARL) and how it differs from single-agent RL. Discuss the challenges associated with coordination and competition in multi-agent environments.		CO5	U

Seat No.:

	May-June 2025 B. Tech. Program: Artificial Intelligence and Data Science, Scheme: II Regular Examination: Last Year, Semester: VIII Course Code: AIC801 and Course Name: Reinforcement Learning			
Date of Exam: 19/05/2025	Duration: 02.5 Hours		Max. Marks: 60	

Q.4	Solve any two questions out of three. (10 marks each)	20		
a)	What is Q-Learning, and how does it differ from other reinforcement learning algorithms? Explain the Q-value update equation in Q-learning.		CO2	U
b)	Explain the concept of learning a policy directly in reinforcement learning. What are the potential advantages of learning a policy directly compared to value-based methods?		CO4	U
c)	Discuss the role of reinforcement learning in solving the elevator dispatching problem. Explain the potential benefits of using RL in elevator dispatching systems.		CO6	U

****

Seat No.: