Multi-Agent Reinforcement Learning

Code: CS31225CS	Course Type: Elective	Semester: II
Evolution Scheme (TA-MSE-ESE): 20-30-100	Total Marks: 150	L-T-P (Credits): 3-0-0 (3)

Pre-Requisites: Reinforcement Learning, Probability and Statistics, Linear Algebra, Optimization, Python Programming, and basic Deep Learning.

Course Objective: To develop understanding of reinforcement learning, game-theoretic multi-agent systems, modern MARL algorithms, and their applications in robotics, autonomous systems, traffic control, games, and distributed decision-making environments.

Course Outcomes: At the end of the course, the student will be able to

CO-1	Explain the foundations of reinforcement learning, multi-agent systems, matrix games, and stochastic games in the context of MARL.
CO-2	Analyze cooperative, competitive, and mixed multi-agent learning settings using suitable MARL paradigms and solution concepts.
CO-3	Apply and compare modern MARL algorithms such as independent learning, CTDE-style methods, value decomposition, MADDPG, QMIX, and self-play.
CO-4	Evaluate MARL methods for communication, coordination, and real-world applications such as robotics, traffic control, games, and autonomous systems.

Course Articulation Matrix:

PO	PEO-1	PEO-2	PEO-3	PEO-4	PEO-5
CO-1	2	3	1	1	2
CO-2	2	3	1	1	2
CO-3	3	3	1	1	3
CO-4	2	2	1	2	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Unit-1: Reinforcement Learning and Multi-Agent Foundations
Review of reinforcement learning: Markov decision processes, value functions, Bellman equations, policy evaluation and improvement, Q-learning, policy gradients, and actor-critic methods. Multi-agent systems: agents, environments, cooperation, competition, and coordination. Matrix and stochastic games, partially observable settings, and Dec-POMDP intuition. Challenges in MARL: non-stationarity, partial observability, scalability, and credit assignment.

Unit-2: Core Multi-Agent Reinforcement Learning Algorithms
Independent learners, centralized and decentralized learning, centralized training with decentralized execution (CTDE), joint action learning, reward shaping, parameter sharing, and value decomposition. MARL algorithms: MADDPG, QMIX, shared policies, self-play, and adversarial training

Unit-3: Game-Theoretic and Deep MARL Perspectives

Extensive-form and imperfect-information games, equilibrium concepts, no-regret learning, and stochastic games. Deep reinforcement learning for large-scale games, deep fictitious play, Deep CFR, double oracle methods, and policy optimization in games. Applications to poker, Hanabi, Stratego, and Diplomacy.

Unit-4: Communication, Coordination, and Applications

Learning to communicate, coordination under partial observability, and task allocation in cooperative and competitive settings. Evaluation and benchmarking of multi-agent systems. Applications in traffic control, robotics, autonomous systems, distributed control, and resource allocation.

Learning Resources:

1. Michael L. Littman, Selected Readings on Markov Games and Multi-Agent Reinforcement Learning.
2. Yoav Shoham and Kevin Leyton-Brown, Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations, Cambridge University Press, 2008

Text Books:

1. Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer, Multi-Agent Reinforcement Learning: Foundations and Modern Approaches MIT Press, 2024, First Edition, ISBN: 978-0262049375
2. Richard S. Sutton and Andrew G. Barto, Reinforcement Learning: An Introduction, Second Edition, MIT Press, 2018, ISBN: 978-0262039246