



National University of Engineering (UNI)
School of Artificial Intelligence
Syllabus 2027-I

1. COURSE

AI265. Reinforcement Learning (Mandatory)

2. GENERAL INFORMATION

2.1 Course	: AI265. Reinforcement Learning
2.2 Semester	: 7 th Semester.
2.3 Credits	: 4
2.4 Hours	: 2 HT; 4 HP;
2.5 Duration of the period	: 16 weeks
2.6 Type of course	: Mandatory
2.7 Learning modality	: Face to face
2.8 Prerequisites	: AI263. Introduction to Machine Learning. (6 th Sem)

3. PROFESSORS

Meetings after coordination with the professor

4. INTRODUCTION TO THE COURSE

Reinforcement Learning (RL) is fundamental for autonomous systems. This course covers key algorithms such as Q-Learning, DQN, and Policy Gradients, with applications in robotics and games.

5. GOALS

- Design RL agents for stochastic environments.
- Implement RL algorithms using modern frameworks (OpenAI Gym, TensorFlow/PyTorch).
- Analyze ethical problems in autonomous systems based on RL.

6. COMPETENCES

AG-C08) Problem Analysis: Identifies, formulates, and analyzes complex computing problems. (Usage)

AG-C09) Design and Development of Solutions: Designs, implements, and evaluates solutions for complex computing problems. (Usage)

7. TOPICS

Unit 1: Foundations of RL (18 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
• Introduction to the Reinforcement Learning paradigm: agent, environment, actions, states, rewards. • Formalization of sequential problems: Markov Decision Processes (MDPs) and Partially Observable MDPs (POMDPs). • Concept of horizon and discount factor (γ). • The Bellman Equation: foundation of optimality for state-value (V) and action-value (Q) functions. • Dynamic Programming Methods: Iterative Policy Evaluation and Improvement. • Monte Carlo Methods: learning from complete episodes. • Temporal Difference (TD) Algorithms: TD(0), SARSA, and the bias-variance trade-off. • Q-Learning (off-policy): derivation, convergence properties, and parameters (learning rate, epsilon-greedy).	• Model problems as MDPs [Evaluar] • Implement Q-Learning [Usar]
Readings : [sutton2018rl]	

Unit 2: Advanced Classical RL Methods (14 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
• Limitations of tabular methods and the curse of dimensionality. • Value Function Approximation: replacing tables with parametric models (linear, non-linear). • Temporal Difference Gradient Algorithms: minimization of the Bellman error. • Eligibility Traces: unification of Monte Carlo and TD ($TD(\delta)$). • Integrated Planning and Learning: Environment models and Dyna-Q algorithms. • Policy-Based Methods (Policy Search): introduction to policy gradient. • Exploration vs. Exploitation Trade-off: advanced strategies (UCB, entropy-based exploration).	• Model problems as MDPs [Evaluar] • Implement Q-Learning [Usar]
Readings : [sutton2018rl]	

Unit 3: Deep Reinforcement Learning (20 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
• The Deep Learning revolution in RL: training stabilization. • Deep Q-Networks (DQN): architecture, loss function, and stabilization mechanisms (Target Network, Experience Replay). • DQN extensions: Double DQN, Dueling DQN, Prioritized Experience Replay. • Stochastic Policy for continuous action spaces: introduction to Actor-Critic methods. • Actor-Critic Methods: policy gradient decomposition, advantage function (A2C). • Advanced Actor-Critic Methods: A3C (Asynchronous Advantage Actor-Critic) and its parallel architecture. • Direct Policy Gradient: REINFORCE with and without baseline. • Introduction to maximum entropy methods (Soft Actor-Critic - SAC) and offline RL.	• Train DQN agents in complex environments [Evaluar] • Optimize policies with A3C [Usar]
Readings : [mnih2015dqn], [silver2017a3c]	

Unit 4: Ethics, Safety, and Applications in RL (12 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
• Biases in reward function design: misspecification, misaligned rewards, and reward hacking. • Challenges of generalization and robustness in RL-trained agents. • Safety in autonomous systems: avoiding catastrophic failures, Safe RL, and AI Alignment. • Case studies: analysis of failures in simulations and real-world robots. • Adversarial attacks on RL: attacks on sensors, policy function, and agent's reward function. • Social implications: automation, decision-making, and transparency (interpretable agents). • Ethical frameworks for the development and deployment of autonomous agents.	• Identify ethical risks in RL applications [Familiarize] • Propose mitigation strategies [Usar]
Readings : [sutton2018rl]	

8. WORKPLAN

8.1 Methodology

Individual and team participation is encouraged to present their ideas, motivating them with additional points in the different stages of the course evaluation.

8.2 Theory Sessions

The theory sessions are held in master classes with activities including active learning and roleplay to allow students to internalize the concepts.

8.3 Practical Sessions

The practical sessions are held in class where a series of exercises and/or practical concepts are developed through problem solving, problem solving, specific exercises and/or in application contexts.

9. EVALUATION SYSTEM

***** EVALUATION MISSING *****

10. BASIC BIBLIOGRAPHY