Reinforcement Learning

Academic Year: (2023 / 2024) **Review date: 08-05-2023**

Department assigned to the subject: Computer Science and Engineering Department

Coordinating teacher: FERNANDEZ REBOLLO, FERNANDO

Type: Electives ECTS Credits : 3.0

Year : 1 Semester : 1

REQUIREMENTS (SUBJECTS THAT ARE ASSUMED TO BE KNOWN)

Machine Learning contents are recommended for the Reinforcement Learning course

DESCRIPTION OF CONTENTS: PROGRAMME

Introduction to Reinforcement Learning

- Introduction to reinforcement learning
- Markov Decision Processes
- Policies and optimality: discounted infinite horizon
- Value Functions

Dynamic Programming
- Problem Solving on MDP: model-free, model-based and dynamic programming methods
- Policy Iteration Algorithm
- Value Iteration Algorithm

Direct reinforcement learning

- Monte Carlo methods: and Monte Carlo with exploratory start
- Model-free methods: Q-Learning
- Example of execution of Q-Learning
- On-policy methods vs. off-policy: SARSA
- Exploration and exploitation: e-greedy and softmax

Model-Based Methods
- Model Learning
- Dyna-Q

Representation in Reinforcement Learning
- Representation of the space of states, actions and Q
- State space discretization: uniform and adaptive methods
- Approximate methods to represent the function Q: Batch Q-Learning

Generalization Through Function Approximation
- Approximation through neural networks
- Deep reinforcement learning

Policy Search Methods
- Policy Approximation
- Actor-critic methods
- Proximal Policy Optimization (PPO)

Other Reinforcement Learning topics
- Hierarchical Reinforcement Learning
- Transfer of learning learned
- Multi-agent Reinforcement Learning
- Safe reinforcement learning
- Offline Reinforcement Learning
- Multi-objective Reinforcement Learning
- Partially observable Reinforcement Learning

Reinforcement Learning in the real world:
- Applications of reinforcement learning
- Reinforcement learning frameworks and software

% end-of-term-examination:	30
% of continuous assessment (assigments, laboratory, practicals…):	70

BASIC BIBLIOGRAPHY

- Richard Sutton and Andrew Barto Reinforcement Learning: an Introduction, The MIT Press.

BASIC ELECTRONIC RESOURCES

- DeepMind . MuJoCo: https://mujoco.org/
- Open AI . Open AI Proximal Policy Optimization: https://openai.com/research/openai-baselines-ppo
- Open AI . Gymnasium: https://gymnasium.farama.org/