

Reinforcement Learning (AI644)

Prepared by Dr. Ashraf Elsayed

Orientation Show

عرض توجيهي



Reinforcement Learning

Orientation Show

Prepared by
Dr. Ashraf Elsayed

About me

Associate Professor of Computer Science, Faculty of Science, Alexandria University, Egypt.

My research Interests include Artificial Intelligence, Machine Learning, Deep Learning, Data Science, Big data Analytics, Medical Image Mining.

Qualifications:


received B.Sc. degree from the Alexandria University, Alexandria, Egypt, in 1995, the M.Sc. degree from the Alexandria University, Alexandria, Egypt, in 2004, both in Computer science. And the Ph.D. degree from the University of Liverpool, UK, in 2012, in Computer science.

Reinforcement Learning

What is Reinforcement Learning?

- RL is a type of machine learning paradigm where an agent learns to make decisions by interacting with an environment. The agent takes actions in the environment, and based on these actions, it receives feedback in the form of rewards or penalties. The goal of the agent is to learn a policy, which is a mapping from states to actions, that maximizes the cumulative reward over time.

Machine Learning




Machine Learning Areas

- **Supervised Learning:** Data and corresponding labels are given
- **Unsupervised Learning:** Only data is given, no labels provided
- **Semi-supervised Learning:** Some (if not all) labels are present
- **Reinforcement Learning:** An agent interacting with the world makes observations, takes actions, and is rewarded or punished; it should learn to choose actions in such a way as to obtain a lot of reward


Supervised Learning

- Learning a discrete function: **Classification**
 - Boolean classification:
 - Each example is classified as true(positive) or false(negative).
- Learning a continuous function: **Regression**


Steps



- Given $x_1, x_2, \dots, x_n$ (without labels)
- Output hidden structure behind the x 's
 - E.g., clustering



Unsupervised Learning



Social network analysis




Market segmentation



Image credit: NASA/JPL-Caltech/E. Churchwell (Univ. of Wisconsin, Madison)

Astronomical data analysis


Reinforcement Learning




What is Reinforcement Learning?

- Learning from interaction with an *environment* to achieve some long-term *goal* that is related to the *state of the environment*
- The goal is defined by *reward* signal, which must be maximized
- *Agent* must be able to partially/fully sense the environment state and take *actions* to influence the environment state
- The *state* is typically described with a feature-vector

Agent, Environment, States, Actions and Rewards



Example (Tic-Tac-Toe)



- Operating a drone or autonomous vehicle
- Manipulating a robot to navigate the environment and perform various tasks
- Managing an investment portfolio and taking trading decisions
- Playing games such as Go, Chess, video games

Primary Goal of This Course

- This course has been designed to give students an extensive overview of Reinforcement Learning issues, tools and techniques that are critical in solving problems in RL domain.
- **This course aims to:** Provide students with concepts of:
 - a solid understanding of key concepts, algorithms, and techniques used in Reinforcement Learning.
 - Understanding the formal framework of Markov Decision Processes, which is a mathematical model used to describe decision-making problems in Reinforcement Learning.
 - Understanding policies as strategies that map states to actions. Students learn how to formulate and evaluate different policies in the context of RL.
 - Grasping the trade-off between exploration (trying new actions to discover their effects) and exploitation (choosing actions that are known to yield high rewards). This is a crucial aspect of RL.
 - Applying RL concepts to real-world problems and case studies, such as game playing, robotics, and autonomous systems.

Units of Course Contents



Lecture #	Course unit	Week #
1	Introduction to Reinforcement Learning	1 st
2	Exploration and Exploitation	2 nd
3	Markov Decision Processes	3 rd
4	Value Functions & Bellman Equations	4 th
5	Dynamic Programming	5 th
6	Sample-based Learning Methods	6 th
7	Temporal Difference Learning Methods for Prediction	7 th
8	Temporal Difference Learning Methods for Control	8 th
9	Expected SARSA & Planning	9 th
10	Dyna-Q architecture, Dyna-Q algorithm	10 th

Course Learning Outcomes (CLOs)

By the end of the course, the students should be able to:

- Recognize the advanced paradigms of Reinforcement learning.
- Demonstrate use of appropriate theory, science, mathematics and computational technology to analyze and solve real-life problems
- Formalize a task as a reinforcement learning problem.
- Apply suitable algorithms to tackle different reinforcement learning problems.
- Solve problems relevant to **applications of Reinforcement Learning** in computer science using ideas and techniques some of which are at the forefront of the discipline

Recommended Textbooks





Thank You