

Reinforcement Learning


Model-Free Reinforcement Learning

Presentation

The goal of this lab is to implement model-free reinforcement learning algorithms. We first consider very simple environments and then move to more complex ones.

Exercise 1.

In this first exercise, you will implement two simple environments. The dynamic of the each environment is described by the corresponding figure.



Create a class `Env` with the following methods :

- `__init__` : the constructor of the class. It should take as input the number of states and the number of actions.
- `reset` : this method should reset the environment to its initial state.
- `play` : this method should take as input a state and an action and returns the next state and the reward.

Create three classes `Env1`, `Env2` and `Env3` that inherit from the class `Env` and implement the corresponding dynamic.

Exercise 2. (Model-Free Methods)

In this exercise, you will implement different model-free reinforcement learning algorithms. For this purpose, you will create a file `model_free_agents.py` and implement the following agents :

- A class `MonteCarlo` that implements a Monte Carlo agent.
- A class `SARSAAgent` that implements a SARSA agent.
- A class `QLearningAgent` that implements a Q-learning agent.

Exercise 3. (Application)

In this exercise, you will apply the agents implemented in the previous exercise to the environments implemented in the first exercise.