AI Lab - Exam

2018

Riccardo Sartea

University of Verona
Department of Computer Science

June 7^{th} 2018

UNIVERSITÀ
di VERONA
Dipartimento
di INFORMATICA

Upgrade Gym AI Lab

Listing 1: Upgrade the repository

```
cd gym-ai-lab
git pull
pip3 install --user -I --no-deps -e .
```

In case of problems you can just delete the "gym-ai-lab" folder and clone it again.

Listing 2: Upgrade the repository

```
rm -rf gym-ai-lab
git clone https://github.com/SaricVr/gym-ai-lab
cd gym-ai-lab
pip3 install --user -I --no-deps -e .
```

How to Submit Your Solutions

Create a .zip file named MATRICOLA_AI_Exam.zip, e.g., VR487172_AI_Exam.zip containing the following:

- exercise1.txt with:
  - The output printed by your code: path and stats for each algorithm
  - The answer to the questions
- exercise2.txt with:
  - The output printed by your code: the final policy for each algorithm
  - The answer to the questions
- rewards.png: the chart representing the rewards obtained by the algorithms during learning
- lenghts.png: the chart representing the lenghts of the learning episodes for the algorithms
- a folder named "code" containing the code you run to obtain the results

Note: the code you provide must execute without any errors! Therefore, include any additional file/procedure needed

How to Submit Your Solutions II

Send the zip file to:

`alessandro.farinelli@univr.it`

with *subject*: esame laboratorio IA <MATRICOLA>

Guidelines

- Each answer must explicitly refer to the corresponding exercise
- Each answer must be explicitly motivated referring to the results obtained by your code
- Results must be reproducible: if you refer to a result your code is not producing in output, or that is different from your code's output, the answer will be invalidated
- Remember to report the output of your code in the .txt along with the answers

Exercise 1 - Environment

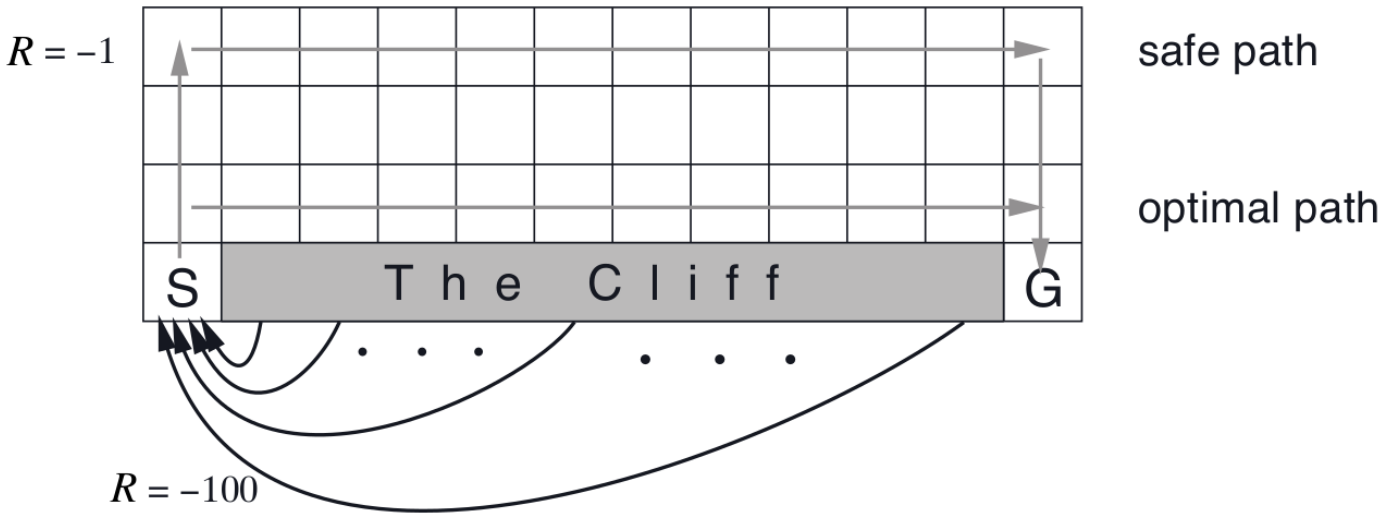

Actions

- 0 - U
- 1 - R
- 2 - D
- 3 - L

Environment taken from Sutton and Barto

Environment name in gym: "CliffWalkingExam2018-v0"

Exercise 1 - Tasks

State which search-based algorithm[1] (considering both tree search and graph search versions) you would choose to solve this problem. Motivate your choice and show a print-out of the path and statistics (max number of node in memory and number of expanded node) for the algorithms you considered.

Warning

Given the size of the environment, DO NOT USE the tree search version of BFS, UCS and Greedy.

[1] Use the L_1 norm as distance heuristic

Exercise 2 - Environment

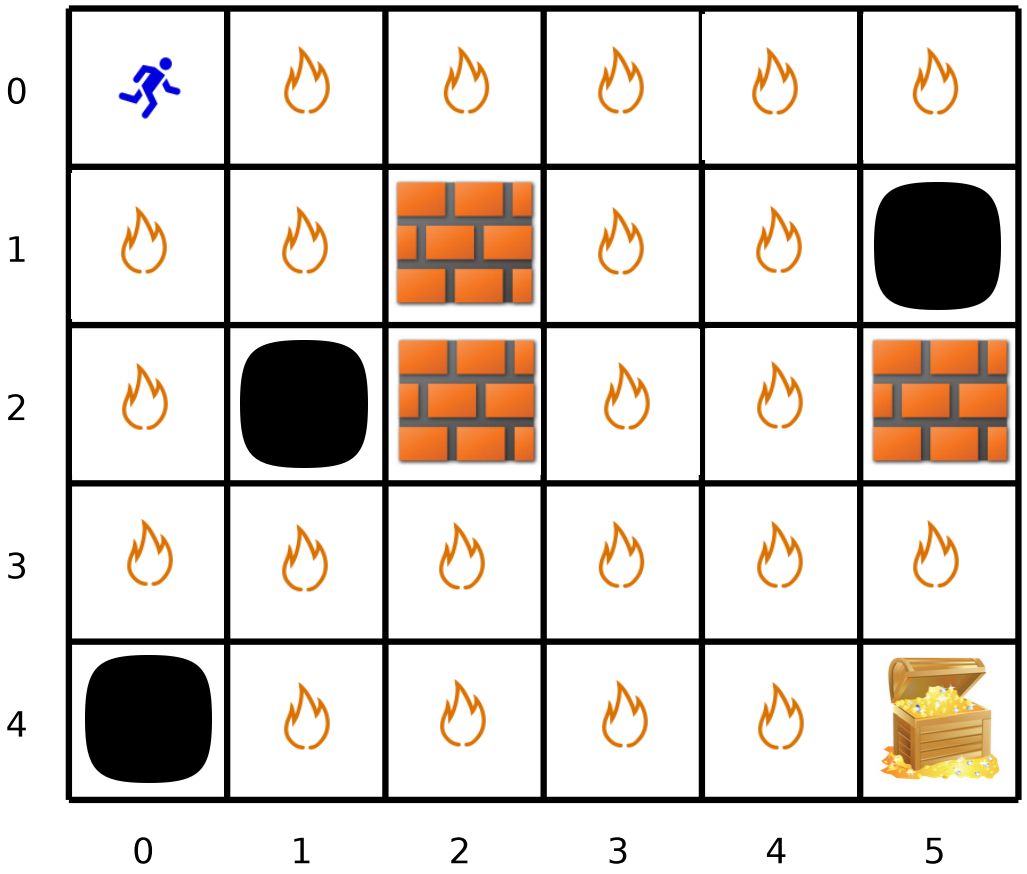

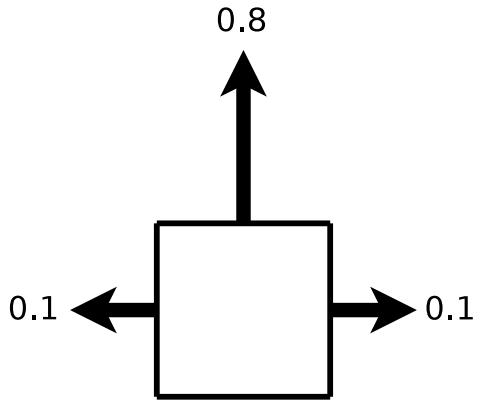

Action dynamics:

- 0.8 chosen direction
- 0.1 directions at 90° w.r.t. the chosen one

Rewards:

- -0.04 for each lava cell
- -10 for the black pit
- +10 for the gold treasure

Environment name in gym: "LavaFloorExam2018-v0"

Exercise 2 - Tasks

State which Reinforcement Learning algorithm you would choose to solve this problem. Motivate your choice and show the following charts:

1. rewards obtained during the learning phase for each algorithm
2. lengths of the learning episodes for each algorithm

Consider only the following versions for the RL algorithms

- Model-Based with initial policy of 0 - "L" for every state;
- Q-Learning $epsilon$-greedy;
- SARSA $epsilon$-greedy.

Note: the environment is stochastic! Execute your code multiple times in order to verify the consistency of the results

Exercise 2 - Parameters

Use the following parameters (where they apply):

1. number of episodes $= 1000$
2. max number of steps for each episode (ep_limit) $= 200$
3. max iterations for value iteration $= 200$
4. convergence threshold for value iteration $(\delta) = 1^{-3}$
5. discount factor $(\gamma) = 0.9$, learning rate $(\alpha) = 0.6$
6. probability of sub-optimal action $(\epsilon) = 0.1$
7. rolling window[2] $= 20$

[2] The smoothing function in [session3/rl.utils.rolling]