ULiège
Prof. Damien Ernst

INFO8003-1
Optimal decision making
for complex problems

Assignment 3

Deep Reinforcement Learning with Images for the Car on the Hill Problem

1 DOMAIN

We consider the *car on the hill* problem described in Assignment 2, with a simple modification: the agent does not directly have access to the state $x \in X$. Instead, the information to which the agent has access to is only made by the images created by the visualization routine implemented in Assignment 2[1].

2 DEEP Q-LEARNING (8 POINTS)

Build a deep neural network architecture which is able to approximate Q-functions. You may need to consider Convolutional Neural Networks. We also recommend you to read scientific articles related to deep learning applied to image analysis. Show and motivate your deep neural network architecture, e.g. by referring to relevant scientific literature.
You may use programming libraries such as Keras, TensorFlow or equivalent which provide routines to build and use your deep learning architectures.
Display $\widehat{Q}$ - obtained after training your deep neural network - in a 3D grid using the true state space with the state resolution described in Assignment 2. Estimate and display the expected return of $\widehat{\mu}^*$ derived from $\widehat{Q}$.

3 DEEQ Q-LEARNING VS FQI-TREES (5 POINTS)

Design an experimental protocol to compare your deep Q-learning implementation with (i) FQI with Extremely Randomized Trees where the input is also an image and (ii) FQI (with Extremely Randomized Trees) and Q-learning (with a neural network) and direct access to the state space. Explain the difference in results (e.g., by referring to the scientific literature).

[1] But it does know whether he is located in a terminal state

4 DEEQ Q-LEARNING VS DEEP Q-NETWORK (7 POINTS)

Design an experimental protocol to compare your deep Q-learning implementation with an existing implementation of Deep Q-Network[2] where inputs are images for both implementations. Explain the difference in results (e.g., by referring to the scientific literature).

[2] e.g., https://github.com/openai/baselines