ULiège
Prof. Damien Ernst

INFO8003-1
Optimal decision making
for complex problems

Assignment 3

Deep Reinforcement Learning with Images

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 $x \in X$, except if it is a terminal state. 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.

2 Deep Q-Learning (8 points)

Build a deep neural network architecture which is able to approximate efficiently 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 considering less complex architectures or referring to relevant scientific literature.
You may use (one of) the programming libraries Keras and TensorFlow which provide routines to build and train your deep learning architectures.
Display $\widehat{Q}$ in a 3D grid using the true state space with a reasonable resolution. Estimate and display the expected return of $\widehat{\mu}^*$ derived from $\widehat{Q}$.

3 DeeQ Q-Learning VS FQI-Trees (12 points)

Design an experiment 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 direct access to the state space.