
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

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. Motivate your deep neural network architecture.

The programming libraries Keras and TensorFlow provide routines to build and train your deep learning architectures.

Display \hat{Q} . Derive the policy $\hat{\mu}^*$ from \hat{Q} . Estimate and display the expected return of $\hat{\mu}^*$. 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.