

Check Points: Reinforcement Learning

If you are not able to answer some of these questions, first check the slides and lecture notes. If you still do not understand the question, send me an e-mail at Dario.Flozano@epfl.ch

- Mathematically define the Total Return
- Mathematically define the Q function
- Is the Q value of a state-action pair equal to the reward?
- Given a Q-function, how do you choose the action for a given state?
- What are the outputs of a Deep Q-Learning network?
- What are the limitations of Q-learning algorithms?
- What are the outputs of a Policy learning network for discrete action space?
- What are the outputs of a Policy learning network for continuous action space?