

Lecturer: Emanuele PANIZON (Area Science Park, Italy)

*Reinforcement learning concerns finding the optimal strategy of interaction with an environment to solve a task. Its theoretical framework is grounded in the concept of Markovian Decision Processes, and at the level of discrete and finite state/action spaces RL provides well defined guarantees of convergence. However, most cases of interest - such for example robotics - are to be found beyond, where states and/or actions are continuous. In this short tutorial we will see how deep learning and the Deterministic Policy Gradient Theorem can be combined to tackle such continuous tasks with success.*