



A mini-course on decision-making and active learning

Machine Learning has come of age and is found everywhere, in applications and basic sciences alike. While its practical successes are spectacular, understanding reasons (including for failure) are harder to cope with. It is therefore important to have a clear view of the fundamentals of the field to grasp its foundations and basic aspects. The scope of this mini-course is to provide such an introduction by stressing aspects related to active learning and decision-making before going to reinforcement learning. Examples from life sciences will illustrate various aspects of the problems.

Lecture 3

Friday, 24 May 2024 - h. 14:30

Aula Grassano (Department of Physics)

Prof. Massimo VERGASSOLA

Ecole Normale Supérieure, Paris

“An introduction to reinforcement learning.”

Abstract

Dynamic programming provides a systematic framework to approach problems of sequential decisions. Bellman Equation provides conditions for optimality and resulting policies of action for Markov Decision Processes. A crucial property of the equation is its contraction property, which ensures that iterative methods and successive approximations are guaranteed to converge. Generalizations to Partially Observable Decision Processes and the functional space of beliefs are then discussed. The methods provide the foundations for the more general (and difficult) cases where policies ought to be learned from experience, which we shall illustrate by an elementary introduction to Q-Learning.