

Reinforcement learning

ECTS : 4

Volume horaire : 27

Description du contenu de l'enseignement :

- Models: Markov decision processes (MDP), multiarmed bandits and other models
- Planning: finite and infinite horizon problems, the value function, Bellman equations, dynamic programming, value and policy iteration
- Basic learning tools: Monte Carlo methods, temporal-difference learning, policy gradient
- Probabilistic and statistical tools for RL: Bayesian approach, relative entropy and hypothesis testing, concentration inequalities
- Optimal exploration in multiarmed bandits: the explore vs exploit tradeoff, lower bounds, the UCB algorithm, Thompson sampling
- Extensions: Contextual bandits, optimal exploration for MDP

Compétence à acquérir :

Reinforcement Learning (RL) refers to scenarios where the learning algorithm operates in closed-loop, simultaneously using past data to adjust its decisions and taking actions that will influence future observations. Algorithms based on RL concepts are now commonly used in programmatic marketing on the web, robotics or in computer game playing. All models for RL share a common concern that in order to attain one's long-term optimality goals, it is necessary to reach a proper balance between exploration (discovery of yet uncertain behaviors) and exploitation (focusing on the actions that have produced the most relevant results so far).

The methods used in RL draw ideas from control, statistics and machine learning. This introductory course will provide the main methodological building blocks of RL, focussing on probabilistic methods in the case where both the set of possible actions and the state space of the system are finite. Some basic notions in probability theory are required to follow the course. The course will imply some work on simple implementations of the algorithms, assuming familiarity with Python.

Mode de contrôle des connaissances :

- Individual homework (in Python)
- Final exam

Bibliographie, lectures recommandées :

Bibliographie, lectures recommandées

- M. Puterman. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons, 1994.
- R. Sutton and A. Barto. Introduction to Reinforcement Learning. MIT Press, 1998.
- C. Szepesvari. Algorithms for Reinforcement Learning. Morgan & Claypool Publishers, 2010.
- T. Lattimore and C. Szepesvari. Bandit Algorithms. Cambridge University Press. 2019.