Introduction to Reinforcement Learning

Erwan Le Pennec
Erwan.Le-Pennec@polytechnique.edu

Fall 2024

Reinforcement Learning

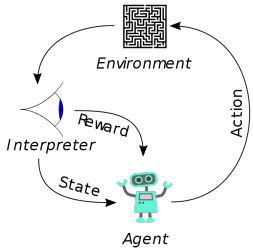

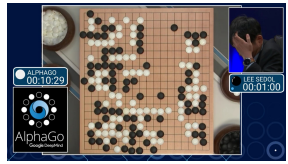

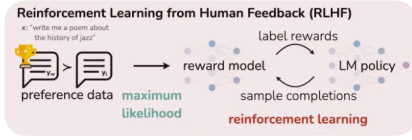

Art of learning to take sequential decisions while discovering the world!

- Intersection of Machine Learning, Operations Research and Control Theory.
- Strong mathematical foundations...
- But a lot of heuristics in the implementations.

- **Understanding the heuristics requires understanding the math!**

Source: Wikipedia/Shelly Palmer/Manish Chablani

(A Mathematically Structured) Introduction to RL

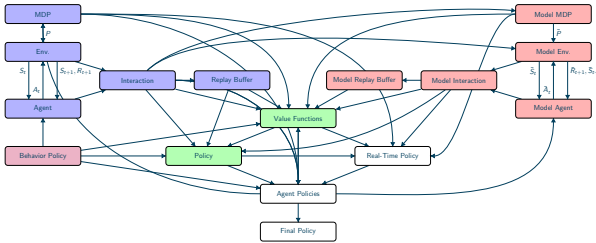

By the end of the course, you should

- understand the math of the reinforcement learning setting,
- have a comprehensive view of most RL techniques and their heuristics,
- be able to read research articles,
- be able to implement them.

Syllabus

- 6 lectures of 3h30 hours mixing several sources !

Outline

1. 08/11: Sequential Decisions, MDP and Policies
2. 15/11: Operations Research: Prediction and Planning
3. 22/11: Reinforcement Learning: Prediction and Planning in the Tabular Setting
4. 29/11: Reinforcement Learning: Advanced Techniques in the Tabular Setting
5. 06/12: Reinforcement Learning: Approximation of the Value Functions
6. 13/12: Reinforcement Learning: Policy Approach

Grade

- Article reading with ou without implementation

- The course *Stochastic approximation and reinforcement learning* by P. Bianchi focus on stochastic approximation, a central tool that will only be used (and not proved) in my course.

References

R. Sutton and A. Barto.
Reinforcement Learning, an Introduction (2nd ed.)
MIT Press, 2018

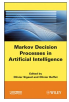

O. Sigaud and O. Buffet.
Markov Decision Processes in Artificial Intelligence.
Wiley, 2010

M. Puterman.
Markov Decision Processes. Discrete Stochastic Dynamic Programming.
Wiley, 2005

D. Bertsekas and J. Tsitsiklis.
Neuro-Dynamic Programming.
Athena Scientific, 1996

W. Powell.
Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions.
Wiley, 2022

S. Meyn.
Control Systems and Reinforcement Learning.
Cambridge University Press, 2022

V. Borkar.
Stochastic Approximation: A Dynamical Systems Viewpoint.
Springer, 2008

T. Lattimore and Cs. Szepesvári.
Bandit Algorithms.
Cambridge University Press, 2020