

facebook

Artificial Intelligence Research

Reinforcement Learning

shorturl.at/fpJZ5

Alessandro Lazaric

Facebook AI Research


Acknowledgments

Most of this material is extracted from a course held by A. Lazaric and M. Pirotta at the Master MVA

Why: Important Problems

Why: Important Problems

- Autonomous robotics



Why: Important Problems

- Autonomous robotics



- *Elder care*

Why: Important Problems


- Autonomous robotics



- *Elder care*
- *Exploration of dangerous environments*

Why: Important Problems

- Autonomous robotics



- *Elder care*
- *Exploration of dangerous environments*
- *Robotics for entertainment*

Why: Important Problems

- Robotics
- Finance



Why: Important Problems

- Robotics
- Finance



- *Trading execution algorithms*

Why: Important Problems

- Robotics
- Finance



- *Trading execution algorithms*
- *Portfolio management*

Why: Important Problems

- Robotics
- Finance



- *Trading execution algorithms*
- *Portfolio management*
- *Option pricing*

Why: Important Problems

- Robotics
- Finance
- Resource management



Why: Important Problems

- Robotics
- Finance
- Resource management



- *Energy grid integration*

Why: Important Problems

- Robotics
- Finance
- Resource management



- *Energy grid integration*
- *Energy market regulation*

Why: Important Problems

- Robotics
- Finance
- Resource management



- *Energy grid integration*
- *Energy market regulation*
- *Energy production management*

Why: Important Problems

- Robotics
- Finance
- Resource management
- Recommender systems



Why: Important Problems

- Robotics
- Finance
- Resource management
- Recommender systems



- *Web advertising*

Why: Important Problems

- Robotics
- Finance
- Resource management
- Recommender systems



- Web advertising
- Product recommendation

Why: Important Problems


- Robotics
- Finance
- Resource management
- Recommender systems



- Web advertising
- Product recommendation
- MOOCs / ITS


Why: Important Problems

- Robotics
- Finance
- Resource management
- Recommender systems
- Games



Why: Important Problems


- Robotics
- Finance
- Resource management
- Recommender systems
- Games



- *Board games*


Why: Important Problems

- Robotics
- Finance
- Resource management
- Recommender systems
- Games



- *Board games*
- *Computer games*

What: Reinforcement Learning



“**Reinforcement learning** is learning how to map states to actions so as to **maximize** a numerical **reward** signal in an unknown and **uncertain** environment.

In the most interesting and challenging cases, **actions** affect not only the immediate reward but also the **next situation** and all subsequent rewards (**delayed reward**).

The agent is not told which actions to take but it must discover which actions yield the most reward by trying them (**trial-and-error**).”

— Sutton and Barto [1998]

What: the Structure of the Course

How to *model* an RL problem *[30min + Hands-on 15min]*

What: the Structure of the Course

How to *model* an RL problem *[30min + Hands-on 15min]*

How to solve *exactly* an RL problem *[30min]*

What: the Structure of the Course

How to *model* an RL problem *[30min + Hands-on 15min]*

How to solve *exactly* an RL problem *[30min]*

How to solve *incrementally* an RL problem *[55min]*

What: the Structure of the Course

How to *model* an RL problem *[30min + Hands-on 15min]*

How to solve *exactly* an RL problem *[30min]*

How to solve *incrementally* an RL problem *[55min]*

How to solve *approximately* an RL problem *[55min]*

What: the Structure of the Course

How to *model* an RL problem *[30min + Hands-on 15min]*

How to solve *exactly* an RL problem *[30min]*

How to solve *incrementally* an RL problem *[55min]*

How to solve *approximately* an RL problem *[55min]*

How to trade off *exploration and exploitation* *[55min]*

What: the Structure of the Course

How to *model* an RL problem *[30min + Hands-on 15min]*

How to solve *exactly* an RL problem *[30min]*

How to solve *incrementally* an RL problem *[55min]*

How to solve *approximately* an RL problem *[55min]*

How to trade off *exploration and exploitation* *[55min]*

Hands-on sessions *[40min]*

- Inventory problem [\[link\]](#)
- RL-Sim [\[link\]](#)
- Q-learning [\[link\]](#)
- REINFORCE [\[link\]](#)
- A2C [\[link\]](#)
- Bandit [\[link\]](#)



Thank you!

facebook

Artificial Intelligence Research



. \ |

Bibliography

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. Bradford Book, MIT Press, Cambridge, MA, 1998.