

Solving the Tragedy of the Commons Using Multi-Agent Reinforcement Learning

Supervisor: Panayiotis Danassis

February 1, 2019

1 Project Overview & Goal Description

The focus of this project is to acclimate the student in the field of multi-agent reinforcement learning. Reinforcement learning [1] is based on the concept of learning through the interactions with the environment. An agent takes an action, observes some feedback from the environment, and updates his policy so as to maximize some notion of cumulative reward. The most eminent example of such an algorithm is Q-learning [2] which solves Bellman's optimality equation [3] using an iterative approximation procedure. A detailed taxonomy of multi-agent reinforcement learning algorithms can be found in [4].

In recent years, reinforcement learning has been reinvigorated through multiple exciting experimental results (AlphaGo Zero [5], mastering Atari games [6], modeling social dilemmas for rational agents [7], etc.). This project aims to explore various multi-agent reinforcement learning techniques and apply them in the common pool resource (CPR) appropriation domain.

2 Project Steps

- Get acquainted with the state-of-the-art reinforcement learning techniques.
- Model the world and define the state representation.
- Implement one or more multi-agent reinforcement learning algorithms, and an evaluation platform.
- Empirically evaluate the properties of the implemented techniques.

3 Required Skills

Good programming skills are required.

General knowledge of reinforcement learning & neural networks is a big plus.

Being passionate about the topic and good English skills are a must.

References

- [1] Wikipedia, “Reinforcement learning — wikipedia, the free encyclopedia,” 2017, [Online; accessed 20-November-2017]. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Reinforcement_learning&oldid=810477823
- [2] C. J. C. H. Watkins and P. Dayan, “Q-learning,” *Machine Learning*, vol. 8, no. 3, pp. 279–292, May 1992. [Online]. Available: <https://doi.org/10.1007/BF00992698>
- [3] R. Bellman, *Dynamic programming*. Courier Corporation, 2013.
- [4] L. Busoniu, R. Babuska, and B. De Schutter, “A comprehensive survey of multiagent reinforcement learning,” *Trans. Sys. Man Cyber Part C*, vol. 38, no. 2, pp. 156–172, Mar. 2008. [Online]. Available: <http://dx.doi.org/10.1109/TSMCC.2007.913919>
- [5] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, “Mastering the game of go without human knowledge,” *Nature*, vol. 550, no. 7676, pp. 354–359, 2017.
- [6] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, “Playing atari with deep reinforcement learning,” *arXiv preprint arXiv:1312.5602*, 2013.
- [7] J. Z. Leibo, V. Zambaldi, M. Lanctot, J. Marecki, and T. Graepel, “Multi-agent reinforcement learning in sequential social dilemmas,” in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2017, pp. 464–473.