




SA/MA:

Byzantine Reinforcement Learning

Reinforcement learning is a powerful tool to investigate environments that are described by Markov Decision Processes (MDPs). Given an action of the agent in some environment, the agent changes to a new state of the system according to the underlying MDP. The new state can give either positive or negative reward to the agent. The aim of the agent is to maximize its total reward.

Just as every other system, reinforcement learning algorithms are prone to failures. The worst possible failures that can happen in a system are the Byzantine failures, that is, agents who are controlling parts of the system are deliberately disturbing the decision process.

In this work we want to investigate how Byzantine behavior affects reinforcement learning algorithms. We will assume that a Byzantine agent might control any step in the protocol - the action, observation, or they might even embody the agents themselves. Your task would be to simulate Byzantine behavior and derive protocols which are robust to this kind of malicious behavior.



Requirements: An interested student should be familiar with Byzantine behavior, have knowledge in Deep Learning or a solid background in Machine Learning. Implementation experience is an advantage. The student should be able to work independently!

Interested? Please contact us for more details!

Contacts

- Darya Melnyk: dmelnyk@ethz.ch, ETZ G97
- Oliver Richter: richtero@ethz.ch, ETZ H93