



Prof. R. Wattenhofer

Topics on Distributed Reinforcement Learning

Reinforcement learning (RL) has recently been applied to many real-world decision-making problems such as gaming, robotics, and yes, Large Language Models (LLMs). However, despite its impressive performances in simulation, RL often suffers from poor sample efficiency, which hinders its success in real-world applications. For example, when RL is applied to provide clinical decision support, its performance is limited by the number (i.e., sample size) of admission records possessed by a hospital, which cannot be synthetically generated. A natural solution is to encourage multiple RL agents (i.e., different hospitals) to share their trajectories, to collectively build a better decision-making policy that one single agent can not obtain by itself.



Figure 1: Illustration of our federated policy gradient framework [1]

However, in many applications, raw RL trajectories contain sensitive information (e.g., the medical records contain sensitive information about patients), and thus sharing them is prohibited. How can we collectively learn a better RL policy from distributed agents without sharing their trajectories? How do we ensure that the privacy of agents is protected? What happens if adversarial agents are present? In this project, we aim to leverage distributed computing to develop innovative techniques for distributed and decentralized reinforcement learning, and address many interesting open problems, such as privacy and robustness in RL.

Interested? Please contact us for more details!

Contact

• Flint Xiaofeng Fan: xiafan@ethz.ch, ETZ G97

References

 Xiaofeng Fan, Yining Ma, Zhongxiang Dai, Wei Jing, Cheston Tan, and Bryan Kian Hsiang Low. Fault-tolerant federated reinforcement learning with theoretical guarantee. Advances in neural information processing systems, 34:1007–1021, 2021.