Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Federated Reinforcement Learning from Distributed Expert Knowledge

Modern learning systems increasingly operate across distributed agents, such as robots, autonomous systems, user devices, or organizations, where useful expert knowledge is naturally generated and stored locally. Pooling such data at a central location is often undesirable or infeasible because of privacy, ownership, communication, and operational constraints.

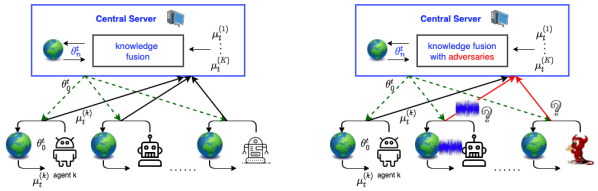

Figure 1: Illustration of our federated policy gradient framework [1]

Federated learning provides a natural framework for collaborative learning without directly sharing raw local data. However, extending this paradigm to sequential decision-making and reinforcement learning introduces challenges that do not arise in standard supervised federated learning. In particular, different agents may interact with different environments, observe different behaviors, and possess highly heterogeneous local experience.

This project will investigate how distributed agents can collaboratively learn from expert knowledge while keeping their local experience decentralized. The student will study learning algorithms and experimental methodologies at the intersection of federated learning, reinforcement learning, and learning from demonstrations, with particular emphasis on heterogeneous agents and generalization across environments.

Interested? Please contact us for more details!

Contact

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

References

[1] 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.