Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Meta-Federated Reinforcement Learning: Learning to Learn from Distributed Expert Knowledge

Federated Reinforcement learning (FedRL) has demonstrated impressive capabilities in collaborative policy optimization across distributed agents while preserving data privacy. Current FedRL approaches excel at aggregating policy gradients and value functions from multiple agents to achieve faster convergence than independent learning. However, FedRL suffers from a critical limitation: inability to leverage heterogeneous learning experiences across agents. While agents can share model parameters, they cannot effectively share their diverse learning strategies, adaptation techniques, or task-specific optimization approaches. This leads to suboptimal learning efficiency, especially when agents face different environments, have varying data distributions, or possess different local capabilities.

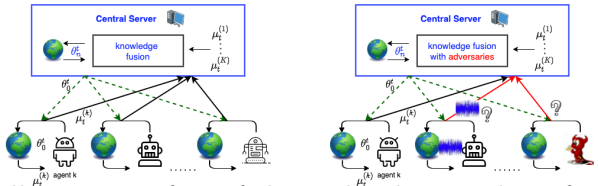

Figure 1: Illustration of our federated policy gradient framework [1]

Meta-learning, on the other hand, excels at learning optimal learning strategies and rapid adaptation across diverse tasks. Meta-learning algorithms like MAML have shown remarkable success in learning generalizable optimization procedures that can quickly adapt to new environments and tasks. Meta-learning is specifically designed for learning how to learn efficiently and transferring learning strategies across heterogeneous scenarios—exactly what FedRL lacks. We propose Meta-Federated Reinforcement Learning (Meta-FedRL), which combines the collaborative benefits of FedRL with the learning strategy transfer capabilities of meta-learning. In Meta-FedRL, agents collaboratively learn optimal learning strategies, adaptation techniques, and optimization approaches that can be effectively applied across heterogeneous environments and tasks. This enables better utilization of diverse agent experiences while maintaining the privacy and collaboration benefits of federated learning.

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.