



## Distributed Reinforcement Learning for Optimization (ETH-MIT Joint Project)

Reinforcement learning is increasingly used as a general-purpose approach to solve sequential decision-making problems that can be cast as optimization. Compared to classical solvers, RL is particularly attractive when objectives are noisy, non-stationary, or only observable through interaction, and when the solution must adapt online under changing conditions.

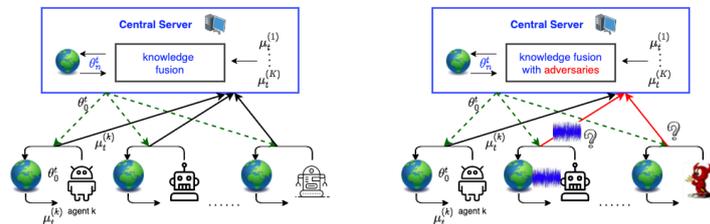


Figure 1: Illustration of our federated policy gradient framework [1]

Many real-world optimization problems, however, are distributed by nature: data, constraints, and feedback signals are generated across multiple sources (e.g., machines, users, sensors, institutions), each with only partial visibility. This motivates studying distributed RL: how to learn policies or optimization strategies when experience is collected across multiple agents/clients, communication is limited, and the system must remain robust to heterogeneity, partial observability, and even potential adversarial attacks, as illustrated in Figure [1]. The goal of this thesis is to explore how distributed RL can be used as a principled, scalable tool for optimization in such settings, and to develop reusable insights and prototypes that generalize beyond a single application.

We are looking for an exceptional and highly motivated student to join this research project, in official collaboration with MIT (Dr. Yining Ma and Prof. Cathy Wu). You will conduct your thesis at ETH, while also participating in an MIT summer research camp as part of the joint program. This is an opportunity to work at the intersection of distributed systems, optimization, and modern RL, with close mentorship and a clear path toward a publishable outcome.

**Interested? Please contact us with your transcript and CVs for more details!**

### Contact

- Flint Xiaofeng Fan: [xiafan@ethz.ch](mailto: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.