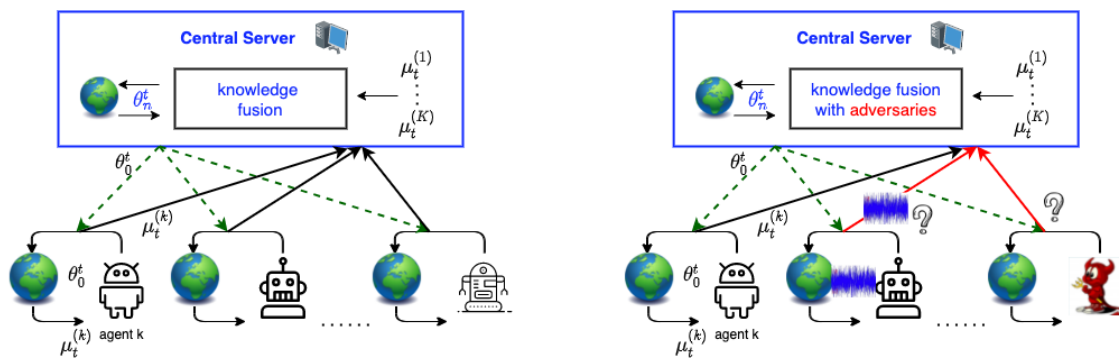




Distributed and Federated Reinforcement Learning

Reinforcement learning (RL) has recently been applied to many real-world decision-making problems such as gaming, robotics, healthcare, etc. 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.



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 federated reinforcement learning, and address many interesting open problems¹, such as privacy and robustness in RL.

Requirements: Knowledge in reinforcement learning or a solid background in machine learning. Experience with PyTorch (or TensorFlow) and OpenAI Gym is an advantage. We will have weekly meetings to address questions, discuss progress and think about future ideas. Motivation to conduct independent research is a strong plus, as this project is expected to lead to a conference publication.

Interested? Please contact us for more details!

Contact

- Flint Xiaofeng Fan: xiafan@ethz.ch, ETZ G93
- Florian Grötschla: fgroetschla@ethz.ch, ETZ G63

¹The topic details can be discussed and specified with the student.