SCIENCE HUB FOR HUMANITY AND ARTIFICIAL INTELLIGENCE

UCLA amazon

UCLA Samueli Computer Science

Multi-Agent Reinforcement Learning: Asynchronous Communication, Robustness and Privacy

Jiafan He

Advisor: Prof. Quanquan Gu

01/25/2024

Reinforcement Learning

Sequential Decision-Making Problems

- Interact with environment
- Sequence of decisions repeatedly
- Adjust the policy based on the past information

Goal: maximize the cumulative reward

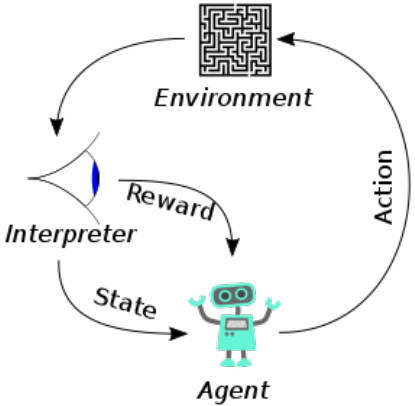

Image Credit: https://commons.wikimedia.org/w/index.php?curid=57895741

Cooperative Multi-Agent Reinforcement Learning

Various of successful applications: game, autonomous driving, dialogue system

Large application scale

- Enormous number of states in practice (more than 10^{100})
- Massive # of samples (more than 10^{6})
- Single agent: all data in central server

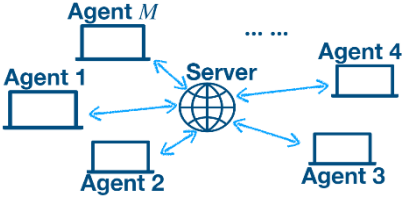

Cooperative Multi-Agent RL (Federated RL)

- Keep private data decentralized
- Model learning in agent side
- Cooperative learning (communication)

Image Credit: https://en.wikipedia.org/wiki/Go_(game)

Topic 1: Asynchronous Communication

Synchronous Environment

- Agents: full participation
- Server: global synchronization
- Impractical (offline, unavailable agent)

My research: Asynchronous Communication Protocol

- Agent : decide whether or not to participate
- Communication: independent for different agent
- Efficiency: communication, performance

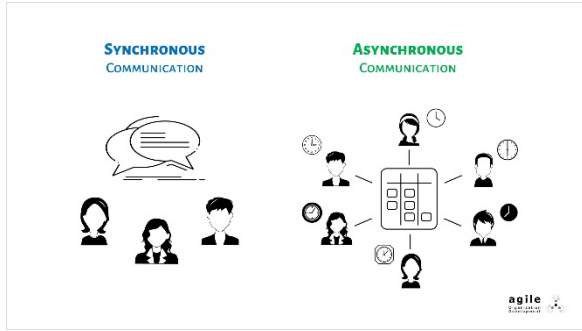

Image Credit: https://agile-od.com/mmdojo/11303/remote-and-hybrid-teamwork-success-asynchronous-communication-is-the-key

Topic 2: Robust Reinforcement Learning

Adversarial attack for Communication

- Channel: Influenced by adversary
- Agent: Provided adversarial sample
- Identify the trustworthiness: expensive, challenging

My research: Learning the model

- When the data may not accurate
- Efficiency: performance (low attack level)

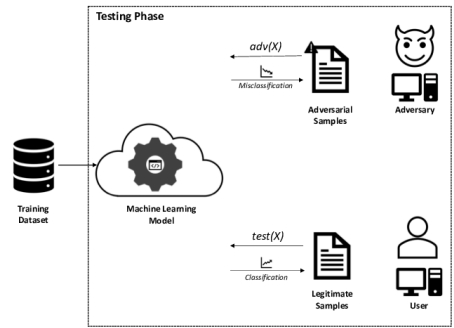

Image Credit: https:Diagram-of-the-model-evasion-attack-against-a-trained-machine-learning-model-during-the//www.researchgate.net/figure/_fig1_340998746

Topic 3: Differentially Private Reinforcement Learning

Multi-Agent Reinforcement Learning

- Update the model with personal data
- Data: sensitive personal information
- Danger: potential privacy leakage

My research: safe data sharing and privacy guarantee

- Differential Privacy Guarantee
- Efficiency: performance

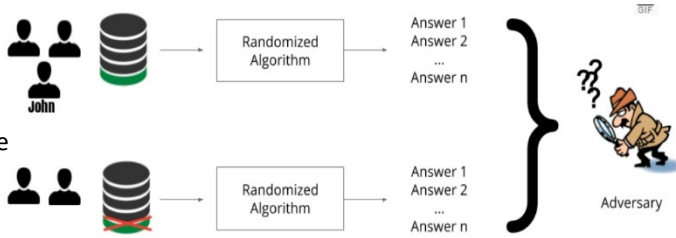

Image Credit: https://blog.openmined.org/maintaining-privacy-in-medical-data-with-differential-privacy

SCIENCE HUB FOR HUMANITY
AND ARTIFICIAL INTELLIGENCE
UCLA amazon

UCLA Samueli Computer Science

Thanks