Hierarchical Multi-Agent Deep Reinforcement Learning

Summary: While current deep reinforcement learning (DRL) systems can achieve super-human performance in various domains, they possess problems that prohibit their real-world applicability. Hierarchical reinforcement learning is a promising technique that has gained a resurgent attention from DRL community in recent years. It promises a solution to some long-standing problems of DRL: It allows crediting long-term actions with rewards in the future, more efficient learning by decomposing a task into its subtasks, and makes possible transferring what agents learn between different tasks that share similar structures – with the goal of creating agents that are more adaptive, more efficient, and better suited to a variety of tasks.

This idea of hierarchical organization and problem decomposition also has the potential to mitigate problems encountered by cooperative multi-agent systems as well. We know that from the natural organization of organisms and groups of intelligent biological entities, which evolved to be inherently hierarchical structures that allow complex and adaptive behaviors. In this regard, recently, Hierarchical RL were applied to multi-agent systems. These methods aim to create a hierarchical organization structure between multiple reinforcement-learning agents to realize efficient, adaptive organization and collaboration.

This project will begin by exploring the novel hierarchical multi-agent reinforcement learning (MARL) methods implemented in the literature in simple scenarios. We will move forward by applying the same methods in more complex problems, with different hierarchical architectures and different interaction possibilities (inter-agent communication, reward shaping, etc.). The properties of the resulting system will be analyzed: Not only in terms of performance and learning compared to non-hierarchical architectures, but also aspects including controllability, transfer of abilities between tasks, specialization, adaptability, and explainability. (The exact direction of the projects will depend on the initial results obtained.)

Who is this project aimed for?

- Students with interest in machine learning or reinforcement learning techniques; their applications, and their limitations,
- Students with interest in multi-agent systems, game theory, complexity, emergent organization, and related topics.
- **Proposed as:** Semester or Master project

Requirements:

- Good understanding of basic machine learning concepts
- Familiarity with reinforcement learning framework (literacy about basic concepts is sufficient)
- Familiarity with Python (the more experience the better, but not obligatory)
- Basic game theory understanding would help, but not obligatory

Note: In case of sufficiently early contact, the student can familiarize himself/herself with the missing required areas before the beginning of the project.

References with brief explanations:

Barto, A. G., & Mahadevan, S. (2003). Recent advances in hierarchical reinforcement learning. Discrete event dynamic systems, 13(1-2), 41-77.
An overview of Hierarchical RL and its paradigms, but it is old and long. You can read if you are patient and really want to learn the details of paradigms; otherwise, move to the next one.

Yannis Flet-Berliac, "The Promise of Hierarchical Reinforcement Learning", The Gradient, 2019.
An easy-to-read, popular-language introduction to Hierarchical RL and why it is important for the future of AI.

Ahilan, S., & Dayan, P. (2019). Feudal multi-agent hierarchies for cooperative reinforcement learning. arXiv preprint arXiv:1901.08492.
An example hierarchical multi-agent reinforcement learning application. This paper will probably be our starting point.

Kumar, S., Shah, P., Hakkani-Tur, D., & Heck, L. (2017). Federated control with hierarchical multi-agent deep reinforcement learning. arXiv preprint arXiv:1712.08266.
Another example of hierarchical MARL.

Vezhnevets, A. S., Wu, Y., Leblond, R., & Leibo, J. Z. (2019). Options as responses: Grounding behavioural hierarchies in multi-agent RL. arXiv preprint arXiv:1906.01470.
Yet another example of hierarchical MARL by Google DeepMind.

Contact person
Zeki Doruk Erden | *zeki.erden@epfl.ch*