Duplication-Differentiation Based Hierarchical Reinforcement Learning

Summary: While current reinforcement learning (RL) systems can achieve good performance in various domains, they possess two major interconnected drawbacks for applicability: First, they tend to specialize within the task domain that they are trained in, and perform poorly even in only slightly different tasks, usually requiring a whole training process from scratch within the new task domain. Second, mainstream RL systems require large amounts of experience to learn, to the extent that it becomes practically impossible to train in real-world scenarios.

A well-structured hierarchical policy can, in principle, help alleviate these limitations of RL. However, structure of the hierarchical policy in such systems are often specified by the designer and hence limited in terms of their learnability and adaptability. As a potential solution, mechanisms that structure complex organisms can help us design hierarchies that can be developed flexibly without designer intervention. One of such mechanisms is duplication and differentiation combined with switching, which at higher scales can develop a fractal-like graph that can represent a hierarchical structure.

The aim of this project is, taking into account this mechanism, to experiment with and develop a simple framework for learning the hierarchical structures, in principle with unbounded depth and width, using a duplication-differentiation-based mechanism. Expectations from the student include using a machine learning framework of his choice to implement this proposed mechanism, experimenting with various parameters and processes in its design space, and to devise and perform experiments to evaluate its performance.

Requirements:

- Interest in machine learning or reinforcement learning techniques
- Familiarity with reinforcement learning framework (literacy about basic concepts)
- Familiarity with Python (the more experience the better, but not obligatory)
- **Proposed as:** Semester Project.

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:

Yannis Flet-Berliac, "The Promise of Hierarchical Reinforcement Learning", The Gradient, 2019.
A brief overview of hierarchical reinforcement learning.

Contact person

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