Forward-Forward algorithm evaluation in Deep Reinforcement Learning

Laboratory:

Swiss Data Science Center

Type:

Semester Project

Description:

Reinforcement learning (RL) is a branch of machine learning that deals with learning how to make decisions in an environment through trial and error. Deep reinforcement learning (DRL) combines RL with deep neural networks to enable learning from high-dimensional state spaces and solve more complex problems.

The training process in DRL algorithms is essential for obtaining good performance, which involves updating the parameters of the neural network to maximize the reward signal received from the environment. Backpropagation is the most widely used method for training neural networks, but it has several drawbacks, such as the vanishing gradient problem and high computational cost. Alternative methods, such as Evolution Strategies, have been shown to be a valid alternative to backpropagation in simple environments.

The goal of this project is to implement Geoffrey Hinton's recently proposed "Forward-Forward" algorithm as an alternative to backpropagation in the context of DRL. Although the algorithm has shown promising results in calibrating classification networks on the MNIST dataset, its application in DRL has not been widely explored.

In this project, the student will train an agent on environments provided by OpenAI gym (games) and CityLearn (energy management), and evaluate the potential of the "Forward-Forward" algorithm as a new method for training neural networks in DRL by comparing its performance and convergence speed to those achieved by backpropagation.

Goals/Benefits:

- Development of the Forward-Forward algorithm
- Exposure to several DRL algorithms
- Experience with DRL

Prerequisites:

- Python (Advanced)
- Reinforcement Learning (Basic)
- Experience with Pytorch

Deliverables:

- Well-documented code

- Written report and oral presentation

References:

[1] Geoffrey Hinton. "The Forward-Forward Algorithm: Some Preliminary Investigations." arXiv:2212.13345v1 (2022)

[2] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, Ilya Sutskever. "Evolution Strategies as a Scalable Alternative to Reinforcement Learning." arXiv:1703.03864v2 (2017)

Contact:

giulio.romanelli@epfl.ch