

Ultra-low Power Evolutionary Reinforcement Learning PhDs

Abstract

This project will recruit 2 PhDs in the domain of frugal Machine Learning (ML). The aim of the PhDs is to propose full-stack methods and open-source tools to train and infer ultra-lightweight AIs, by extending, implementing, and optimizing a new ML technique that relies on the light-by-construction and adaptive Tangled Program Graph (TPG) model. The two PhD students will work in tandem to optimize energy during the whole life cycle, from training to inference.

Keywords: **Artificial Intelligence**, **Reinforcement Learning**, **Ultra Low Power**, **Tangled Program Graphs**, **Open-Source**.

Context

Reinforcement Learning with Tangled Program Graphs

Reinforcement Learning (RL) is a branch of Machine Learning (ML) techniques where an autonomous artificial intelligence learns how to interact with its environment. As depicted in the figure below, using a trial and error mechanism, the artificial intelligence learns from its own experience, by interacting with the environment, and by getting rewards for each attempt. The purpose of the artificial intelligence is to maximize this reward.




Fig 1. Reinforcement Learning

Proposed in 2017, Tangled Program Graphs (TPGs) are a new way to power reinforcement learning AI, based on evolutionary concepts. The main strength of TPGs, compared to state-of-the-art deep learning-based techniques, is the lightweightness of their model, which confers them a low computational complexity, and very high performance on regular desktop CPU. Compared to deep-learning at equivalent accuracy, TPGs execute 1000x faster with 100x less memory.

GEGELATI

Generic Evolvable Graphs for Efficient Learning of Artificial Tangled Intelligence



GEGELATI [dʒedʒelati] is a fresh open-source reinforcement learning framework for training artificial intelligence based on TPGs. The purpose of this framework, developed as a C++ shared library, is to make it as easy and as fast as possible and to train an agent on a new learning environment. The C++ library is developed to be portable, fully documented, and thoroughly unit tested to ensure its maintainability. GEGELATI is developed at the Institut d'Electronique et des Technologies du numéRique (IETR).

Video presentation of GEGELATI: <https://youtu.be/t0Ta5Vo5h7s>

Objectives

The main scientific objectives pursued by the PhD students are to:

- Integrate energy optimizations at the core of the TPG training process: overall energy optimization will minimize the energy consumption of the computing system for TPG training and inference, as well as, when relevant, the energy consumption of physical actuators of the controlled cyber-physical systems.
- Extend the TPG learning capabilities: on top of improving the TPG efficiency on existing environments, model extensions will unlock new types of learning environments like continuous action spaces or non-reinforcement learning environments required for extending TPGs to real-world use cases.
- Propose highly-efficient implementation techniques: in order to find the most suited hardware platform for TPGs, implementation will be pursued and compared on multiple state-of-the-art hardware. Implementation will be studied for both efficient training, and inference on battery-powered ULP devices and reconfigurable devices for nanoseconds reaction time on the factory floor.

The foreseen focus of the two PhDs are:

- PhD 1: Energy-aware low-complexity reinforcement learning.
- PhD 2: HW-SW co-optimization for ultra-low power reinforcement learning.

Characteristics

- Duration: 3 years
- Start: Oct. 2023
- Stipend: 2300€ per month before tax
- Funded by: ANR FOUTICS Research Project

Location

VAADEr team, IETR laboratory - INSA Rennes,
20 Avenue des Buttes de Coësmes
35700 Rennes - France

Supervisors

- Karol Desnos (IETR, Equipe Vaader, Rennes) - kdesnos@insa-rennes.fr
- Mickaël Dardaillon (IETR, Equipe Vaader, Rennes) - mdardail@insa-rennes.fr

Applications

Send resume and application letter to kdesnos@insa-rennes.fr