SpikeRL

FIND OUT MORE AT https://github.com/icl-utk-edu/spike-rl

INTRODUCTION

In this modern era of AI revolution, there have been massive and rapid investments in large-scale AI systems. This proliferation of AI brings new optimization challenges for sustainability without losing scalability and performance in real world applications. In this work, a novel model called SpikeRL is introduced, which is a scalable and efficient framework for Deep Reinforcement Learning (DeepRL) based Spiking Neural Networks (SNNs) for complex continuous control tasks.

SpikeRL ARCHITECTURE

The SpikeRL framework consists of three major components.

I. A DeepRL-based SNN model utilizing population encoding and decoding which optimizes the representation capacity and computational efficiency of the network.

II. Distributed training with both Message Passing Interface (MPI) and NVIDIA Collective Communications Library (NCCL) backend, implemented through the PyTorch Distributed package.

III. Mixed-precision training with BFLOAT16 for optimizing the parameter updates during model training.

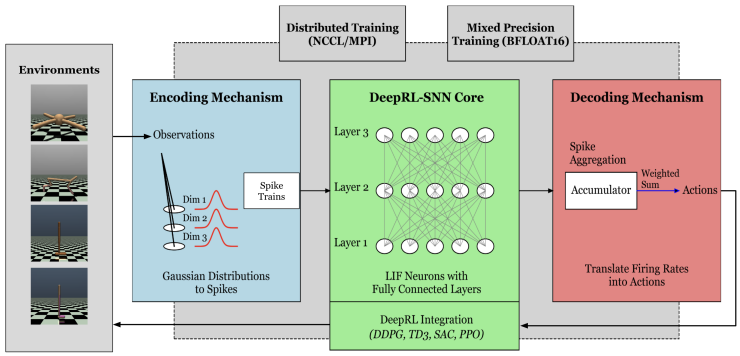

SpikeRL Architecture

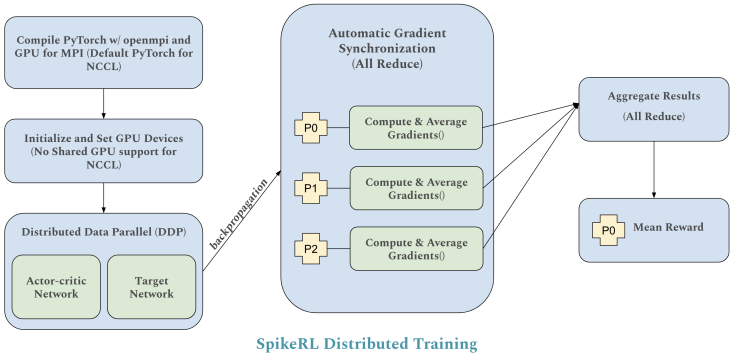

SpikeRL Distributed Training

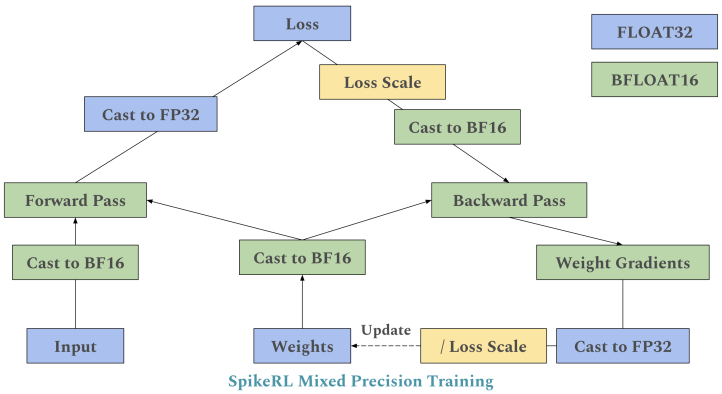

SpikeRL Mixed Precision Training

EXPERIMENTAL RESULTS

The experiments were conducted on ICL's cluster using Guyot's 8 Nvidia A100 80GB GPUs which are well suited for mixed precision and distributed training tasks in deep learning. SpikeRL was tested in various challenging continuous control environments from the MuJoCo physics simulator such as the Ant-v4, HalfCheetah-v4, Hopper-v4, and Wallker2D-v4. Results presented here are for the HalfCheetah-v4 environment.

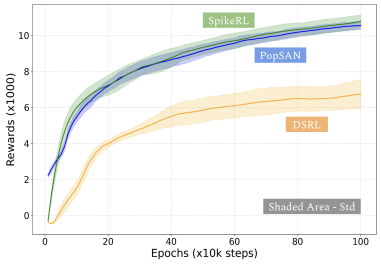

HalfCheetah-v4: Mean and Std (5 runs)

Performance results comparison demonstrated that SpikeRL achieved 3.72% more average rewards over the Population Coded Spiking Actor Network (PopSAN) and 45.47% more than the Deep Spiking Reinforcement Learning (DSRL) method.

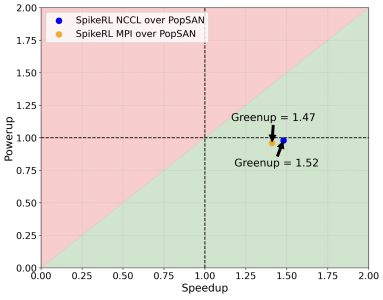

HalfCheetah-v4 GPS-UP Quadrant Graph

The GPS-UP quadrant graph shows the energy efficiency of SpikeRL over PopSAN using the GreenUp metric which is measured as Greenup = Speedup / Powerup. SpikeRL achieved on average 1.48x Speedup, 0.98x Powerup, and 1.52x Greenup over PopSAN.

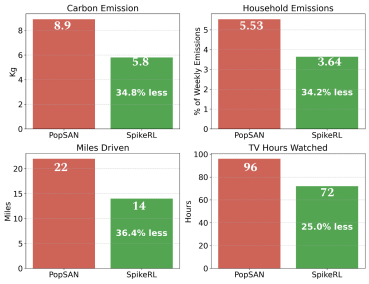

HalfCheetah-v4 Carbon Footprint Analysis

The energy efficiency of SpikeRL shows a direct impact in the carbon footprint analysis where SpikeRL reduces carbon emission by 34.8% compared to PopSAN which is equivalent to 34.2% less weekly household emissions, 36.4% less miles driven of a gasoline car, and 25% less tv hours watched. The distributed training along with mixed-precision optimization makes SpikeRL more energy efficient than other state-of-the-art SNN methods.