RL-X: A Deep Reinforcement Learning Libray (not only) for RoboCup

Nico Bohlinger, Klaus Dorer

Hochschule Offenburg, Germany, contact: klaus.dorer@hs-offenburg.de

Over the last years, many frameworks for Deep Reinforcement Learning (DRL) have been developed. Each of them focusing on different kinds of algorithms, implementation complexities or underlying Deep Learning frameworks. Stable-Baselines3 (SB3) [1] is a framework often used by practitioners, as it is easy to use and supports the most important model-free algorithms. Unfortunately, SB3 comes with a complex code structure, which makes it difficult to extend with new algorithms or modify existing ones for DRL researchers. SB3 is based on PyTorch, which is a well established Deep Learning framework in the Machine Learning community, but is not the computationally fastest one available.

To combat these issues and support the development of DRL algorithms (not only) for RoboCup, we present the RL-X library. RL-X provides a flexible and easy-to-extend codebase with self-contained single directory implementations. Some algorithms are implemented in PyTorch and TorchScript but all of them have a JAX-based version with the Flax framework. JAX enables huge computational speedups with the help of XLA, JIT-compilation and vectorization for GPU and TPU.

We showcase RL-X's capabilities by benchmarking it in a custom RoboCup Soccer Simulation 3D environment and the Humanoid MuJoCo environment against SB3. The results show that RL-X is faster than SB3 in all environment and algorithm combinations. Especially the JAX-based implementation shows up to 4.5x speedups in fast environments with compute extensive algorithms.

RL-X is fully open-source and available at `https://github.com/nico-bohlinger/RL-X`.

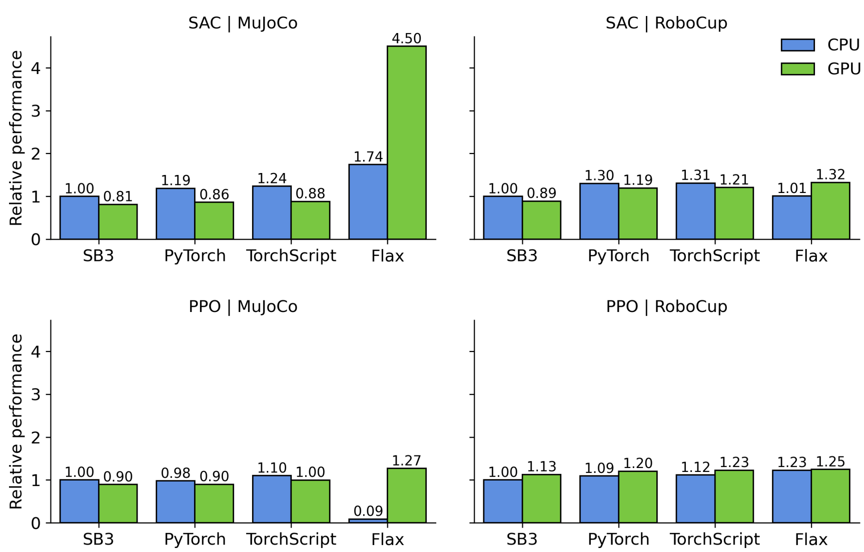

Fig. 1. Comparison of the relative computational performance of RL-X and SB3.

References

1. Raffin A, Hill A, Gleave A, Kanervisto A, Ernestus M and Dormann N: *Stable-Baselines3: Reliable Reinforcement Learning Implementations* Journal of Machine Learning Research, Volume 22, pp. 268, 2021.