

Integrating Distributed Architectures in Highly Modular Reinforcement Learning Libraries

Authors : Albert Bou, Sebastian Dittert, Gianni de Fabritiis

Abstract : Advancing reinforcement learning (RL) requires tools that are flexible enough to easily prototype new methods while avoiding impractically slow experimental turnaround times. To match the first requirement, the most popular RL libraries advocate for highly modular agent composability, which facilitates experimentation and development. To solve challenging environments within reasonable time frames, scaling RL to large sampling and computing resources has proved a successful strategy. However, this capability has been so far difficult to combine with modularity. In this work, we explore design choices to allow agent composability both at a local and distributed level of execution. We propose a versatile approach that allows the definition of RL agents at different scales through independent, reusable components. We demonstrate experimentally that our design choices allow us to reproduce classical benchmarks, explore multiple distributed architectures, and solve novel and complex environments while giving full control to the user in the agent definition and training scheme definition. We believe this work can provide useful insights to the next generation of RL libraries.

Keywords : deep reinforcement learning, Python, PyTorch, distributed training, modularity, library

Conference Title : ICLR 2023 : International Conference on Learning and Reinforcement

Conference Location : Barcelona, Spain

Conference Dates : August 10-11, 2023