Mimicking Evolution with Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

In nature, there are two processes driving the development of the brain: evolution and learning. Evolution acts slowly, across generations, and amongst other things, it defines what agents learn by changing their internal reward function. Learning acts fast, within one's lifetime, and it quickly updates agents' policies to maximise the evolved reward function. Although previous work has emulated both of these processes working in tandem, it is very computationally expensive. This work introduces Evolution via Evolutionary Reward (EvER) - a reward function that aims to maximise the number of current (and future) genetically similar agents; in a finite environment with limited resources this inadvertently drives improved survival mechanisms and reproductive success . Given that this reward function is not evolved, we avoid the computationally intense process of performing evolution on top of lifetime learning. We demonstrate the viability of this proposed reward function by testing it in two bio-inspired, open-ended environments and monitoring a number of metrics such as population size and life expectancy. We compare EvER with the state-of-the-art evolutionary algorithm: CMA-ES, and show the superiority of EvER at producing agents that maximise the number of its genes across time.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.