Parallel reinforced learning: an all-in-one AI solution

Teng Liu
Bin Tian
Yunfeng Ai
Li Li
Dongpu Cao
Fei-Yue Wang

Video Abstract

Keywords: Parallel reinforced learning: an all-in-one AI solution; deep reinforcement learning; self-driving cars; detection algorithms

Posted Date: September 20th, 2019

DOI: https://doi.org/10.21203/rs.2.15107/v1

Abstract

Researchers from China have pooled together some of the most powerful techniques in machine learning to create the ultimate control system. Successfully deployed in AI-regulated hybrid electric vehicles, the framework could grant other autonomous systems unprecedented levels of control and foresight. Machine learning is booming. And arguably the most popular technique in this branch of artificial intelligence is deep reinforcement learning. Loosely modeled after our brains' reward system, deep reinforcement learning has enabled machines to reach or even surpass human-level performance in various tasks. Those tasks range from the trivial, like playing Go or video games, to the possibly life-saving, such as detecting firearms from video. But deep reinforcement learning algorithms have their limitations. For one, they generally lack the ability to take lessons learned in one task and apply them to another. Models also tend to be inefficient, requiring enormous amounts of data to enable machines to interact with their environment. To shore up these deficiencies, the research team proposed bulking up with other machine learning techniques—hoping to exploit the unique advantages of each one when run in parallel. Transfer learning focuses on storing knowledge gained in solving one problem and applying to solve a different but related problem. It's similar to how toddlers learn to recognize shapes and animals given their limited life experience. This approach grants algorithms the ability to adapt to new settings when operating on limited data. Predictive learning enables self-guidance. This approach uses prior knowledge to model an environment, test outcomes, and forecast the best route. Finally, parallel learning brings the different machine learning styles together. Parallel learning ensures autonomous behavior by interfacing a machine system operating in the real world with a digitized version of that system operating virtually. As a proof of concept, the team applied their so-called parallel reinforcement learning approach to manage the energy use of a hybrid electric vehicle. Transfer learning made the computed management strategies adaptive to real-world driving conditions. While reinforcement learning generated the corresponding controls. In a separate system involving a hybrid tracked vehicle, predictive learning was incorporated to predict future power demand. Notably, both systems outperformed the conventional reinforcement approach in both calculation speed and control performance. If adapted to self-driving cars, the team envisions their framework capable of doing more than managing energy use. There it could help onboard AI plan trajectories and make decisions on the fly. Much more work is needed to get there. But the researcher's approach could prove a killer combo for that and many other AI applications to come.