

The AI Arena: A Framework for Distributed Multi-Agent Reinforcement Learning

Authors : Edward W. Staley, Corban G. Rivera, Ashley J. Llorens

Abstract : Advances in reinforcement learning (RL) have resulted in recent breakthroughs in the application of artificial intelligence (AI) across many different domains. An emerging landscape of development environments is making powerful RL techniques more accessible for a growing community of researchers. However, most existing frameworks do not directly address the problem of learning in complex operating environments, such as dense urban settings or defense-related scenarios, that incorporate distributed, heterogeneous teams of agents. To help enable AI research for this important class of applications, we introduce the AI Arena: a scalable framework with flexible abstractions for distributed multi-agent reinforcement learning. The AI Arena extends the OpenAI Gym interface to allow greater flexibility in learning control policies across multiple agents with heterogeneous learning strategies and localized views of the environment. To illustrate the utility of our framework, we present experimental results that demonstrate performance gains due to a distributed multi-agent learning approach over commonly-used RL techniques in several different learning environments.

Keywords : reinforcement learning, multi-agent, deep learning, artificial intelligence

Conference Title : ICMASAI 2021 : International Conference on Multi-Agent Systems and Artificial Intelligence

Conference Location : Berlin, Germany

Conference Dates : May 20-21, 2021