CALIFORNIA STATE SCIENCE FAIR
2017 PROJECT SUMMARY

Name(s)

Kevin Frans; Alex Gao

Project Number

S0808

Project Title

Coordinated Multi-Agent Control Utilizing Deep Reinforcement Learning

Abstract

Objectives/Goals

Due to the necessity for agents to learn precise coordination, control in cooperative multi-agent domains has traditionally been difficult. In previous studies, tasks such as robotic control have been represented as single agent problems, using a central neural network policy to control all joints. We examine the effects of taking a decentralized, multi-agent viewpoint, allowing each individual joint to make its own decisions. We present a novel multi-agent reinforcement learning algorithm based on the policy gradient method, by representing each agent's policy as distinct deep neural network. By training these agents simultaneously, they learn to expect the other agents' behaviors and respond accordingly, leading to cooperation. In addition, our algorithm reduces compute time by training multiple smaller networks in parallel, rather than iterating over a single large network.

Methods/Materials

We demonstrate the success and robustness of the algorithm by applying it to various tasks: 1) a swimming robot 2) a robotic arm; 3) a hopping robot; and 4) a package delivery problem. In each task, ideal performance is defined differently: the agents may be rewarded based on moving forward, or by reaching certain locations.

Our experiment procedure consists of training the algorithm with each environment, then measuring the average return for each training episode. We then repeat multiple trials for each task and compare the results the centralized method.

Results

The agents learned to cooperate in the absence of explicit communication, achieving up to two times the performance of the centralized method, while learning at a faster rate. The performance of our algorithm scaled up to more complex tasks, making it more practical for real life problems.

Conclusions/Discussion

Our work provides a reliable basis for practical cooperative learning in a multitude of environments, and paves the way for future research in the emergent field of multi-agent control. There are countless real-world applications of our algorithm, such as drone package delivery, autonomous car traffic management, and search and rescue systems.

Summary Statement

We developed a novel algorithm that autonomously teaches multiple agents how to act cooperatively.

Help Received

None