

Study of sample efficiency improvements for reinforcement learning algorithms

Tianyue Cao, Princeton International School of Mathematics and Science, USA

Abstract

Machine learning is the study of how programmed algorithms can learn useful knowledge from data automatically. As a sub-field of machine learning, reinforcement learning (RL) focuses on problems that require sequential decision making. In particular, it is about interacting with the environment and taking action according to the environment information sequentially to maximizing some rewards.

Reinforcement learning attracts many interests due to its recent successes in robotics as well as playing video games, GO, and poker. However, the fundamental challenges in reinforcement learning still limit its applications to real-world, cost and risk sensitive applications. One major challenge is relatively low sample efficiency in most systems. Sample efficiency is a term used to describe how well the samples are used to train the model. Because of low sample efficiency, it requires a huge number of samples to reach a certain level of performance. In most algorithms of reinforcement learning, methods such as experience replay are used to increase the sample efficiency. In the experience replay, a certain number of samples are saved in a buffer and new data will replace the oldest data in the set. When training, data will be randomly selected from the buffer. However, this will generate the problem of distribution mismatch, as the data chosen this way may not match the current model. In my research, methods are designed so that the samples collected from the past can reflect the current model. That will allow the model to use the data more effectively and thus increase its training efficiency.