

Master's Thesis

Exploring while Exploiting Workers for Efficient Distributed Machine Learning

We consider the problem of running a distributed machine learning algorithm with the help of workers that have different performance/speed. To fully leverage the performance of the workers, we ideally want to assign a computation load proportional to the speed of each worker. However, the performance of the workers is not known a priori. To that end, we investigate the use of multi-arm bandits theory that allows learning the performance of the workers (exploring them) while assigning useful computations (exploiting them).

Advisors

Rawad Bitar