



Improving Sample Efficiency of Online Temporal Difference Learning



Mr Yangchen Pan

Ph.D. Candidate
Department of Computing Science
University of Alberta
Canada

Date : 24 November 2020 (Tuesday)
Time : 11:00 a.m. - 12:00 noon

► Abstract

Reinforcement Learning (RL) achieved several remarkable successes in recent years, such as playing Atari games at the human-level, power station control, and finance portfolio management, etc. However, it is still far away from developing RL's full potential. One of the most important scientific hurdles is that RL algorithms suffer from low sample efficiency. That is, an RL agent typically needs to have many physical interactions with the real-world to achieve a reasonably good policy. Such interactions are typically quite expensive. I will introduce my efforts in improving the sample efficiency of online RL algorithms for both policy evaluation and control problems. Specifically, I have been making efforts in the following directions: 1) bringing in second-order optimization methods for policy evaluation algorithms in a linear function approximation setting; 2) designing a special regularization method to leverage the intrinsic structure of the concerned problem; 3) designing an efficient, scalable sparse representation learning activation function for a broad class of deep RL algorithms; 4) investigating efficient sampling distribution for model-based control problems. All of my developed methods are supported by strong empirical evidence.

► About the Speaker

Yangchen Pan is currently a Ph.D. candidate at the University of Alberta working in reinforcement learning. He is co-supervised by Dr Martha White from the University of Alberta and Dr Amir-massoud Farahmand from the University of Toronto. Yangchen's long-term research goal is to develop RL agents that interactively learn from data to solve complex real-world tasks. During his Ph.D. program, he has been working on fundamental research in reinforcement learning, aiming at improving the sample efficiency of online reinforcement learning algorithms. His research covers a broad range of topics: policy evaluation problems, model-based reinforcement learning control problems, sparse representation learning methods, extremely high dimensional continuous control problems, etc. He has published refereed papers at several well-known conferences such as ICML, ICLR, NeurIPS, AAAI, etc. He also serves as a committee member for those conferences and as a reviewer for the Journal of Machine Learning Research (JMLR).

ALL are welcome!

Enquiries : Professor George Baciu
Email : csgeorge@polyu.edu.hk
Tel : 2766 7272

We drive **innovation** through
SMART COMPUTING

