

On the Convergence of Reinforcement Learning Algorithms

SPEAKER: Keith Ross, NYU Shanghai

TIME: 5:00pm-6:00pm, Thursday, April 15, 2021

VENUE: Via Zoom (Members of NYU Shanghai Community can join from Room 613 at Pudong campus)

2021 APRIL

ABSTRACT

Much of the excitement in AI is due to recent advances in Deep Reinforcement Learning (DRL). By using DRL to overcome the curse of dimensionality in dynamic programming, researchers have recently made surprising breakthroughs in learning to play Go and the Atari video games, and learning to control high-dimensional robotic locomotion. However, when formulating and proving theorems in reinforcement learning, researchers normally work with tabular reinforcement learning due to tractability.

In the first half of this talk, I will briefly review the AlphaZero algorithm, which learns strictly from self-play and now easily beats master Go players. AlphaZero uses “Monte Carlo DRL”.

In the second half of the talk I will discuss an open problem in tabular reinforcement learning, namely, proving the almost-sure convergence of the “Monte Carlo Exploring States (MCES)” algorithm. We will provide a proof for an important special case. We will then discuss what is and is not known outside of this special case.

BIOGRAPHY

Keith Ross is the Dean of Engineering and Computer Science at NYU Shanghai and the Leonard J. Shustek Chair Professor of Computer Science at NYU Tandon. Previously he was a professor at University of Pennsylvania (13 years) and a professor at Eurecom Institute in France (5 years). He received a Ph.D. in Computer and Control Engineering from The University of Michigan. He is an ACM Fellow and an IEEE Fellow.

His current research interests are in reinforcement learning. He has also worked in Internet privacy, peer-to-peer networking, Internet measurement, stochastic modeling of computer networks, queuing theory, and Markov decision processes. At NYU Shanghai he has been teaching Machine Learning, Reinforcement Learning, and Introduction to Computer Programming.



RSVP HERE