



Distributed Stochastic Approximation and Multi-Agent Reinforcement Learning



We will discuss two mathematical frameworks for distributed reinforcement learning. These problems combine stochastic approximation, an iterative method for finding the fixed point of a function from noisy observations, and consensus, a general averaging technique for multiple agents to cooperatively solve a distributed problem.

In the first part of the talk, we will discuss the policy evaluation problem in multi-agent reinforcement learning. In this problem, a set of agents operate in a common environment under a fixed control policy, working together to discover the value (accumulative reward) associated with each environmental state. We give a finite-time analysis on the performance of the well-known "TD-lambda" algorithm that depends on the connectivity of the agents and the intrinsic properties of the Markov decision process driving the agents decisions.

In the second part of the talk, we discuss two approaches to multi-agent policy optimization. In this setting, the agents are immersed in different environments, and we are trying to learn a control policy that is effective across all of them. Our first approach is a distributed version of the venerable Q-learning algorithm, which we again analyze through the lens of distributed stochastic approximation. The second approach is based on a gradient descent algorithm on the policy itself. This problem is nonconvex, but we give conditions under which gradient descent can find the globally optimal policy.

Justin Romberg, Ph.D.

Schlumberger Professor and
Associate Director for the Center for
Machine Learning

School of Electrical & Computer Engineering
Georgia Institute of Technology

Date: Friday, October 30

Time: 1:00 – 2:00 p.m. US Central Time

Zoom Meeting ID: 998 4499 3279

Passcode: 724615

Faculty host: Simon Foucart, MATH

Biography

Dr. Romberg received the B.S.E.E, M.S. and Ph.D. degrees from Rice University in Houston, Texas. From 2003 until 2006, he was a Postdoctoral Scholar in Applied and Computational Mathematics at the California Institute of Technology. He spent the Summer of 2000 as a researcher at Xerox PARC, the Fall of 2003 as a visitor at the Laboratoire Jacques-Louis Lions in Paris, and the Fall of 2004 as a Fellow at UCLA's Institute for Pure and Applied Mathematics. In the Fall of 2006, he joined the Georgia Tech ECE faculty. In 2008 he received an ONR Young Investigator Award, in 2009 he received a PECASE award and a Packard Fellowship, and in 2010 he was named a Rice University Outstanding Young Engineering Alumnus. In 2006-2007 he was a consultant for the TV show "Numb3rs". He was an Associate Editor for the IEEE Transactions on Information Theory from 2008-2011, for the SIAM Journal on Imaging Science from 2013-2018, and the SIAM Journal on the Mathematics of Data Science from 2018-present. In 2018, he was named an IEEE Fellow.

You can also click this link to join the seminar <https://tamu.zoom.us/j/99844993279?pwd=TkJodWFVRURyMmkwaki4SWZGeVJTQT09>

