

UNIVERSITY OF CHICAGO
DEPARTMENT OF COMPUTER SCIENCE

PRESENTS:

“Foundations of Learning Systems with (Deep) Function Approximators”



Simon Du

Institute for Advanced Study

Abstract:

Function approximators, such as deep neural networks, play a crucial role in building learning systems that make predictions and decisions. In this talk, I will discuss my work on understanding, designing, and applying function approximators.

First, I will focus on understanding deep neural networks. The main result is that the over-parameterized neural network is equivalent to a new kernel, Neural Tangent Kernel. This equivalence implies two surprising phenomena: 1) the simple algorithm gradient descent provably finds the global optimum of the highly non-convex empirical risk, and 2) the learned neural network generalizes well despite being highly over-parameterized. Furthermore, this equivalence helps us design a new class of function approximators: we transform (fully-connected and graph) neural networks to (fully-connected and graph) Neural Tangent Kernels, which achieve superior performance on standard benchmarks.

In the second part of the talk, I will focus on applying function approximators to decision-making, aka reinforcement learning, problems. In sharp contrast to the (simpler) supervised prediction problems, solving reinforcement learning problems requires an exponential number of samples, even if one applies function approximators. I will then discuss what additional structures that permit statistically efficient algorithms.

Bio

Simon S. Du is a postdoc at the Institute for Advanced Study of Princeton, hosted by Sanjeev Arora. He completed his Ph.D. in Machine Learning at Carnegie Mellon University, where he was co-advised by Aarti Singh and Barnabás Póczos. Previously, he studied EECS and EMS at UC Berkeley. He has also spent time at Simons Institute and research labs of Facebook, Google, and Microsoft. His research interests are broadly in machine learning, with a focus on the foundations of deep learning and reinforcement learning.

Wednesday, February 12, 2020

10:30 am

Crerar 390

Host: Michael Maire