## ENGINEERING SCIENCES & APPLIED MATH--FEATURED SPEAKER



Neural networks as interacting particle systems: Understanding global convergence of parameter optimization dynamics

Grant Rotskoff, Ph.D. James S. McDonnell Fellow, Courant Institute, New York University

February 11, 2019 • 11 am – 12 pm M416 (ESAM Conference Room), Tech

The performance of neural networks on high-dimensional data distributions suggests that it may be possible to parameterize a representation of a target high-dimensional function with controllably small errors, potentially outperforming standard interpolation methods. We demonstrate, both theoretically and numerically, that this is indeed the case. We map the parameters of a neural network to a system of particles relaxing with an interaction potential determined by the loss function. This mapping gives rise to a deterministic partial differential equation that governs the parameter evolution under gradient descent dynamics. We also show that in the limit that the number of parameters n is large, the landscape of the mean-squared error becomes convex and the representation error in the function scales link n^{-1}. In this limit, we prove a dynamical variant of the universal approximation theorem showing that the optimal representation can be attained by stochastic gradient descent, the algorithm ubiquitously used for parameter optimization in machine learning. This conceptual framework can be leveraged to develop algorithms that accelerate optimization using non-local transport. I will conclude by showing that using neuron birth/death processes in parameter optimization guarantees global convergence and provides a substantial acceleration in practice.