



Optimization Program WISO Workshop February 8-10, 2017

Lecture: *On Gradient-Based Optimization: Accelerated, Distributed, Asynchronous and Stochastic*

Speaker: Michael I. Jordan

Abstract:

Many new theoretical challenges have arisen in the area of gradient-based optimization for large-scale statistical data analysis, driven by the needs of applications and the opportunities provided by new hardware and software platforms.

I discuss several recent results in this area, including: (1) a new framework for understanding Nesterov acceleration, obtained by taking a continuous-time, Lagrangian/Hamiltonian perspective, (2) a general theory of asynchronous optimization in multi-processor systems, (3) a computationally-efficient approach to variance reduction, and (4) a primal-dual methodology for gradient-based optimization that targets communication bottlenecks in distributed systems.