


2016 Distinguished Lecture Series UCLA Department of Mathematics

Donald Goldfarb Columbia University

Optimization for Learning and Big Data

Abstract: Many problems in both supervised and unsupervised machine learning (e.g., logistic regression, support vector machines, deep neural networks, robust principal component analysis, dictionary learning, latent variable models) and signal processing (e.g., face recognition and compressed sensing) are solved by optimization and related algorithms. In today's age of big data, the size of these problems is often formidable. E.g., in logistic regression the objective function may be expressed as the sum of 10^9 functions (one for each data point) involving 10^6 variables (features). In this series of talks, we will review current optimization approaches for addressing this challenge from the following classes of methods: first-order (and accelerated variants), stochastic gradient and second-order extensions, alternating direction methods for structured problems (including proximal and conditional gradient and multiplier methods), tensor decomposition, randomized methods for linear systems, and parallel and distributed variants.



Donald Goldfarb
Columbia University

Lecture 1

Tuesday, May 17, 2016

3:00 - 3:50 pm

MS 6627

Lecture 2

Wednesday, May 18, 2016

3:00 - 3:50 pm

MS 6627

Lecture 3

Thursday, May 19, 2016

3:00 - 3:50 pm

MS 6627