

Incremental methods for additive convex cost optimization

Keynote speaker: Asuman Ozdaglar

Motivated by machine learning problems over large data sets and distributed optimization over networks, we consider the problem of minimizing the sum of a large number of convex component functions. We study incremental gradient methods for solving such problems, which process component functions sequentially one at a time. We first consider deterministic cyclic incremental gradient methods (that process the component functions in a cycle) and provide new convergence rate results under some assumptions. We then consider a randomized incremental gradient method, called the random reshuffling (RR) algorithm, which picks a uniformly random order/permutation and processes the component functions one at a time according to this order (i.e., samples functions without replacement in each cycle). We provide the first convergence rate guarantees for this method that outperform its popular with-replacement counterpart stochastic gradient descent (SGD). We finally consider incremental aggregated gradient methods, which compute a single component function gradient at each iteration while using outdated gradients of all component functions to approximate the global cost function gradient, and provide new linear rate results.

This is joint work with Mert Gurbuzbalaban and Pablo Parrilo.