

A review of of the alternating direction method of multipliers for distributed optimization

Recently, the alternating direction method of multipliers (ADMM) regained significant attention, especially in the machine learning community, because it allows to solve convex optimization problems in a distributed setting that involve a large number of parameters. In a typical machine learning method, like the least squares method for regression problems, the parameters are just the data points. Hence, ADMM is often the method of choice for Big Data problems, where the data does not fit into the memory of a single compute node.

In my talk I will review ADMM and discuss an extension that allows to solve convex Optimization problems with a large number of constraints.

Such problems typically also arise as Big Data problems, but instead of contributing a term to the objective function each data point now contributes a constraint to the problem.