Fast Alternating Direction Augmented Lagrangian Methods for Solving Convex Optimization Problems

Donald Goldfarb
Columbia University
IEOR Department, New York, NY 10027

Abstract

We present in this talk several fast alterrnating direction augmnented Lagrangian methods for solving large-scale convex optimization problems that arise in a wide range of areas from imaging and signal processing to machine learning.

Our first class of methods apply to problems involving the sum of two functions, each of which is easy to minimize but together are not, subject to relatively simple convex constraints. Our basic methods require at most $O(\frac{1}{\epsilon})$ iterations to obtain an ϵ-optimal solution, while our accelerated (i.e., fast) versions require at most $O(\frac{1}{\sqrt{\epsilon}})$ iterations with little change in the computational effort at each iteration. These methods require only one of the two functions in the objective to be smooth and can incorporate a full line search while preserving their worst-case iteration complexities.

Our second class of methods, can be applied to objective functiolns involving any finite sum of smooth convex functions. These methods are Jacobi-like, in contrast with the first class of methods which are Gauss-Seidel-like, but have similar iteration-complexity bounds.

Our third class of methods are accelerated versions of linearized Bregman methods. The later are augmented Lagrangian methods in whcih the quadratic penalty term in the augmented Lagrangian is linearized while adding a prox term to it.

Finally, we consider alternating direction augmented Lagrangian methods that can be applied to problems whose objective function, none of whose parts are smooth. While we have not been able to derive iteration-complexity bounds for these methods, we have been able to provide strong convergence results for them. Moreover, they are extremely fast in practice.

Numerical results on very large (tens of millions of variables) compressed sensing, total-variation image denoising and deblurring, matirtx completion, and robust and stable principal component pursuit problems on real data are presented that validate our theoretical results and demonstrate the practical potential of our methods.

This research is supported in part by DOE Grant DE-FG02-08ER25856