An algorithm for Unconstrained Quadratically Penalized Convex Optimization

Steven Ellis[1]

1. NYSPI at Columbia University

Estimators are often defined as the solutions to data dependent optimization problems. So if a statistician invents a new estimator, perhaps for an unconventional application, he/she may be faced with a numerical optimization problem. In looking for software to solve that problem the statistician may find the options few, confusing, or both.

A common form of objective function (i.e., function to be optimized) that arises in statistical estimation is the sum of a convex function and a (known) quadratic complexity penalty. A standard paradigm for creating kernel-based estimators leads to exactly such an optimization problem. Suppose that the particular optimization problem of interest is of this sort and unconstrained. Unfortunately, even generic off-the-shelf software specifically written for unconstrained convex optimization is difficult to find. So the statistician may have to fall back upon a general optimizer like BFGS, which may or may not deliver good performance on the particular problem he/she is facing.

This paper describes an optimization algorithm designed for unconstrained optimization problems in which the objective function is the sum of a non-negative convex function and a known quadratic penalty. The algorithm is described and compared with BFGS on some penalized logistic regression and penalized L$\hat{(}$3/2) regression problems.