Least-Squares Covariance Matrix Adjustment

S. Boyd and L. Xiao

- psd_cone_proj.pdf

We consider the problem of finding the smallest adjustment to a given symmetric n by n matrix, as measured by the Euclidean or Frobenius norm, so that it satisfies some given linear equalities and inequalities, and in addition is positive semidefinite. This least-squares covariance adjustment problem is a convex optimization problem, and can be efficiently solved using standard methods when the number of variables (*i.e.*, entries in the matrix) is modest, say, under 1000. Since the number of variables is $n(n+1)/2$, this corresponds to a limit around n=45. Malik [2005] studies a closely related problem, and calls it the semidefinite least-squares problem. In this paper we formulate a dual problem that has no matrix inequality or variables, and a number of (scalar) variables equal to the number of equality and inequality constraints in the original least-squares covariance adjustment problem. This dual problem allows us to solve far larger least-squares covariance adjustment problems than would be possible using standard methods. Assuming a modest number of constraints, problems with n=1000 are readily solved by the dual method. The dual method coincides with the dual method proposed by Malik when there are no inequality constraints, and can be obtained as an extension of his dual method when there are inequality constraints. Using the dual problem, we show that in many cases the optimal solution is a low rank update of the original matrix. When the original matrix has structure, such as sparsity, this observation allows us to solve very large least-squares covariance adjustment problems.