
Optimization solutions for Non-Negative Least Squares problems

(Bounded–Variable Least Squares)

A search on the web will quickly reveal numerous applications for a routine which finds the best fit vector x to a system of linear equations where the components of x are constrained to be non-negative. For example statisticians may wish to fit a linear regression model:

$$y = Xc + e$$

where y is a vector of observations, X a design matrix and e a vector of noise. If c represents inherently non-negative quantities, such as prices, growth rates or pixel intensity perhaps, then a constraint might be that the elements of c be non-negative. Such a problem is termed a non-negative (linear) least squares problem.

We may similarly find situations where, additionally, the ‘solution’ should not have components exceeding a given value. Such a requirement defines a bounded-variable (linear) least squares problem.

The NAG optimization chapters contain very powerful and general routines, so such problems could always be cast into a form where the problems could be addressed by these routines. When we did this at the request of a user we found, not unsurprisingly, that these non-specialist routines were not as efficient as a routine designed specifically for such problems.

Consequently, we decided to add to the functionality of the NAG Library, at Mark 24, by including a specialist routine for bounded-variable least squares problems. This routine is E04PCx and is designed for problem sizes that are not too big. (Dense matrices and techniques are employed throughout. Depending upon user-demand, we will consider specialist sparse algorithms for larger problems at future marks.)

One feature of these types of problems is that the solution may not be unique. This occurs when the matrix (X above) is not of full rank. This may, or may not, be a problem, since the actual minimum will be the same, merely achieved by different solution values, c . If uniqueness is required then NAG provides a facility to return that solution which minimises the Euclidean norm of c i.e. the minimal length solution. This facility distinguishes the NAG routine from others on the market, but, for efficiency reasons, the option should not be employed unless a unique solution is really necessary.