Nearest Correlation Matrix

In Mark 23 of the NAG Library we have extended the functionality in the area of computing the nearest correlation matrix. In this mini-article we take look at Nearest Correlation Matrix problems, giving some background and introducing the routines that solve them.

Introduction

A correlation matrix is characterised as being a real, square symmetric matrix with ones on the diagonal and with non-negative eigenvalues. A matrix with non-negative eigenvalues is called positive semidefinite. If a matrix C is a correlation matrix then the elements of C, $C(I,J)$, represent the pair-wise correlation of entity I with entity J, that is, the strength and direction of a linear relationship between the two.

In the literature there are numerous examples illustrating the use of correlation matrices but the one we have encountered the most arises in finance where the correlation between various stocks is used to construct sensible portfolios. Unfortunately, for a variety of reasons, an input matrix which is supposed to be a correlation matrix may fail to be semidefinite. The correlations may be between stocks measured over a period of time and some data may be missing for example. Treated incorrectly the missing data problem can give rise to an indefinite matrix. Still drawing from finance, a practitioner may wish to explore the effect on a portfolio of assigning correlations between certain assets different from those computed from historical values. This can destroy the semidefiniteness of the matrix too.

In such situations the user has a matrix which approximates a correlation matrix but which fails to be so. Since subsequent analysis relies upon the matrix being a valid correlation matrix for the results to be valid, it is natural to seek a neighbouring matrix which differs least from the input matrix to act in its stead.

The Basic Nearest Correlation Matrix Problem

The NAG routine g02aaX implements a Newton algorithm to solve the basic problem we outlined in the introduction. It finds a true correlation matrix X that is closest to the approximate input matrix, G, in the Frobenius norm, that is we find the minimum of

$$||G - X||_F$$

The algorithm was described in a paper by Qi & Sun that had superior rate of convergence properties over previously suggested approaches. A research student at the University of Manchester, Rüdiger Borsdorf, with Professor Higham as supervisor, looked at this in greater detail and offered further improvements. These included a different iterative solver (MINRES was preferred to Conjugate Gradient) and a means of pre-conditioning the linear equations. It is this enhanced algorithm that has been incorporated into our library releases.

Since Mark 22 of the Library we have made some performance enhancements to the routine, by better exploiting the structure of the problem.

The Weighted Norm Problem and Forcing a Positive Definite Correlation Matrix

In the new NAG routine g02abX we have enhanced the functionality provided by g02aaX. If we have an approximate correlation matrix it is reasonable to suppose that not all of the matrix is actually approximate, but only part of it. For example, we may know that correlations are true between a subset of the entities we are measuring.

In this algorithm we apply the original work of Qi & Sun, to now use a *weighted* norm. Thus we find the minimum of

$$\left\| W^{\frac{1}{2}}(G-X)W^{\frac{1}{2}} \right\|_F$$

Here W is a diagonal matrix of weights. This means that we are seeking to minimise the elements $W(I,I)(G(I,J)-X(I,J))W(J,J)$. Thus by choosing elements in W appropriately we can favour some elements in G, forcing the corresponding elements in X to be closer to them.

The routine also allows the user to specify that the computed correlation matrix is positive definite, that is its eigenvalues are greater than zero. This is required in some applications to improve the condition of the matrix and to increase stability.

The Nearest Correlation Matrix with Factor Structure

A correlation matrix with *factor structure* is one where the off-diagonal elements agree with some matrix of rank k. That is, a correlation matrix C can be written as

$$C = \text{diag}(I - XX^T) + XX^T$$

where X here is an $n \times k$ matrix, often referred to as the factor loading matrix, and k is generally much small than n. These correlation matrices arise in factor models of asset returns, collateralized debit obligations and multivariate time series.

The routine g02aeX computes the nearest factor loading matrix, X, that gives the nearest correlation matrix (as defined above) for an approximate one, G, by finding the minimum of

$$\|G - XX^T + \text{diag}(XX^T - I)\|_F$$

We have implemented the spectral projected gradient method of Birgin, Martinez and Raydan as suggested by Borsdorf and Higham.