Derivative Free Optimization Solver for Calibration Problems

Bound constrained nonlinear least squares solver

Benjamin Marteau (Numerical Algorithms Group)

Why derivative free optimization (DFO)?

Most standard optimization methods need to evaluate the derivatives of the objective function via either

- explicitly written derivatives
- finite differencing (bumping), $\frac{\partial f}{\partial x_i} \approx \frac{f(x+he_i)-f(x)}{h}$
- algorithmic differentiation (see NAG AD tools)

Advantages of derivative free optimization

- **black box** models – algorithmic differentiation is not possible
- **noisy** models – finite differences are inaccurate or completely wrong
- **expensive** models – finite differences are impractical
- if high-accuracy is not required, DFO requires fewer function evaluations

Model-based derivative free optimization outperforms finite differences in those cases.

Derivative free optimization for data fitting

One of the most common mathematical tasks in the financial industry is fitting a model to a set of observed data, which is commonly expressed as a nonlinear least squares optimization problem

$$\min_{\alpha\in\mathbb{R}^n} \sum_{i=1}^{m} (F(x_i,\alpha) - y_i)^2 \quad \text{s.t. } l \le \alpha \le u \tag{1}$$

where α are the model parameters and $r_i(\alpha) = F(x_i,\alpha) - y_i$ are residuals expressing the quality of the fit of the model at the given point.

NAG introduces at Mark 26 a DFO solver (e04ff) able to exploit the structure of calibration problems:

1. **Initialization:** Sample the objective and build quadratic models for the residuals
2. **Repeat until convergence criteria are satisfied:**
   (a) build a quadratic of the objective from the residuals models
   (b) minimize this quadratic model in a local neighbourhood (trust region) to obtain a new estimate
   (c) replace a sample point with the new estimate
   (d) reduce or move the trust region if desirable and update the quadratic models of the residuals

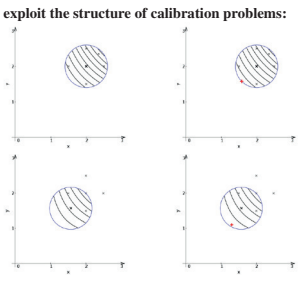

Minimizing noisy nonlinear least squares

Noise can have undesired effects for optimization solvers. Consider the unbounded problem:

$$\min_{x\in\mathbb{R}^n} \sum_{i=1}^{m} (r_i(x) + \epsilon)^2 \tag{2}$$

where ϵ is some uniform noise in the interval $[-\nu,\nu]$.

Comparison between quasi-Newton method combined with finite differences (e04fc) and the new derivative free solver (e04ff) on the Rosenbrock test function:

level of noise ν	0.0e00	1.0e-10	1.0e-08	1.0e-06	1.0e-04	1.0e-02	1.0e-01
e04fc	89	92	221	∞	∞	∞	∞
e04ff	29	29	29	29	29	31	∞

Table 1: the number of objective evaluations required to reach a point within 10^{-5} of the solution without noise

Without an efficient way to compute derivatives, the derivative free optimization framework:

- is more efficient in terms of function evaluations
- brings robustness with respect to noise

Use cases in finance

Many calibration problems in finance are very well suited to derivative free optimization:

- Sometimes F is computed with low precision algorithms (such as low precision quadrature) in the interests of speed. This introduces noise in the objective function.
- Sometimes F is computed via Monte Carlo (for example hybrid models): as paths move in and out of the money, noise is introduced into the objective function.
- Often the market prices y_i have finite precision, for example prices are quoted to one basis point. Calibration is not performed to high accuracy: it can be stopped when model prices are within one basis point of observed prices.

The table below shows a calibration of the Heston with term structure model to an accuracy of one basis point.

	average number of evaluations	average CPU time (s)
e04un	1027	43
e04ff	851	36

Table 2: comparison of DFO (e04ff) and derivative based solvers (e04un) on 1358 Heston calibrations

When exact derivatives are lacking, or a high-accuracy calibration is not required, DFO outperforms derivative based solvers

Results Matter. Trust NAG.

Nearest Correlation Matrices

The University of Manchester

1. Introduction

A correlation matrix C, has elements c_{ij} representing the pair-wise correlation of entity i with entity j, that is, the strength and direction of a linear relationship between the two. The matrix:

(a) is real, square and symmetric
(b) has unit diagonal and $|c_{ij}| \le 1$
(c) is positive semidefinite, its eigenvalues are positive or zero

Not all matrices having properties (a) and (b) have property (c):

$$\begin{pmatrix} 1.0 & 0.1 & -0.4 \\ 0.1 & 1.0 & 0.9 \\ -0.4 & 0.9 & 1.0 \end{pmatrix}$$

this matrix has eigenvalues $\lambda \approx \{-0.02, 1.07, 1.95\}$.

Approximate correlation matrices and an application in finance:

- correlations between stocks are used to construct portfolios
- some data may be missing, leading to constructed matrices being *approximate* correlation matrices that are not positive semidefinite
- a matrix that is semidefinite is required for analysis, so we seek a *true* correlation matrix that is *near* to the input matrix

This poster discusses the techniques to compute nearest correlation matrices and their implementation in the NAG Library.

2. The Basic Problem

Find 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:

$$\|G - X\|_F.$$

Alternating Projections [3,6] is one approach where we project on to the sets of semidefinite and unit diagonal matrices. Easy to implement but slow to converge.

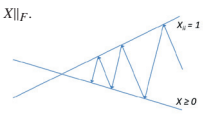

A **Newton Method**, with superior rate of convergence, was described by Qi and Sun [7]. This was improved at the University of Manchester by Borsdorf and Higham [1] using a different iterative solver and a means of preconditioning the linear equations. It has been implemented as g02aa in the NAG Library.

3. Weighting and Fixing Elements

For an approximate correlation matrix we may have greater confidence in some correlations over others or know some correlations are exact. This leads to the desire to *weight* or *fix* some elements.

In the NAG routine g02ab we have extended the functionality provided by g02aa to include **weights** [7], we find the minimum of

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

Here W is a diagonal matrix of weights. This means that we are seeking to minimize the elements $\sqrt{w_{ii}}(g_{ij} - x_{ij})\sqrt{w_{jj}}$. Thus by choosing elements in W appropriately we can favour some elements in G.

Here we show $G - X$ after a call to g02aa (left) and g02ab on an 8×8 problem. We chose weights for g02ab to attempt to fix the 3×3 leading block of true correlations. It is inexpensive, but too aggressive in its weighting.

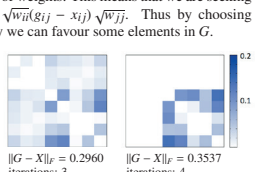

$\|G - X\|_F = 0.2960$ iterations: 3
$\|G - X\|_F = 0.3537$ iterations: 4

Rather than whole rows and columns of G being weighted, g02aj allows a preferable **element-wise weighting**, by minimizing

$$\|H \circ (G - X)\|_F.$$

Here, by choosing appropriate values in H we can emphasise *individual elements* in G and leave the others unweighted [5].

Fixing a block of correlations was introduced in Mark 25 of the NAG Library in g02an. Using the shrinking method of Higham, Strabić and Šego [4], the routine finds a true correlation matrix of the form

$$\alpha\begin{pmatrix} G_{11} & 0 \\ 0 & I \end{pmatrix} + (1-\alpha)G, \qquad G = \begin{pmatrix} G_{11} & G_{12} \\ G_{12}^T & G_{22} \end{pmatrix}.$$

We find the smallest $\alpha \in [0,1]$ that gives us a positive semidefinite result, preserving G_{11} which needs to be positive definite.

On the left is $G - X$ after a call to g02aj on our problem. We have only weighted the 3×3 leading block, but it is **very** expensive. On the right g02an fixes the block cheaply but the result is further away from our input.

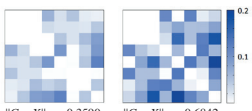

$\|G - X\|_F = 0.3590$ iterations: 414
$\|G - X\|_F = 0.6842$ iterations: 27

In Mark 26 of the NAG Library we generalized the shrinking algorithm to allow the **fixing of arbitrary elements**. The user defines a positive definite *target* matrix, T, and g02ap finds a solution of the form

$$\alpha T + (1-\alpha)G, \qquad T = H \circ G.$$

Setting elements of H to 1 will fix corresponding elements in G.

4. 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. It can be written as

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

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

The routine g02ae computes the nearest factor loading matrix, X, for an approximate matrix, G, by minimizing

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

We have implemented the spectral projected gradient method of Birgin, Martinez and Raydan as suggested by Borsdorf *et al.* [2].

5. Summary of Available Routines

NAG offer these routines to compute nearest correlation matrices:

- g02aa basic problem using the method of Qi and Sun
- g02ab incorporates weights and bounds on eigenvalues
- g02ae matrices with k-factor structure
- g02aj provides element-wise weighting of input
- g02an shrinking method with fixed submatrix
- g02ap shrinking method with user supplied target

References

To find out more, see the following papers, and references therein:
[1] R. Borsdorf and N. J. Higham. A preconditioned (Newton) algorithm for the nearest correlation matrix. *IMA J. of Numer. Anal.*, 30(1):94-107, 2010.
[2] R. Borsdorf, N. J. Higham and M. Raydan. Computing a nearest correlation matrix with factor structure. *SIAM J. Matrix Anal. Appl.*, 31(5):2603-2622, 2010.
[3] N. J. Higham. Computing the nearest correlation matrix - A problem from finance. *IMA J. Numer. Anal.*, 22(3):329-343, 2002.
[4] N. J. Higham, N. Strabić and V. Šego. Restoring definiteness via shrinking, with an application to correlation matrices with a fixed block. *SIAM Review*., 58(2) 245-263, 2016.
[5] K. Jiang, D. Sun and K-C. Toh. An inexact accelerated proximal gradient method for large scale linearly constrained convex SDP. *SIAM J. Optim.*, 22(3):1042-1064, 2012.
[6] C. Lucas. Computing Nearest Covariance and Correlation Matrices. M.Sc. Thesis, University of Manchester, 2001.
[7] H. Qi and D. Sun D. A quadratically convergent Newton method for computing the nearest correlation matrix. *SIAM J. Matrix Anal. Appl.*, 29(2):360-385, 2006.

Craig Lucas craig.lucas@nag.co.uk