

*A Preconditioned Newton Algorithm for the
Nearest Correlation Matrix*

Borsdorf, Rudiger and Higham, Nicholas J.

2008

MIMS EPrint: **2008.50**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

Reports available from: <http://eprints.maths.manchester.ac.uk/>

And by contacting: The MIMS Secretary
School of Mathematics
The University of Manchester
Manchester, M13 9PL, UK

ISSN 1749-9097

A PRECONDITIONED NEWTON ALGORITHM FOR THE NEAREST CORRELATION MATRIX*

RÜDIGER BORSODORF[†] AND NICHOLAS J. HIGHAM[‡]

Abstract. Various methods have been developed for computing the correlation matrix nearest in the Frobenius norm to a given matrix. We focus on a quadratically convergent Newton algorithm recently derived by Qi and Sun. Various improvements to the efficiency and reliability of the algorithm are introduced. Several of these relate to the linear algebra: the Newton equations are solved by minres instead of the conjugate gradient method, as it more quickly satisfies the inexact Newton condition; we apply a Jacobi preconditioner, which can be computed efficiently even though the coefficient matrix is not explicitly available; an efficient choice of eigensolver is identified; and a final scaling step is introduced to ensure that the returned matrix has unit diagonal. Potential difficulties caused by rounding errors in the Armijo line search are avoided by altering the step selection strategy. These and other improvements lead to a significant speedup over the original algorithm and allow the solution of problems of dimension a few thousand in a few tens of minutes.

Key words. correlation matrix, positive semidefinite matrix, Newton's method, preconditioning, rounding error, Armijo line search conditions, alternating projections method

AMS subject classifications. 65F10, 65F15, 90C25

1. Introduction. Correlation matrices are real, symmetric positive semidefinite matrices with ones on the diagonal. They arise in situations where correlations between pairs of random variables are computed, but also when pairwise similarity measures between objects are formed and suitably scaled, for example in machine learning [18]. It is common in practice to be faced with an *approximate* correlation matrix: a matrix that is supposed to be a correlation matrix but for a variety of possible reasons is not. In finance, for example, the correlations may be between stocks measured over a period of time and missing data (perhaps due to a company not trading for the whole period) may compromise the correlations and lead to a non-positive semidefinite matrix. Again in finance, a practitioner may wish to explore the effect on a portfolio of assigning correlations between certain assets differently from the historical values, but this again can destroy the semidefiniteness of the matrix. The use of approximate correlation matrices in these applications can render the methodology invalid and lead to negative variances and volatilities being computed [5], [16], [21].

The prevalence of approximate correlation matrices has led to much interest in the problem of computing the nearest correlation matrix to a given matrix $A \in \mathbb{R}^{n \times n}$, that is, solving the problem

$$(1.1) \quad \min\{\|A - X\|_F : X = X^T, X \geq 0, \text{Diag}(X) = e\},$$

where for symmetric matrices X and Y , $X \geq Y$ denotes that $X - Y$ is positive semidefinite, $\text{Diag}(X)$ is the vector of diagonal elements of X , e is the vector of 1s, and the Frobenius norm $\|X\|_F = \text{trace}(X^T X)^{1/2}$. The objective function in (1.1) is convex and the constraints are closed convex sets, so there is a unique minimizer.

*Version of April 21, 2008. This work was supported by a Royal Society-Wolfson Research Merit Award to the second author.

[†]Department of Mathematics, Chemnitz University of Technology, D-09107 Chemnitz, Germany (ruediger.borsdorf@s2003.tu-chemnitz.de). This author was supported by an EPSRC Studentship held in the School of Mathematics, The University of Manchester.

[‡]School of Mathematics, The University of Manchester, Manchester, M13 9PL, UK (higham@ma.man.ac.uk, <http://www.ma.man.ac.uk/~higham>).

The nearness problem (1.1) has been extensively studied over the last twenty years. Much of the literature is concerned with ad hoc methods that are not guaranteed to solve the problem. An early example is a method of Knol and ten Berge [10] that writes $X = Y^T Y$ and iteratively minimizes the objective function over each unit 2-norm column of Y . More recently, Lurie and Goldberg [11] use the Gauss–Newton method to minimize $\|A - R^T R\|_F^2$, where R is upper triangular with columns of unit 2-norm.

Higham [9] uses convex analysis to give a characterization of the solution and describes an alternating projections algorithm that converges linearly to the solution. This algorithm has several attractive features. First, it is very simple to implement, requiring just matrix additions and computation of eigendecompositions. Second, it can take advantage of the property, proved in [9], that if $a_{ii} \geq 1$ for all i and A has many negative eigenvalues (as is likely in finance applications) then the solution has at least as many zero eigenvalues (so is of low rank). Third, it is readily adapted to solve the problem with additional constraints that require X to belong to a convex set, such as a constraint that holds any set of elements of X fixed [1, Chap. 7]. The main drawback of the alternating projections algorithm is its possibly slow convergence.

Malick [12] studies a problem more general than (1.1) in which the positive semidefinite matrices are replaced by a convex set and the constraints on the diagonal of X are replaced by general linear constraints. He dualizes the linear constraints and applies a quasi-Newton method to the dual problem. Boyd and Xiao [2] explore similar ideas. When applied to the problem (1.1) the methods in both these papers are linearly convergent, because the dual objective function is not twice continuously differentiable. A breakthrough was subsequently made by Qi and Sun [15], who derive a quadratically convergent Newton method for (1.1), again by working with the dual problem. The proof of quadratic convergence relies heavily on the theory of semismooth optimization. Interior point methods can also be applied to classes of problems containing (1.1). Toh [20] develops such a method that requires the solution of dense linear systems of dimension about $n^2/2$, constructs preconditioners for the systems, and applies the method to (1.1).

Qi and Sun [15] build from their theory a globally convergent Newton algorithm for finding the nearest correlation matrix and illustrate its performance on a small set of artificial test problems. The purpose of our work is to improve the efficiency and reliability of the algorithm through a careful analysis of its component steps. The main improvements we make are as follows.

- The Newton equations are solved by minres instead of the conjugate gradient (CG) method, since minres minimizes the residual that appears in the inexact Newton condition.
- We show how to efficiently apply a Jacobi preconditioner to the Newton equations—a nontrivial task since the coefficient matrix is not explicitly available.
- The line search is modified so as to perform reliably in finite precision arithmetic when the convergence tolerance is close to the machine precision.
- We show experimentally that the choice of eigensolver can have a significant effect on the computational cost and identify a suitable choice.
- We introduce a final scaling step, justified by a distance bound, that ensures that the returned matrix has unit diagonal.

The outline of the paper is as follows. In Section 2 we present background on the dual problem and the Newton method and we state Qi and Sun’s algorithm. In

Section 3 we develop our refinements to the algorithm. The improved algorithm is described in Section 4 and some numerical experiments are reported in Section 5. Concluding remarks are in Section 6.

2. Qi and Sun's Newton algorithm. In this section we summarize the key results from the analysis of Qi and Sun [15] that will be needed later and we state Qi and Sun's algorithm.

The problem obtained by dualizing the linear constraints in the nearest correlation matrix problem (1.1) is the unconstrained convex optimization problem [12], [15]

$$(2.1) \quad \min_{y \in \mathbb{R}^n} f(y) := \frac{1}{2} \|(A + \text{diag}(y))_+\|_F^2 - e^T y.$$

Here, $\text{diag}(y)$ for $y \in \mathbb{R}^n$ denotes the diagonal matrix whose diagonal elements are those of the vector y , while $\text{diag}(A)$ for $A \in \mathbb{R}^{n \times n}$ denotes $\text{diag}(a_{11}, a_{22}, \dots, a_{nn})$. (Recall that Diag , introduced in Section 1, maps matrices onto vectors.) The operator $(\cdot)_+$ projects onto the positive semidefinite matrices: for symmetric $C \in \mathbb{R}^{n \times n}$ with spectral decomposition $C = Q\Lambda Q^T$ ($Q^T Q = I$, $\Lambda = \text{diag}(\lambda_i)$), $C_+ = Q \text{diag}(\max(\lambda_i, 0)) Q^T$ is the nearest positive semidefinite matrix to C in the Frobenius norm [7]. The following lemma collects some key properties of the dual problem obtained by Malick [12].

LEMMA 2.1. *The dual problem (2.1) has the following properties.*

- (a) *f is convex and continuously differentiable and has a unique minimizer.*
- (b) *The gradient ∇f is given by*

$$(2.2) \quad \nabla f(y) = \text{Diag}((A + \text{diag}(y))_+) - e$$

and is Lipschitz continuous with Lipschitz constant 1.

- (c) *The solutions y_* of the dual problem (2.1) and X_* of the primal problem (1.1) are related by*

$$(2.3) \quad X_* = (A + \text{diag}(y_*))_+.$$

Note that the lemma shows that the original constrained problem with $(n^2 - n)/2$ variables is equivalent to an unconstrained problem with just n variables.

To find y_* we need to solve $g(y_*) = 0$, where $g(y) = \nabla f(y)$. Noting that g is not differentiable, we denote by ∂g the generalized Jacobian, which is defined since ∇f is Lipschitz continuous. Qi and Sun apply the generalized Newton iteration

$$(2.4) \quad y_{k+1} = y_k - V_k^{-1} g(y_k), \quad V_k \in \partial g(y_k), \quad k = 0: \infty.$$

For a general g this iteration need not converge. However, by exploiting the strong semismoothness of the operator $(\cdot)_+$, Qi and Sun are able to prove quadratic convergence of the iteration for $g = \nabla f$.

THEOREM 2.2 (Qi and Sun). *Let y_* denote the minimizer of (2.1). All $V \in \partial g(y_*)$ are positive definite and the generalized Newton method (2.4) converges quadratically to y_* for any choice of V_k if y_0 is sufficiently close to y_* .*

In order to implement the method we need to be able to compute a generalized Jacobian $V \in \partial g(y)$. Qi and Sun show that such a matrix is given implicitly by

$$(2.5) \quad V_y h = \text{Diag}(P_y(W_y \circ (P_y^T H P_y)) P_y^T).$$

Here, \circ denotes the Hadamard product ($X \circ Y = (x_{ij}y_{ij})$); $h \in \mathbb{R}^n$ and $H = \text{diag}(h)$; P_y is an orthogonal matrix calculated from the spectral decomposition of $A + \text{diag}(y)$:

$$(2.6) \quad A + \text{diag}(y) = P_y \text{diag}(\lambda(y)) P_y^T,$$

with $\lambda(y)$ the vector of eigenvalues; and W_y depends on the eigenvalues $\lambda(y)$ in the way we now describe. Let $\lambda(y)$ be in descending order and define the sets $\alpha = \{i : \lambda_i(y) > 0\}$, $\beta = \{i : \lambda_i(y) = 0\}$, and $\gamma = \{i : \lambda_i(y) < 0\}$. Then the matrix W_y is defined by

$$(2.7) \quad W_y = \begin{bmatrix} E_{\alpha\alpha} & E_{\alpha\beta} & \mathcal{T} \\ E_{\beta\alpha} & 0 & 0 \\ \mathcal{T} & 0 & 0 \end{bmatrix}, \quad \mathcal{T} = \left(\frac{\lambda_i(y)}{\lambda_i(y) - \lambda_j(y)} \right)_{i \in \alpha, j \in \gamma},$$

where $E_{\alpha\beta}$ denotes the matrix of ones of dimension $|\alpha| \times |\beta|$. It is easy to show that $V_y \geq 0$. We have

$$\begin{aligned} h^T V_y h &= \text{trace}(H P_y (W_y \circ (P_y^T H P_y)) P_y^T) \\ &= \text{trace}(\tilde{H} (W_y \circ \tilde{H})) \quad \text{where } \tilde{H} = P_y^T H P_y, \\ &= \|\tilde{W} \circ \tilde{H}\|_F^2 \quad \text{where } \tilde{W} \circ \tilde{W} = W_y, \\ &\geq 0, \end{aligned}$$

using the symmetry of \tilde{H} .

The matrix V_y can be explicitly computed by setting $h = e_i$ in (2.5), $i = 1:n$, where e_i is the i th unit vector. But since evaluating (2.5) for a single h costs $O(n^3)$ operations, obtaining V_y costs a prohibitively expensive $O(n^4)$ operations. We are therefore restricted to solving the Newton equation by methods that require matrix-vector products only.

The following algorithm implements the method above as an inexact Newton method (the linear system (2.4) is solved only approximately) and it uses a line search strategy and globalization techniques. The algorithm is globally convergent and is essentially the same as [15, Alg. 5.1].

ALGORITHM 2.3. *Given a symmetric matrix $A \in \mathbb{R}^{n \times n}$ and a convergence tolerance tol this algorithm computes the nearest correlation matrix X to A in the Frobenius norm. On termination $\|\nabla f(y_k)\|_2 \leq \text{tol}$ (see (2.2)). The algorithm is quadratically convergent.*

- Step 1: Initialization: $y_0 \in \mathbb{R}^n$, $\eta \in (0, 1)$, $\rho, \sigma \in (0, 1/2]$, and $k = 0$.
- Step 2: Calculate $\nabla f(y_k)$. If $\|\nabla f(y_k)\|_2 \leq \text{tol}$, set $X = (A + \text{diag}(y_k))_+$ and quit.
- Step 3: Compute a spectral decomposition (2.6) of $A + \text{diag}(y_k)$ and form the matrix W_{y_k} from (2.7).
- Step 4: Determine the new direction d_k by applying an iterative method (using (2.5) to compute $V_k d_k$) to

$$(2.8) \quad V_k d_k = -\nabla f(y_k),$$

terminating when both the conditions

$$(2.9) \quad \|\nabla f(y_k) + V_k d_k\|_2 \leq \eta_k \|\nabla f(y_k)\|_2,$$

$$(2.10) \quad -\frac{\nabla f(y_k)^T}{\|d_k\|_2} \cdot \frac{d_k}{\|d_k\|_2} \geq \eta_k,$$

are satisfied, where $\eta_k = \min(\eta, \|\nabla f(y_k)\|_2)$. If either one of these conditions cannot be satisfied, let

$$(2.11) \quad d_k = -B_k^{-1} \nabla f(y_k)$$

where B_k is any symmetric positive definite matrix with $\{\|B_k\|_2\}$ and $\{\|B_k^{-1}\|_2\}$ uniformly bounded.

Step 5: Choose an appropriate step length α_k by applying Armijo backtracking: find the smallest nonnegative integer m_k such that

$$(2.12) \quad f(y_k + \rho^{m_k} d_k) \leq f(y_k) + \sigma \rho^{m_k} \nabla f(y_k)^T d_k$$

is satisfied.

Step 6: Set $\alpha_k = \rho^{m_k}$, $y_{k+1} = y_k + \alpha_k d_k$ and $k \leftarrow k + 1$. Goto Step 2.

In the next section we develop several refinements that improve the efficiency and robustness of the basic algorithm.

3. Refinements.

3.1. Linear equation solver. Qi and Sun [15] take the CG method as the solver for the Newton system (2.8). The stopping criterion (2.9) is based on the norm of the residual $r_k = \nabla f(y_k) + V_k d_k$, but the CG method minimizes the *error* $y_* - y_k$ in the V_k norm ($\|x\|_{V_k} = (x^T V_k x)^{1/2}$) rather than the residual, and it can produce very non-monotonic residuals. Moreover, the possible singularity of the coefficient matrix V_k can cause problems for the CG method.

Instead of CG, we use the minres method of Paige and Saunders [14]. This method minimizes the residual norm $\|r_k\|_2$ at every iteration and so produces a monotonically decreasing sequence of residuals. Moreover, unlike CG, minres is defined for indefinite systems and thus it should be more stable than CG in finite precision arithmetic for nearly singular or numerically indefinite matrices. Minres requires only one matrix-vector product per iteration, but it requires a few more vector operations than CG.

3.2. Preconditioning. As we noted in Section 2, the coefficient matrix V_k is not explicitly available. We know that V_k is positive semidefinite and is positive definite for k sufficiently large, but nothing else is known about the spectrum of V_k . Preconditioning the system (2.8) is therefore a challenge. However, it is possible to compute the diagonal elements of V_k in $O(n^3)$ operations and thereby to apply the Jacobi preconditioner. Recall that the Jacobi preconditioner for a positive definite matrix A is $D = \text{diag}(A)$ and the preconditioned matrix is $D^{-1/2} A D^{-1/2}$, which has 2-norm condition number within a factor n of the minimum over all diagonal congruences, by a result of van der Sluis [22], [8, Cor. 7.6]. The Jacobi preconditioner is a reasonable choice in view of the existence of residual bounds for minres that depend on $\kappa_2(A)$ [4, Chap. 6], [6, Chap. 3].

To see how to compute $\text{diag}(V_k)$, let $h = e_i$, $H = e_i e_i^T$, and $P_k^T = [p_1, p_2, \dots, p_n]$. Then, from (2.5), the (i, i) element of V_k is given by

$$\begin{aligned} v_{ii} &= e_i^T P_k (W_k \circ P_k^T H P_k) P_k^T e_i \\ &= p_i^T (W_k \circ p_i p_i^T) p_i \\ &= p_i^T \text{diag}(p_i) W_k \text{diag}(p_i) p_i \\ &= q_i^T W_k q_i, \end{aligned}$$

where $q_i = p_i \circ p_i$. Thus the diagonal elements v_{ii} can be computed as follows:

$$\begin{aligned} Q_k &= [q_1, q_2, \dots, q_n] = P_k \circ P_k && n^2 \text{ flops,} \\ M_k &= [m_1, m_2, \dots, m_n] = W_k Q_k && \leq 2n^3 \text{ flops,} \\ v_{ii} &= q_i^T m_i, \quad i = 1:n && 2n^2 \text{ flops.} \end{aligned}$$

The dominant cost is therefore the matrix-matrix multiplication giving M_k . In forming M_k the zero and ee^T blocks of W_k (see (2.7)) can be exploited to reduce the cost.

To allow for a possibly singular V_k and the effects of rounding errors we set all diagonal entries less than a predefined positive tolerance to that tolerance.

3.3. Armijo backtracking. We are aiming for an algorithm that is capable of computing the nearest correlation matrix to full machine accuracy, so we wish to allow the convergence tolerance tol in Algorithm 2.3 to be of the order of the unit roundoff, u . However, the Armijo backtracking can break down for small tolerances. To see why, consider a twice continuously differentiable function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ and the expansion

$$\phi(x+p) = \phi(x) + \nabla\phi(x)^T p + \frac{1}{2} p^T \nabla^2\phi(x+tp)^T p, \quad t \in (0, 1).$$

If $|\phi(x)| = 1$, $\|p\|_2 < u^{1/2}$, $\|\nabla\phi(x)\|_2 < u^{1/2}/2$, and $\|\nabla^2\phi(x+tp)\|_2 < 1$ then $|\phi(x+p) - \phi(x)| < u|\phi(x)|$ and so $fl(\phi(x+p)) = fl(\phi(x))$. In this situation, x may still be some distance from a minimizer of ϕ —albeit perhaps only one or two steps away for a quadratically converging method—yet the Armijo condition cannot be verified because the function values it needs to compare are indistinguishable in floating point arithmetic. In numerical experiments we have found that this problem with the Armijo condition can cause Algorithm 2.3 to fail to converge in finite precision arithmetic.

To avoid this problem, when $fl(f(y_k + \rho^{m_k} d_k))$ and $fl(f(y_k))$ are close enough that rounding errors are dominating we take the inexact Newton direction with step length 1, provided that the resulting y_{k+1} satisfies

$$(3.1) \quad \frac{\|\nabla f(y_{k+1})\|_2}{\|\nabla f(y_k)\|_2} < 1 - \mu \quad \text{for some } \mu > 0,$$

where the latter condition ensures that useful progress is made towards the minimizer. Or, if (3.1) is not satisfied we take the steepest descent direction with step length 1.

The next result provides support for the test (3.1) by showing that it is satisfied for sufficiently large k .

LEMMA 3.1. *For sufficiently large k in Algorithm 2.3, $y_{k+1} = y_k + d_k$ with d_k the inexact Newton direction and*

$$\frac{\|\nabla f(y_{k+1})\|_2}{\|\nabla f(y_k)\|_2} = O(\|d_k\|_2).$$

Proof. From the proof of [15, Thm. 5.3] we know that for all sufficiently large k the inexact Newton step is taken with step $\alpha_k = 1$, d_k satisfies (2.9) and (2.10), that

$$(3.2) \quad \|y_{k+1} - y_*\|_2 = O(\|y_k - y_*\|_2^2),$$

$$(3.3) \quad \|y_k - y_*\|_2 \leq \|d_k\|_2 + O(\|d_k\|_2^2),$$

and also that there exists a $\rho > 0$ so that

$$(3.4) \quad \|\nabla f(y_k)\|_2 \|d_k\|_2 \geq -\nabla f(y_k)^T d_k \geq \rho \|d_k\|_2^2.$$

It follows from (3.3) and (3.4) that for sufficiently large k

$$(3.5) \quad \frac{\|y_k - y_*\|_2^2}{\|\nabla f(y_k)\|_2} \leq \frac{\|y_k - y_*\|_2^2}{\rho \|d_k\|_2} \leq \frac{1}{\rho} \|d_k\|_2 + O(\|d_k\|_2^2) = O(\|d_k\|_2).$$

From (3.2), using the Lipschitz property in Lemma 2.1 (b) of $\nabla f(y)$ and (3.5), we deduce that

$$\begin{aligned} \frac{\|\nabla f(y_{k+1})\|_2}{\|\nabla f(y_k)\|_2} &= \frac{\|\nabla f(y_{k+1}) - \nabla f(y_*)\|_2}{\|\nabla f(y_k)\|_2} \leq \frac{\|y_{k+1} - y_*\|_2}{\|\nabla f(y_k)\|_2} \\ &= O\left(\frac{\|y_k - y_*\|_2^2}{\|\nabla f(y_k)\|_2}\right) = O(\|d_k\|_2), \end{aligned}$$

which completes the proof. \square

3.4. Accuracy of the solution. A subtle problem with Algorithm 2.3 is that it does not yield a matrix with unit diagonal, because the constraints $\text{Diag}(X) = e$ are not explicitly enforced. Indeed if $\nabla f(y) \neq 0$ on termination then it is clear from (2.2) and (2.3) that the returned matrix does not have unit diagonal. If we simply set the diagonal elements to 1 then we may destroy the definiteness. We could then restore definiteness by projecting onto the nearest positive semidefinite matrix, which changes the diagonal. Iterating this procedure essentially gives the alternating projections algorithm [9].

We will adopt a simpler and less expensive approach. We replace the final iterate X by

$$\tilde{X} = D^{-1/2} X D^{-1/2}, \quad D = \text{diag}(X),$$

which has unit diagonal. Since this transformation is a congruence it preserves the definiteness of X . However, it can increase the distance from A . The next lemma provides a bound on the increase.

LEMMA 3.2. *If $X \geq 0$ is the output of Algorithm 2.3 and $D = \text{diag}(X) > 0$ then*

$$(3.6) \quad \|A - D^{-1/2} X D^{-1/2}\|_F \leq \|A - X\|_F + \frac{\text{tol}}{1 - \text{tol}} \|X\|_F.$$

Proof. We have

$$(3.7) \quad \|A - D^{-1/2} X D^{-1/2}\|_F \leq \|A - X\|_F + \|X - D^{-1/2} X D^{-1/2}\|_F.$$

Our aim is to bound the second term of (3.7) by using $\|\nabla f(y_k)\|_2 \leq \text{tol}$, where y_k is the final iterate of Algorithm 2.3. From (2.2) and (2.3) it follows that

$$(3.8) \quad D = \text{diag}(\nabla f(y_k)) + I.$$

With $G = X - D^{-1/2} X D^{-1/2}$, we have

$$(3.9) \quad \begin{aligned} g_{ij}^2 &= \left(x_{ij} - (\nabla f(y_k)_i + 1)^{-\frac{1}{2}} x_{ij} (\nabla f(y_k)_j + 1)^{-\frac{1}{2}}\right)^2 \\ &= x_{ij}^2 \left(1 - (\nabla f(y_k)_i + 1)^{-\frac{1}{2}} (\nabla f(y_k)_j + 1)^{-\frac{1}{2}}\right)^2. \end{aligned}$$

Using $|\nabla f(y_k)_i| \leq \text{tol}$ for all i yields

$$(3.10) \quad \frac{1}{1 + \text{tol}} \leq (\nabla f(y_k)_i + 1)^{-\frac{1}{2}} (\nabla f(y_k)_j + 1)^{-\frac{1}{2}} \leq \frac{1}{1 - \text{tol}}.$$

Hence, in order to find an upper bound for g_{ij}^2 it is enough to maximize the function $f(s) = (1 - 1/(1 + s))^2 = s^2/(1 + s)^2$ over $s \in [-\text{tol}, \text{tol}]$. The maximum is attained at $s = -\text{tol}$ and so we obtain from (3.9) $g_{ij}^2 \leq x_{ij}^2 \text{tol}^2 / (1 - \text{tol})^2$, giving $\|G\|_F \leq \|X\|_F \text{tol} / (1 - \text{tol})$. The required bound then follows. \square

The bound (3.6) is very satisfactory: it says that the increase in the distance $\|A - X\|_F$ induced by the normalization of the diagonal is at most about $\text{tol}\|X\|_F \lesssim n\text{tol}$, and we expect $n\text{tol} \ll \|A - X\|_F$ in applications.

3.5. Choice of eigensolver. Algorithm 2.3 requires a full eigenvalue decomposition of the symmetric matrix $A + \text{diag}(y)$ for every evaluation of $f(y)$ and $\nabla f(y)$, thus at least one eigenvalue decomposition per iteration. This is a major part of the total cost of the method, so it is essential to minimize its cost.

There are three main algorithmic options, for which the NAG Library and LAPACK codes are `f08fa/dsyev`, `f08fc/dsyevd`, and `f08fd/dsyevr`. All three algorithms reduce the symmetric matrix to a tridiagonal matrix but then proceed differently: `f08fa` uses the QR algorithm, `f08fc` uses the divide and conquer algorithm, and `f08fd` uses the dqds algorithm and multiple relatively robust representations (MRRR). In our numerical experiments we compare these algorithms.

4. The modified algorithm. The following modification of Algorithm 2.3 incorporates the improvements described in the previous section.

ALGORITHM 4.1. *Given a matrix $A \in \mathbb{R}^{n \times n}$ and a convergence tolerance tol this algorithm computes the nearest correlation matrix X to A in the Frobenius norm. On termination $\|\nabla f(y_k)\|_2 \leq \text{tol}$ (see (2.2)). The algorithm is quadratically convergent.*

- Step 1: Initialization: $\eta = 0.5$, $\varphi = 10^{-6}$, $\mu \in (0, 1)$, $\rho, \sigma \in (0, 1/2]$, and $k = 0$.
- Step 2: Set $A \leftarrow (A + A^T)/2$ if A is nonsymmetric. Set $a_{ii} = 1$, $i = 1:n$, and $y_0 = 0$.
- Step 3: Calculate $\nabla f(y_k)$. If $\|\nabla f(y_k)\|_2 \leq \text{tol}$, set $X = D^{-1/2} \tilde{X} D^{-1/2}$ where $\tilde{X} = (A + \text{diag}(y_k))_+$ and $D = \text{diag}(\tilde{X})$, and quit.
- Step 4: Compute a spectral decomposition of $A + \text{diag}(y_k)$ and form the matrix W_{y_k} from (2.7) and the Jacobi preconditioner D_k (see Section 3.2).
- Step 5: Determine the new direction d_k by applying minres to the preconditioned linear system (using (2.5) to compute $V_k d_k$)

$$(4.1) \quad D_k^{-1/2} V_k D_k^{-1/2} \cdot D_k^{1/2} d_k = -D_k^{-1/2} \nabla f(y_k),$$

terminating when both the conditions

$$(4.2) \quad \|\nabla f(y_k) + V_k d_k\|_2 \leq \min(\eta, \|\nabla f(y_k)\|_2) \|\nabla f(y_k)\|_2,$$

$$(4.3) \quad -\frac{\nabla f(y_k)^T}{\|d_k\|_2} \cdot \frac{d_k}{\|d_k\|_2} \geq \min(\varphi, \|\nabla f(y_k)\|_2),$$

are satisfied. If either of these conditions cannot be satisfied, let $d_k = -\nabla f(y_k)$.

- Step 6: (Choice of step using Armijo backtracking.)
for $m = 0: \infty$

If $f(y_k + \rho^m d_k) \leq f(y_k) + \sigma \rho^m \nabla f(y_k)^T d_k$ set $\alpha_k = \rho^m$ and goto Step 7.

If $f(y_k + \rho^m d_k)$ and $f(y_k)$ are “nearly equal” then if

$$(4.4) \quad \frac{\|\nabla f(y_k + \alpha_k d_k)\|_2}{\|\nabla f(y_k)\|_2} \leq 1 - \mu,$$

set $\alpha_k = 1$ and goto Step 7, else set $d_k = -\nabla f(y_k)$, $\alpha_k = 1$, and goto Step 7.

end

Step 7: Set $y_{k+1} = y_k + \alpha_k d_k$ and $k \leftarrow k + 1$. Goto Step 3.

A few comments are in order.

(a) Since all matrices agreeing with A on the off-diagonal have the same nearest correlation matrix we set the diagonal of A to unity at the start. With $y_0 = 0$ this gives immediate convergence if the resulting matrix is positive semidefinite.

(b) Step 2 projects onto the symmetric matrices [7] and allows the algorithm to work for nonsymmetric inputs A [1, Thm. 4.91].

(c) Our use of $\varphi \ll \eta$ in (4.3) and (4.2) encourages the use of inexact Newton directions over the steepest descent direction, which our numerical experiments have shown leads to faster run times.

(d) In Step 6 a suitable test for two floating point numbers a and b being “nearly equal” is $|a - b| < \gamma u(1 + |a| + |b|)$, where γ is a constant that we set to 100.

(e) Some other straightforward tests that terminate when rounding errors start to dominate are omitted to avoid clutter, but are included in our implementation tested in the next section.

5. Numerical experiments. We now present some numerical experiments that illustrate the behaviour of Algorithm 4.1 and compare it with Algorithm 2.3 and the alternating projections method. Details of more extensive tests can be found in [1]. The tests were carried out in MATLAB R2006b on a 2.4 GHz AMD Athlon under linux (Tables 5.1 and 5.2) and MATLAB R2007b on a 2.2 GHz AMD Athlon under Windows XP (Tables 5.3 and 5.4); the unit roundoff is $u = 2^{-53} \approx 1.1 \cdot 10^{-16}$. We invoked certain NAG Fortran Library (Mark 21) codes via the NAG Toolbox for MATLAB (Beta 1 under linux and Beta 2 under Windows) [13].

The codes tested are as follows.

- **nearcor**: a MATLAB implementation of Algorithm 2.3 written by the authors of [15] and used in the testing in that paper.
- **nearcor_new**: Our MATLAB implementation of Algorithm 4.1. We take $\mu = 0.9$, $\rho = 0.5$, and $\sigma = 10^{-4}$. We use a MATLAB implementation of minres provided by the authors of [4].
- The MATLAB implementation of the alternating projections method used in the testing in [9].

We use four test matrices, all of which are approximate correlation matrices with unit diagonal.

cor1399 A matrix of dimension 1399 of stock data provided by a fund management company. It is highly rank-deficient and its off-diagonal entries are in the interval $[-0.9644, 1.1574]$.

cor3120 A matrix from the same source as the first one. It has dimension $n = 3120$ and has full rank. The off-diagonal elements are in the interval $[-0.6250, 1.0751]$.

TABLE 5.1
Comparison of different iterative methods in `nearcor_new`; $\text{tol} = 10^{-7}n$.

cor1399			cor3120		
	CG	minres		CG	minres
no preconditioning			no preconditioning		
T_{tot}	213.4	170.9	T_{tot}	2799	1736
T_{mvp}	146.3	104.3	T_{mvp}	2100	1138.6
T_{eig}	62.6	61.7	T_{eig}	662.0	562.8
Iters.	7	7	Iters.	7	6
# mvp	42	30	# mvp	57	31
with preconditioning			with preconditioning		
T_{tot}	142.4	111.0	T_{tot}	1197	905.7
T_{mvp}	76.6	45.2	T_{mvp}	624.2	331.1
T_{eig}	52.8	52.8	T_{eig}	468.3	469.2
Iters.	6	6	Iters.	5	5
# mvp	22	13	# mvp	17	9
Time Pre.	8.83	8.9	Time Pre.	72.9	72.89

Risk-daily, Risk-monthly Matrices from the RiskMetrics database [19]. The documentation says that “The data sets contain consistently calculated volatilities and correlation forecasts for use in estimating market risk. The asset classes covered are government bonds, money markets, swaps, foreign exchange and equity indices (where applicable) for 31 currencies, and commodities.” We obtained two matrices for a one day and a one month horizon assigned to July 15, 2006, which have dimension 387 and a smallest eigenvalue of $-7.92 \cdot 10^{-6}$ and $-4.91 \cdot 10^{-6}$, respectively.

First, we investigate the influence of the iterative solver and of preconditioning. With $\text{tol} = 10^{-7}n$, we solved the first two problems using `nearcor_new`, again with `nearcor_new` with the CG method replacing minres, using the NAG CG suite `f11gd/f11ge/f11gf`, and in each case solving both with and without preconditioning. The results are in Table 5.1, where we report T_{tot} : the total run time (in seconds), T_{mvp} : the time taken to compute all the matrix-vector products $V_k h$ (see (2.5)), T_{eig} : the time to compute the spectral decompositions, Iters.: the number of outer (inexact Newton) iterations, # mvp: the number of matrix-vector products required, and Time Pre.: the time to compute the preconditioner.

Several comments can be made on Table 5.1. First, minres leads to a faster algorithm than the CG method both with and without preconditioning, and results in one fewer iteration for the cor3120 matrix without preconditioning; the reduction in time is largely accounted for by the fewer matrix-vector products. Second, preconditioning brings a useful speedup, amounting for cor3120 to a 48% reduction in time with minres and a 58% reduction with CG. Third, the eigenvalue computations take 32% of the total time for cor3120 with unpreconditioned minres, rising to 52% of the time with preconditioning.

We now take a look at the effect of the choice of eigensolver (see Section 3.5.) The results in Table 5.1 are based on the use of NAG code `f08fc` (divide and conquer). Table 5.2 reports the results for solving the same problems as in Table 5.1 using `nearcor_new` with this eigensolver, `f08fa` (the QR algorithm—as used by MATLAB for the `eig` function), and `f08fd` (dqds/MRRR). The results shown are ratios of total

TABLE 5.2

Ratios of total time taken by `nearcor_new` with `f08fa` and `f08fd` to total time for `nearcor_new` with `f08fc`.

	f08fa	f08fc	f08fd
cor1399	2.1	1.0	1.9
cor3120	2.3	1.0	2.2

times spent to run Algorithm 4.1 with each eigensolver normalized by the run time with the fastest eigensolver, `f08fc`. We see that with `f08fc` the algorithm is twice faster than if the other two eigensolvers are used, for these matrices, and indeed `f08fc` is also the fastest in similar tests we have performed with different matrices [1]. We found this difference quite surprising, and it shows the importance of trying different algorithmic variants of the basic linear algebra “black boxes”. The performance of these codes for tridiagonal matrices is investigated by Demmel, Marques, Parlett, and Vömel [3], who observe that the performance of divide and conquer and dqds/MRRR depends strongly on the particular matrices to which they are applied.

Finally, in Tables 5.3 and 5.4 we compare `nearcor_new` with Qi and Sun’s code `nearcor` and with the alternating projections code from [9], with two different convergence tolerances corresponding to half and full machine precision. The main differences between `nearcor_new` and `nearcor` affecting the run time are the different eigensolvers (`f08fc` for `nearcor_new`, MATLAB’s `eig` for `nearcor`), the different iterative method for computing the Newton direction (minres versus CG), the use of the Jacobi preconditioner in `nearcor_new`, and more optimized MATLAB coding in `nearcor_new`, particularly for the gradient evaluations. In the tables “–” denotes that `nearcor` did not converge or that alternating projections was unable to handle the matrix in a reasonable time (we estimate a run time of several days for cor3120). The speedup of `nearcor_new` over `nearcor` of a factor more than 6 on both problems is significant given that both codes are using the same Newton algorithm.

6. Concluding remarks. Our MATLAB implementation `nearcor_new` of Qi and Sun’s Newton method can solve problems of dimension a few thousand in a few minutes. The run time is dominated by the cost of computing spectral decompositions, and most of the remaining time is spent in computing the matrix-vector products (2.5). Although an extensive computational comparison of all available methods for solving the nearest correlation matrix problem is lacking, the available evidence suggests that the Newton method is the best general-purpose method. Toh [20] reports his code having run times of over 1 hour to solve a nearest correlation matrix problem of dimension 1600 (the maximum size reported therein), whereas `nearcor_new` solves a problem of twice the dimension within 30 minutes. The alternating projections method is very easy to implement and is attractive for small problems and modest convergence tolerances, but in general cannot compete with Newton’s method for efficiency. Extensions of the Newton method to incorporate constraints and to Hadamard weighting have recently been developed [16], [17] and we expect that the ideas herein can be profitably employed in those methods.

Acknowledgements. We thank Jörg Liesen, Marcos Raydan, David Sayers, and David Silvester for helpful discussions on this work.

TABLE 5.3

Time in seconds for the new code `nearcor_new`, `nearcor` (Q_i and Sun), and alternating projections; $\text{tol} = 10^{-7}n$.

	<code>nearcor_new</code>		<code>nearcor</code>		Altern. proj.	
	Time	Iter.	Time	Iter.	Time	Iter
cor1399	96.9	5	378.5	5	529.0	62
cor3120	814.0	4	5256.7	4	–	–
Risk-daily	0.39	0	0.47	0	1.02	2
Risk-monthly	0.36	0	0.53	0	1.22	2

TABLE 5.4

Time in seconds for the new code `nearcor_new`, `nearcor` (Q_i and Sun), and alternating projections; $\text{tol} = 2nu$.

	<code>nearcor_new</code>		<code>nearcor</code>		Altern. proj.	
	Time	Iter.	Time	Iter.	Time	Iter
cor1399	171.9	7	537.1	7	4251.0	494
cor3120	1533.8	6	15685.0	9	–	–
Risk-daily	5.73	5	–	–	30.92	55
Risk-monthly	15.94	10	–	–	18.66	27

REFERENCES

- [1] Rüdiger Borsdorf. A Newton algorithm for the nearest correlation matrix. M.Sc. Thesis, The University of Manchester, Manchester, UK, September 2007. 151 pp. MIMS EPrint 2008.49, Manchester Institute for Mathematical Sciences, The University of Manchester, UK.
- [2] Stephen Boyd and Lin Xiao. Least-squares covariance matrix adjustments. *SIAM J. Matrix Anal. Appl.*, 27(2):532–546, 2005.
- [3] James W. Demmel, Osni A. Marques, Beresford N. Parlett, and Christof Vömel. Performance and accuracy of LAPACK’s symmetric tridiagonal eigensolvers. *SIAM J. Sci. Comput.*, 30(3):1508–1526, 2008.
- [4] Howard C. Elman, David J. Silvester, and Andrew J. Wathen. *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics*. Oxford University Press, 2005. xiii+400 pp. ISBN 0-19-852867-1.
- [5] Christopher C. Finger. A methodology to stress correlations. *RiskMetrics Monitor*, Fourth Quarter:3–11, 1997.
- [6] Anne Greenbaum. *Iterative Methods for Solving Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1997. xiii+220 pp. ISBN 0-89871-396-X.
- [7] Nicholas J. Higham. Computing a nearest symmetric positive semidefinite matrix. *Linear Algebra Appl.*, 103:103–118, 1988.
- [8] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Second edition, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2002. xxx+680 pp. ISBN 0-89871-521-0.
- [9] Nicholas J. Higham. Computing the nearest correlation matrix—A problem from finance. *IMA J. Numer. Anal.*, 22(3):329–343, 2002.
- [10] Dirk K. Knol and Jos M. F. ten Berge. Least-squares approximation of an improper correlation matrix by a proper one. *Psychometrika*, 54(1):53–61, 1989.
- [11] Philip M. Lurie and Matthew S. Goldberg. An approximate method for sampling correlated random variables from partially-specified distributions. *Management Science*, 44(2):203–218, 1998.
- [12] Jerome Malick. A dual approach to solve semidefinite least-squares problems. *SIAM J. Matrix Anal. Appl.*, 26(1):272–284, 2004.
- [13] *NAG Toolbox for MATLAB*. NAG Ltd., Oxford. <http://www.nag.co.uk/>.
- [14] C. C. Paige and M. A. Saunders. Solution of sparse indefinite systems of linear equations. *SIAM J. Numer. Anal.*, 12(4):617–629, 1975.
- [15] Hou-Duo Qi and Defeng Sun. A quadratically convergent Newton method for computing the nearest correlation matrix. *SIAM J. Matrix Anal. Appl.*, 28(2):360–385, 2006.
- [16] Hou-Duo Qi and Defeng Sun. Correlation stress testing for Value-at-Risk: An unconstrained convex optimization approach. Manuscript, March 2007. 35 pp.
- [17] Hou-Duo Qi and Defeng Sun. An augmented Lagrangian dual approach for the H-weighted nearest correlation matrix problem. Manuscript, February 2008. 21 pp.
- [18] Hou-Duo Qi, Zhonghang Xia, and Guangming Xing. An application of the nearest correlation matrix on web document classification. *Journal of Industrial and Management Optimization*, 3(4):701–713, 2007.
- [19] Educational datasets. http://www.riskmetrics.com/stddownload_educ.html.
- [20] Kim-Chuan Toh. An inexact primal-dual path following algorithm for convex quadratic SDP. *Math. Programming*, 112:221–254, 2008.
- [21] Saygun Turkay, Eduardo Epperlein, and Nicos Christofides. Correlation stress testing for value-at-risk. *The Journal of Risk*, 5(4):75–89, 2003.
- [22] A. van der Sluis. Condition numbers and equilibration of matrices. *Numer. Math.*, 14:14–23, 1969.