

EFFICIENT RANK REDUCTION OF CORRELATION MATRICES*

IGOR GRUBIŠIĆ[†] AND RAOUL PIETERSZ[‡]

Abstract. Geometric optimisation algorithms are developed that efficiently find the nearest low-rank correlation matrix. We show, in numerical tests, that our methods compare favourably to the existing methods in the literature. The connection with the Lagrange multiplier method is established, along with an identification of whether a local minimum is a global minimum. An additional benefit of the geometric approach is that any weighted norm can be applied. The problem of finding the nearest low-rank correlation matrix occurs as part of the calibration of multi-factor interest rate market models to correlation.

Key words. geometric optimisation, correlation matrix, rank, LIBOR market model

1. Introduction. The problem of finding the nearest low-rank correlation matrix occurs in areas such as finance, chemistry, physics and image processing. The mathematical formulation of this problem is as follows. Let \mathbb{S}_n denote the set of real symmetric $n \times n$ matrices and let C be a symmetric $n \times n$ matrix with unit diagonal. For $X \in \mathbb{S}_n$ we denote by $X \succeq 0$ that X is positive semidefinite. Let the desired rank $d \in \{1, \dots, n\}$ be given. The problem is then given by

$$\begin{aligned} & \textbf{Find} && X \in \mathbb{S}_n \\ & \textbf{to minimize} && \frac{1}{2} \|C - X\|^2 \\ & \textbf{subject to} && \text{rank}(X) \leq d; \quad X_{ii} = 1, \quad i = 1, \dots, n; \quad X \succeq 0. \end{aligned} \tag{1.1}$$

Here $\|\cdot\|$ denotes a semi-norm on \mathbb{S}_n . The most important instance is

$$\frac{1}{2} \|C - X\|^2 = \frac{1}{2} \sum_{i < j} W_{ij} (C_{ij} - X_{ij})^2, \tag{1.2}$$

where W is a weights matrix consisting of non-negative elements. In words: Find the low-rank correlation matrix X nearest to the given $n \times n$ matrix C . The choice of the semi-norm will reflect what is meant by nearness of the two matrices. The semi-norm in (1.2) is well known in the literature, and it is called the *Hadamard semi-norm*, see Horn & Johnson (1990). Note that the constraint set is non-convex for $d < n$, which makes it not straightforward to solve Problem (1.1) with standard convex optimization methods.

For concreteness, consider the following example. Suppose C is

$$\begin{pmatrix} 1.0000 & -0.1980 & -0.3827 \\ -0.1980 & 1.0000 & -0.2416 \\ -0.3827 & -0.2416 & 1.0000 \end{pmatrix},$$

and W is the full matrix, $W_{ij} = 1$. With the algorithm developed in this paper, we solve (1.1) with C as above and $d = 2$. The algorithm takes as initial input a matrix $X^{(0)}$ of rank 2 or less, for example,

$$X^{(0)} = \begin{pmatrix} 1.0000 & 0.9782 & 0.8982 \\ 0.9782 & 1.0000 & 0.9699 \\ 0.8982 & 0.9699 & 1.0000 \end{pmatrix},$$

and then produces a sequence of points on the constraint set that converges to the point

$$X^* = \begin{pmatrix} 1.0000 & -0.4068 & -0.6277 \\ -0.4068 & 1.0000 & -0.4559 \\ -0.6277 & -0.4559 & 1.0000 \end{pmatrix}$$

*Date: 12 January 2005. We are grateful for the comments of Erik Balder, Pieter Eendebak, Andrea Giaccobe, Antoon Pelsser, Jordan Rizov and seminar participants at ECMI Conference 2004 (Eindhoven, The Netherlands), Leiden University, Quantitative Methods in Finance Conference 2003 (Sydney, Australia), Utrecht University, and Winter School on Mathematical Finance 2003 (Lunteren, The Netherlands). The second author is grateful for the financial support of the Erasmus Center of Financial Research.

[†]Mathematical Institute, University of Utrecht, P.O. Box 80010, 3508 TA Utrecht, The Netherlands (email: grubisic@math.uu.nl, tel: +31 30 2531529)

[‡]Erasmus Research Institute of Management, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands (email: pietersz@few.eur.nl, tel: +31 10 4088932) and Product Development Group (HQ7011), ABN AMRO Bank, P.O. Box 283, 1000 EA Amsterdam, The Netherlands




FIG. 1.1. The shell represents the set of 3×3 correlation matrices of rank 2 or less. The details of this representation are given in Section 6.2.

that solves (1.1). The constraint set and the points generated by the algorithm have been represented in Figure 1.1. The details of this representation are given in Section 6.2. The blue point in the center and the green point represent, respectively, the target matrix C and the solution point X^* . As the figure suggests, the algorithm has fast convergence and the constraint set is a curved space.

This novel technique we propose, is based on geometric optimisation that can locally minimize the objective function in (1.1) and which incorporates the Hadamard semi-norm. In fact, our method can be applied to any sufficiently smooth objective function. Not all other methods available in the literature that aim to solve (1.1) can handle an arbitrary objective function, see the literature review in Section 2. We formulate the problem in terms of Riemannian geometry. This approach allows us to use numerical methods on manifolds that are numerically stable and efficient, in particular the Riemannian-Newton method is applied. We show, for the numerical tests we performed, that the numerical efficiency of geometric optimisation compares favourably to the other algorithms available in the literature. The only drawback of the practical use of geometric optimisation is that the implementation is rather involved. To overcome this drawback, we have made available a MATLAB implementation ‘LRCM min’ (low-rank correlation matrices minimization) at www.few.eur.nl/few/people/pietersz.

We develop a technique to instantly check whether an obtained local minimum is a global minimum, by adaptation of Lagrange multiplier results of Zhang & Wu (2003). The novelty consists of an expression for the Lagrange multipliers given the matrix X , whereas until now only the reverse direction (an expression for the matrix X given the Lagrange multipliers) was known. The fact that one may instantly identify whether a local minimum is a global minimum is very rare for non-convex optimisation problems, and that makes Problem (1.1), which is non-convex for $d < n$, all the more interesting.

Problem (1.1) is important in finance, as it occurs as part of the calibration of the multi-factor LIBOR¹ market model of Brace, Gątarek & Musiela (1997), Miltersen, Sandmann & Sondermann (1997), Jamshidian (1997) and Musiela & Rutkowski (1997). This model is an interest rate derivatives pricing model and it is used in some financial institutions for valuation and risk management of their interest rate derivatives portfolio. The number of stochastic factors needed for the model to fit to the given correlation matrix is equal to the rank of the correlation matrix. This rank can be as high as the number of forward LIBORs in the model, i.e., as high as the dimension of the matrix. The number of LIBORs in the model can grow large in practical applications, for example, a model with over 80 LIBORs is not uncommon.

¹London inter-bank offer rate.

This implies that the number of factors needed to fit the model to the given correlation matrix can be high, too. There is much empirical evidence that the term structure of interest rates is driven by multiple factors (three, four, or even more), see the review article of Dai & Singleton (2003). Though the number of factors driving the term structure may be four or more, the empirical work shows that it is certainly not as high as, say, 80. This is one reason for using a model with a low number of factors. Another reason is the enhanced efficiency when estimating the model-price of an interest rate derivative through Monte Carlo simulation. First, a lower factor model simply requires drawing less random numbers than a higher factor model. Second, the complexity of calculating LIBOR rates over a single time step in a simulation implementation is of order $n \times d$, with n the number of LIBORs and d the number of factors, see Joshi (2003).

The importance of Problem (1.1) in finance has been recognized by many researchers. In fact, the literature review of Section 2 refers to seventeen articles or books addressing the problem.

Due to its generality our method finds locally optimal points for a variety of other objective functions subject to the same constraints. One of the most famous problems comes from physics and is called *Thomson's problem*. The Thomson problem is concerned with minimizing the potential energy of n charged particles on the sphere in \mathbb{R}^3 ($d = 3$). Geometric optimisation techniques have previously been applied to the Thomson problem by Depczynski & Stöckler (1998), but these authors have only considered conjugate gradient techniques on a 'bigger' manifold, in which the freedom of rotation has not been factored out. In comparison, we stress here that our approach considers a lower dimensional manifold, which allows for Newton's algorithm (the latter not developed in Depczynski & Stöckler (1998)). An implementation of geometric optimisation applied to the Thomson problem has also been included in the 'LRCM min' package.

Finally, for a literature review of interest rate models, the reader is referred to Rebonato (2004a).

The paper is organized as follows. In Section 2, the literature is reviewed. In Section 3, the constraints of the problem are formulated in terms of differential geometry. We parameterize the set of correlation matrices of rank at most d with a manifold named the *Cholesky manifold*. This is a canonical space for the optimisation of the arbitrary smooth function subject to the same constraints. In Section 4, the Riemannian structure of the Cholesky manifold is introduced. Formulas are given for parallel transport, geodesics, gradient and Hessian. These are needed for the minimization algorithms, which are made explicit. In Section 5, we discuss the convergence of the algorithms. In Section 6, the application of the algorithms to the problem of finding the nearest low-rank correlation matrix is worked out in detail. In Section 7, we numerically investigate the algorithms in terms of efficiency. Finally, in Section 8, we conclude the paper.

2. Literature review. We mention five algorithms available in the literature for rank reduction of correlation matrices, in reverse chronological order.

First, we mention the majorization approach of Pietersz & Groenen (2004a, b). This reference also contains a literature review covering all algorithms except for the alternating projections algorithm (discussed below). Majorization can handle an entry-weighted objective function and is guaranteed to converge to a stationary point. The rate of convergence is sub-linear.

Second, we point out the Lagrange multiplier algorithm of Zhang & Wu (2003) and Wu (2003). The correlation matrices produced by this algorithm are not guaranteed to converge to a stationary point of objective function F in (1.1). This lack of convergence has been pointed out in Pietersz & Groenen (2004b, Section 2) and therefore it is not necessary to repeat the reasons here.

Third, we consider the alternating projections algorithm of Grubišić (2002) and Morini & Webber (2004). The discussion below of alternating projections applies only to the problem of rank reduction of correlation matrices. The method is based on alternating projections onto the set of $n \times n$ matrices with unit diagonal and onto the set of $n \times n$ matrices of rank d or less. Both these projections can be efficiently calculated. For projection onto the intersection of two convex sets, Dykstra (1983) and Han (1988) have shown that convergence to a minimum can be obtained with alternating projections onto the individual convex sets if a normal vector correction is applied. Their results do not automatically hold for an alternating projections algorithm with normal correction for Problem (1.1)², since for $d < n$ the set of $n \times n$ matrices of rank d or less is non-convex. Indeed, Pietersz & Groenen (2004b, Section 2) and Morini & Webber (2004) report that alternating projections with normal correction may fail in solving Problem

²The algorithm with normal correction for rank reduction has also been studied in Weigel (2004).

(1.1). The alternating projections algorithm without normal correction stated in Grubišić (2002) and Morini & Webber (2004) however always converges to a feasible point, but not necessarily to a stationary point. In fact, in general, the alternating projections method without normal correction does not converge to a stationary point. The algorithm thus does not minimize the objective function in (1.1), it only selects a feasible point satisfying the constraints of (1.1).

The fourth important contribution is due to Higham (2002). The algorithm of Higham (2002) is the alternating projection algorithm with normal correction applied to the case $d = n$, i.e., to the problem of finding the nearest (possibly full-rank) correlation matrix. Note that the space of positive semidefinite matrices of rank n is indeed convex, therefore the alternating projections with normal correction method is guaranteed to converge to the global minimum. Since the case $d < n$ is of primary interest in this paper, the alternating projections method with normal correction will not be considered in the remainder.

Fifth, we mention the parametrization method of Rebonato (1999a, 1999b, 1999c (Section 10), 2002 (Section 9), 2004b (Sections 20.1–20.4)), Brigo (2002), Brigo & Mercurio (2001, Section 6.9) and Rapisarda, Brigo & Mercurio (2002). The set of correlation matrices of rank d or less $\{YY^T : Y \in \mathbb{R}^{n \times d}, \text{diag}(YY^T) = I\}$ is parameterized by trigonometric functions through spherical coordinates $Y_i = Y_i(\theta_i)$ with $\theta_i \in \mathbb{R}^{d-1}$. As a result, the objective value $F(Y)$ becomes a function $F(Y(\theta))$ of the angle parameters θ that live in $\mathbb{R}^{n \times (d-1)}$. Subsequently, ordinary non-linear optimisation algorithms may be applied to minimize the objective value $F(Y(\theta))$ over the angle parameters θ . In essence, this approach is the same as geometric optimisation, except for the key difference of optimising over θ versus over Y . The major benefit of geometric optimisation over the parametrization method is as follows. Consider, for ease of exposition, the case of equal weights. The differential F_Y , in terms of Y , is given simply as $2\psi Y$, with $\psi = YY^T - C$, see (6.1) below. Note that $F_Y = 2\psi Y$ can thus be efficiently calculated. The differential F_θ , in terms of θ however, is $2\psi Y$ multiplied by the differential of Y with respect to θ , by the chain rule of differentiation. The latter differential is less efficient to calculate since it involves numerous sums of trigonometric functions.

An initial feasible point can be found by a principal component analysis and a suitable re-scaling. This method is explained in Section 6.5 below, and is due to Flury (1988). It first appeared in a finance setting in Sidenius (2000) and Hull & White (2000). The starting point for the five above mentioned algorithms is the re-scaled PCA initial feasible point.

2.1. Weighted norms. We mention two reasons for assigning *non-constant* or *non-homogeneous* weights in the objective function of (1.2). First, in our setting C has the interpretation of measured correlation. It can thus be the case that we are more confident of specific entries of the matrix C . Second, the weighted norm of (1.2) has important applications in finance, see, for example, Higham (2002), Rebonato (1999c) and Pietersz & Groenen (2004b).

The semi-norm in the objective function F can be (i) a Hadamard semi-norm with arbitrary weights per element of the matrix, as defined in (1.2), or (ii) a weighted Frobenius norm $\|\cdot\|_{F,\Omega}$ with Ω a positive definite matrix. Here $\|X\|_{F,\Omega}^2 = \text{tr}(X\Omega X^T)$. The weighted Frobenius norm is, from a practical point of view, by far less transparent than the Hadamard or weights-per-entry semi-norm (1.2). The geometric optimisation theory developed in this paper, and most of the algorithms mentioned in Section 2, can be efficiently applied to both cases. The Lagrange multipliers and alternating projections methods however can only be efficiently extended to the case of the weighted Frobenius norm. The reason is that both these methods need to calculate a projection onto the space of matrices of rank d or less. Such a projection, for the weighted Frobenius norm, can be efficiently found by an eigenvalue decomposition. For the Hadamard semi-norm, such an efficient solution is not available, to our knowledge, and as also mentioned in Higham (2002, page 336).

3. Solution methodology with geometric optimisation. Note that Problem (1.1) is a special case of the following more general problem:

$$\begin{aligned} & \mathbf{Find} && X \in \mathbb{S}_n \\ & \mathbf{to minimize} && \tilde{F}(X) \\ & \mathbf{subject to} && \text{rank}(X) \leq d; \quad X_{ii} = 1, i = 1, \dots, n; \quad X \succeq 0. \end{aligned} \tag{3.1}$$

In this paper methods will be developed to solve Problem (3.1) for the case when \tilde{F} is twice continuously differentiable. In the remainder of the paper, we assume $d > 1$, since for $d = 1$ the constraint set consists of a finite number (2^{n-1}) of points.

3.1. Basic idea. The idea for solving Problem (3.1) is to parameterize the constraint set by a manifold, and subsequently utilize the recently developed algorithms for optimisation over manifolds, such as Newton's algorithm and conjugate gradient algorithms. Such *geometric optimisation* has been developed by Smith (1993).

In Section 3.2, the constraint set is equipped with a topology, and we make an identification with a certain quotient space. In Section 3.3, it will be shown that the constraint set as such is not a manifold; however a dense subset is shown to be a manifold, namely the set of matrices of rank exactly d . Subsequently, in Section 3.4, we will define a larger manifold (named *Cholesky manifold*), of the same dimension as the rank- d manifold, that maps surjectively to the constraint set. We may apply geometric optimisation on the Cholesky manifold. The connection between minima on the Cholesky manifold and on the constraint set will be established.

3.2. Topological Structure. In this section, the set of $n \times n$ correlation matrices of rank d or less is equipped with the subspace topology from \mathbb{S}_n . We subsequently establish a homeomorphism (i.e., a topological isomorphism) between the latter topological space and the quotient space of n products of the $d-1$ sphere over the group of orthogonal transformations of \mathbb{R}^d . Intuitively the correspondence is as follows: We can associate with an $n \times n$ correlation matrix of rank d a configuration of n points of unit length in \mathbb{R}^d such that the inner product of points i and j is entry (i, j) of the correlation matrix. Any orthogonal rotation of the configuration does not alter the associated correlation matrix. This idea is developed more rigorously below.

DEFINITION 3.1. *The set of symmetric $n \times n$ correlation matrices of rank at most d is defined by*

$$C_{n,d} = \{ X \in \mathbb{S}_n ; \text{diag}(X) = I_n, \text{rank}(X) \leq d, X \succeq 0 \}.$$

Here I_n denotes the $n \times n$ identity matrix and diag denotes the map $\mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$, $\text{diag}(X)_{ij} = \delta_{ij} X_{ij}$, where δ_{ij} denotes the Kronecker delta.

The set $C_{n,d}$ is a subset of \mathbb{S}_n . The latter space is equipped with the Frobenius norm $\|\cdot\|_F$, which in turn defines a topology. We equip $C_{n,d}$ with the subspace topology.

In the following, the product of n unit spheres S^{d-1} is denoted by $T_{n,d}$. Elements of $T_{n,d}$ are denoted as a matrix $Y \in \mathbb{R}^{n \times d}$, with each row vector Y_i of unit length. Denote by O_d the group of orthogonal transformations of d -space. Elements of O_d are denoted by a $d \times d$ orthogonal matrix Q .

DEFINITION 3.2. *We define the following right O_d -action³ on $T_{n,d}$:*

$$T_{n,d} \times O_d \rightarrow T_{n,d}, (Y, Q) \mapsto YQ. \quad (3.2)$$

An equivalence class $\{YQ : Q \in O_d\}$ associated with $Y \in T_{n,d}$ is denoted by $[Y]$ and it is called the orbit of Y . The quotient space $T_{n,d}/O_d$ is denoted by $M_{n,d}$. The canonical projection $T_{n,d} \rightarrow T_{n,d}/O_d = M_{n,d}$ is denoted by π . Define the map⁴ Ψ as

$$M_{n,d} \xrightarrow{\Psi} C_{n,d}, \Psi([Y]) = YY^T.$$

Consider a map Φ in the inverse direction of Ψ ,

$$C_{n,d} \xrightarrow{\Phi} M_{n,d},$$

defined as follows: For $X \in C_{n,d}$ take $Y \in T_{n,d}$ such that $YY^T = X$. Such Y can always be found as will be shown in Theorem 3.3 below. Then set $\Phi(X) = [Y]$. It will be shown in Theorem 3.3 that this map is well defined. Finally, define the map $s : T_{n,d} \rightarrow C_{n,d}$, $s(Y) = YY^T$.

The following theorem relates the spaces $C_{n,d}$ and $M_{n,d}$; the proof has been deferred to Appendix A.

THEOREM 3.3. *Consider the following diagram*

$$\begin{array}{ccc} T_{n,d} & \xrightarrow{s} & C_{n,d} \\ \pi \downarrow & \searrow \Phi & \nearrow \Psi \\ M_{n,d} = T_{n,d}/O_d & & \end{array} \quad (3.3)$$

³It is trivially verified that the map thus defined is indeed an O_d smooth action: $YI_d = Y$ and $Y(Q_1Q_2)^{-1} = (YQ_2^{-1})Q_1^{-1}$. Standard matrix multiplication is smooth.

⁴Although rather obvious, it will be shown in Theorem 3.3 that this map is well defined.

with the objects and maps as in Definitions 3.1 and 3.2. We have the following:

- (i) The maps Ψ and Φ are well defined.
- (ii) The diagram is commutative, i.e., $\Psi \circ \pi = s$ and $\Phi \circ s = \pi$.
- (iii) The map Ψ is a homeomorphism with inverse Φ .

3.3. A dense part of $M_{n,d}$ equipped with a differentiable structure. For an exposition on differentiable manifolds, the reader is referred to do Carmo (1992). It turns out that $M_{n,d}$ is not a manifold, but a so-called *stratified space*, see, e.g., Duistermaat & Kolk (2000). However there is a subspace of $M_{n,d}$ that is a manifold, which is the manifold of equivalence classes of matrices of *exactly* rank d . The proof of the following proposition has been deferred to Appendix B.

PROPOSITION 3.4. Let $T_{n,d}^* \subset T_{n,d}$ be the subspace defined by

$$T_{n,d}^* = \{ Y \in T_{n,d} : \text{rank}(Y) = d \}.$$

Then we have the following:

1. $T_{n,d}^*$ is a sub-manifold of $T_{n,d}$.
2. Denote by $M_{n,d}^*$ the quotient space $T_{n,d}^*/O_d$. Then $M_{n,d}^*$ is a manifold of dimension $n(d-1) - d(d-1)/2$.

As shown in Proposition 3.4, a subset $M_{n,d}^*$ of $M_{n,d}$ is a manifold. In the following, we will study charts given by sections of the manifold $M_{n,d}^*$ that will ultimately lead to the final manifold over which will be optimized.

A *section* on $M_{n,d}^*$ is a map $\sigma : \mathcal{U} \rightarrow T_{n,d}^*$, with \mathcal{U} open in $M_{n,d}^*$, such that $\pi \circ \sigma = \text{id}_{M_{n,d}^*}$. Such a map singles out a unique matrix in each equivalence class. In our case we can explicitly give such a map σ . Let $[Y]$ in $M_{n,d}^*$, and let I denote a subset of $\{1, \dots, n\}$ with exactly d elements, such that $\dim(\text{span}(\{Y_i : i \in I\})) = d$, for $Y \in [Y]$. Note that I is well defined since any two $Y^{(1)}, Y^{(2)} \in [Y]$ are coupled by an orthogonal transformation, see the proof of Theorem 3.3, and orthogonal transformations preserve independence. The collection of all such I is denoted by \mathcal{I}_Y . It is readily verified that \mathcal{I}_Y is not empty. Let \prec denote the lexicographical ordering, then (\mathcal{I}_Y, \prec) is a well-ordered set. Thus we can choose the smallest element, denoted by $J(Y) = (j_1, \dots, j_d)$. Define $\tilde{Y} \in \mathbb{R}^{d \times d}$ by taking the rows of Y from J_Y , thus $\tilde{Y}_i = Y_{j_i}$. Define $\tilde{X} = \tilde{Y}\tilde{Y}^T$. Since \tilde{X} is positive definite, Cholesky decomposition can be applied to \tilde{X} , see for example Golub & van Loan (1996, Theorem 4.2.5), to obtain a unique lower-triangular matrix \bar{Y} such that $\bar{Y}\bar{Y}^T = \tilde{X}$ and $\bar{Y}_{ii} > 0$. By Theorem 3.3, there exists a unique orthogonal matrix $Q \in O_d$ such that $\bar{Y} = \tilde{Y}Q$. Define $Y^* = \bar{Y}Q$. Note that Y^* is lower-triangular, since for $i \notin J_Y$, let p be the largest integer such that $i > j_p$, then Y_i^* is dependent on Y_1^*, \dots, Y_p^* , as J_Y is the smallest element from \mathcal{I}_Y , which implies a lower-triangular form for Y^* . Then define $\mathcal{U}_Y = \{[Z] : J(Y) \in \mathcal{I}_Z\} \subset M_{n,d}^*$. It is obvious that \mathcal{U}_Y and $\pi^{-1}(\mathcal{U}_Y)$ are open in the corresponding topologies. Then

$$\sigma_Y : \mathcal{U}_Y \rightarrow \pi^{-1}(\mathcal{U}_Y), [Z] \mapsto Z^*, \quad (3.4)$$

is a section of \mathcal{U}_Y at Y . The following proposition shows that the sections are the charts of the manifold $M_{n,d}^*$. The proof has been deferred to Appendix C.

PROPOSITION 3.5. The differentiable structure on $M_{n,d}^*$ is the one which makes $\sigma_Y : \mathcal{U}_Y \rightarrow \sigma(\mathcal{U}_Y)$ into a diffeomorphism.

3.4. The Cholesky manifold. In this section, we will show that, for the purpose of optimisation, it is sufficient to perform the optimisation on a compact manifold that contains one of the sections. For simplicity we choose the section σ_Y where $J(Y) = \{1, \dots, d\}$. The image $\sigma_Y(\mathcal{U}_Y)$ is a smooth sub-manifold of $T_{n,d}$ with the following representation in $\mathbb{R}^{n \times d}$

$$\left\{ \left(\begin{array}{ccc} Y_1^T & \dots & Y_n^T \end{array} \right)^T : Y_1 = (1, 0, \dots, 0); Y_i \in S_+^{i-1}, i = 2, \dots, d; Y_i \in S^{d-1}, i = d+1, \dots, n \right\},$$

with S_+^{i-1} embedded in \mathbb{R}^d by the first i coordinates such that coordinate i is bigger than 0 and with the remaining coordinates set to zero. Also, S^{d-1} is similarly embedded in \mathbb{R}^d . We can consider the map $s : T_{n,d} \rightarrow C_{n,d}$ restricted to $\sigma_Y(\mathcal{U}_Y)$, which is differentiable since $\sigma_Y(\mathcal{U}_Y)$ is a sub-manifold of $T_{n,d}$. The map $s|_{\sigma_Y(\mathcal{U}_Y)}$ is a homeomorphism, in virtue of Theorem 3.3.

For the purpose of optimisation, we need a compact manifold which is surjective with $C_{n,d}$. Define the following sub-manifold of $T_{n,d}$ of dimension $n(d-1) - d(d-1)/2$,

$$\text{Chol}_{n,d} = \left\{ Y \in \mathbb{R}^{n \times d} : Y_1 = (1, 0, \dots, 0); Y_i \in S^{i-1}, i = 2, \dots, d; Y_i \in S^{d-1}, i = d+1, \dots, n \right\},$$

which we call the Cholesky manifold. The Cholesky parametrization has been considered before by Rapisarda et al. (2002), but these authors do not consider non-Euclidean geometric optimisation. The map $s|_{\text{Chol}_{n,d}}$ is surjective, in virtue of the following theorem, the proof of which has been relegated to Appendix D.

THEOREM 3.6. *If $X \in C_{n,d}$, then there exists a $Y \in \text{Chol}_{n,d}$ such that $YY^T = X$.*

A function \tilde{F} on $C_{n,d}$ can be considered on $\text{Chol}_{n,d}$, too, via the composition

$$\text{Chol}_{n,d} \xrightarrow{s} C_{n,d} \xrightarrow{\tilde{F}} \mathbb{R}, \quad Y \mapsto YY^T \mapsto \tilde{F}(YY^T).$$

From here on, we will write $F(Y) := \tilde{F}(YY^T)$ viewed as a function on $\text{Chol}_{n,d}$.

For a *global* minimum $F(Y)$ on $\text{Chol}_{n,d}$, we have that YY^T attains a global minimum of \tilde{F} on $C_{n,d}$, since the map $s : \text{Chol}_{n,d} \rightarrow C_{n,d}$ is surjective. For a *local* minimum, we have the following theorem. The proof has been deferred to Appendix E.

THEOREM 3.7. *The point Y attains a local minimum of F on $\text{Chol}_{n,d}$ if and only if YY^T attains a local minimum of \tilde{F} on $C_{n,d}$.*

These considerations on global and local minima on $\text{Chol}_{n,d}$ show that, to optimize \tilde{F} over $C_{n,d}$, we might as well optimize F over the manifold $\text{Chol}_{n,d}$. For the optimisation of \tilde{F} over $C_{n,d}$, there is no straightforward way to use numerical methods such as Newton and conjugate gradient, since they require a notion of differentiability, but for optimisation of F on $\text{Chol}_{n,d}$, we can use such numerical methods.

3.5. Choice of representation. In principle, we could elect another manifold \tilde{M} and a surjective open map $\tilde{M} \rightarrow C_{n,d}$. We insist however on explicit knowledge of the geodesics and parallel transport, for this is essential to obtaining an efficient algorithm. We found that if we choose the Cholesky manifold then convenient expressions for geodesics, etc., are obtained. Moreover, the Cholesky manifold has the minimal dimension, i.e., $\dim(\text{Chol}_{n,d}) = \dim(M_{n,d}^*)$.

In the next section, the geometric optimisation tools are developed for the Cholesky manifold.

4. Optimisation over the Cholesky manifold. For the development of minimization algorithms on a manifold, certain objects of the manifold need to be calculated explicitly, such as geodesics, parallel transport, etc. In this section, these objects are introduced and made explicit for $\text{Chol}_{n,d}$.

From a theoretical point of view, it does not matter which coordinates we choose to derive the geometrical properties of a manifold. For the numerical computations however this choice is essential because the simplicity of formulas for the geodesics and parallel transport depends on the chosen coordinates. We found that simple expressions are obtained when $\text{Chol}_{n,d}$ is viewed as a sub-manifold of $T_{n,d}$, which, in turn, is viewed as a subset of the *ambient space* $\mathbb{R}^{n \times d}$. This representation reveals that, to calculate geodesics and parallel transport on $\text{Chol}_{n,d}$, it is sufficient to calculate these on a single sphere.

The tangent space of the manifold $\text{Chol}_{n,d}$ at a point $Y \in \text{Chol}_{n,d}$ is denoted by $T_Y \text{Chol}_{n,d}$. A tangent vector at a point Y is an element of $T_Y \text{Chol}_{n,d}$ and is denoted by Δ .

4.1. Riemannian structure. We start with a review of basic concepts of Riemannian geometry. Our exposition follows do Carmo (1992). Let M be an m -dimensional differentiable manifold. A *Riemannian structure* on M is a smooth map $Y \mapsto \langle \cdot, \cdot \rangle_Y$, which for every $Y \in M$ assigns an inner product $\langle \cdot, \cdot \rangle_Y$ on $T_Y M$, the tangent space at point Y . A *Riemannian manifold* is a differentiable manifold with a Riemannian structure.

Let F be a smooth function on a Riemannian manifold M . Denote the *differential* of F at a point Y by F_Y . Then F_Y is a linear functional on $T_Y M$. In particular, let $\gamma(t)$, $t \in (-\varepsilon, \varepsilon)$, be a smooth curve on M such that $\gamma(0) = Y$ and $\dot{\gamma}(0)$ expressed in a coordinate chart $(\mathcal{U}, x_1, \dots, x_m)$ is equal to Δ , then $F_Y(\Delta)$ can be expressed in this coordinate chart by

$$F_Y(\Delta) = \sum_{i=1}^m \frac{\partial}{\partial x^i} (F \circ x_i^{-1})(\gamma) \Big|_{t=0} \quad (4.1)$$

The linear space of linear functionals on $T_Y M$ (the dual space) is denoted by $(T_Y M)^*$. A *vector field* is a map on M that selects a tangent $\Delta \in T_Y M$ at each point $Y \in M$. The Riemannian structure induces an isomorphism between $T_Y M$ and $(T_Y M)^*$, which guarantees the existence of a unique vector field on M , denoted by $\text{grad}F$, such that

$$F_Y(\Delta) = \langle \text{grad}F, X \rangle_Y \quad \text{for all } X \in T_Y M. \quad (4.2)$$

This vector field is called the *gradient* of F . Also, for Newton and conjugate gradient methods, we have to use second order derivatives. In particular, we need to be able to differentiate vector fields. To do this on a general manifold, we need to equip the manifold with additional structure, namely the *connection*. A connection on a manifold M is a rule $\nabla \cdot$ which assigns to each two vector fields X_1, X_2 on M a vector field $\nabla_{X_1} X_2$ on M , satisfying the following two conditions:

$$\nabla_{F X_1 + G X_2} X_3 = F \nabla_{X_1} X_3 + G \nabla_{X_2} X_3, \quad \nabla_{X_1} (F X_2) = F \nabla_{X_1} (X_2) + (X_1 F) X_2, \quad (4.3)$$

for F, G smooth functions on M and X_1, X_2, X_3 vector fields on M .

Let $\gamma(t)$ be a smooth curve on M with tangent vector $X_1(t) = \dot{\gamma}(t)$. A given family $X_2(t)$ of tangent vectors at the points $\gamma(t)$ is said to be *parallel transported* along γ if

$$\nabla_{X_1} X_2 = 0 \quad \text{on } \gamma(t), \quad (4.4)$$

where X_1, X_2 are vector fields that coincide with $X_1(t)$ and $X_2(t)$, respectively, on $\gamma(t)$. If the tangent vector $X_1(t)$ itself is parallel transported along $\gamma(t)$ then the curve $\gamma(t)$ is called a *geodesic*. In particular, if $(\mathcal{U}, x_1, \dots, x_m)$ is a coordinate chart on M and $\{X_1, \dots, X_m\}$ the corresponding vector fields then the affine connection ∇ on \mathcal{U} can be expressed by

$$\nabla_{X_i} X_j = \sum_{k=1}^m \Gamma_{i,j}^k X_k. \quad (4.5)$$

The functions $\Gamma_{i,j}^k$ are smooth functions, called the *Christoffel symbols* for the connection. In components, the geodesic equation becomes

$$\ddot{x}_k + \sum_{i,j=1}^m \Gamma_{i,j}^k \dot{x}_i \dot{x}_j = 0, \quad (4.6)$$

where x_k are the coordinates of $\gamma(t)$. On a Riemannian manifold there is a unique torsion free connection compatible with the metric, called the *Levi-Civita* connection. This means that Christoffel symbols can be expressed as functions of a metric on M . Note also that (4.6) implies that, once we have determined the equation for the geodesic we can simply read off Christoffel symbols. With respect to an induced metric the geodesic is the curve of shortest length between two points on a manifold. For a manifold embedded in Euclidean space an equivalent characterization of a geodesic is that the acceleration vector at each point along a geodesic is normal to the manifold so long as the curve is traced with uniform speed.

We start by defining Riemannian structures for $T_{n,d}$ and for the Cholesky manifold $\text{Chol}_{n,d}$. We use the Levi-Civita connection, associated to the metric defined as follows on the tangent spaces. Both tangent spaces are identified with suitable subspaces of the ambient space $\mathbb{R}^{n \times d}$, and subsequently the inner product for two tangents Δ_1, Δ_2 is defined as

$$\langle \Delta_1, \Delta_2 \rangle = \text{tr} \Delta_1 \Delta_2^T, \quad (4.7)$$

which is the Frobenius inner product for $n \times d$ matrices. Note that, in our special case, the inner product $\langle \cdot, \cdot \rangle_Y$ is independent of the point Y ; therefore we suppress the dependency on Y .

4.2. Normal and tangent spaces. An equation determining tangents to $T_{n,d}$ at a point Y can be obtained by differentiating $\text{diag}(Y Y^T) = I_n$ yielding $\text{diag}(Y \Delta^T + \Delta Y^T) = 0$, i.e., $\text{diag}(\Delta Y^T) = 0$. The dimension of the tangent space is $n(d-1)$. The *normal space* at the point Y is defined to be the orthogonal complement of the tangent space at the point Y , i.e., it consists of the matrices N , for which $\text{tr} \Delta N^T = 0$ for all Δ in the tangent space. It follows that the normal space is n dimensional. It is straightforward to verify that if $N = DY$, where D is $n \times n$ diagonal, then N is in the normal space. Since the dimension of the space of such matrices is n , we see that the normal space $N_Y T_{n,d}$ at $Y \in T_{n,d}$ is given by

$$N_Y T_{n,d} = \{ DY ; D \in \mathbb{R}^{n \times n} \text{ diagonal} \}.$$

The projections $\pi_{N_Y T_{n,d}}$ and $\pi_{T_Y T_{n,d}}$ onto the normal and tangent spaces of $T_{n,d}$ are given by

$$\pi_{N_Y T_{n,d}}(\Delta) = \text{diag}(\Delta Y^T) Y \quad \text{and} \quad \pi_{T_Y T_{n,d}}(\Delta) = \Delta - \text{diag}(\Delta Y^T) Y,$$

respectively. The projection $\pi_{T_Y \text{Chol}_{n,d}}$ onto the tangent space of $\text{Chol}_{n,d}$ is given by

$$\pi_{T_Y \text{Chol}_{n,d}}(\Delta) = \zeta(\pi_{T_Y T_{n,d}}(\Delta)) \quad \text{with } \zeta(\Delta) \text{ defined by } \zeta(\Delta)_{ij} = \begin{cases} 0 & \text{for } j > i \text{ or } i = j = 1, \\ \Delta_{ij} & \text{otherwise.} \end{cases} \quad (4.8)$$

4.3. Geodesics. It is convenient to work with the coordinates of the ambient space $\mathbb{R}^{n \times d}$. In this coordinate system, geodesics on $T_{n,d}$ with respect to the Levi-Civita connection obey the second order differential equation

$$\ddot{Y} + \Gamma_Y(\dot{Y}, \dot{Y}) = 0, \quad \text{with } \Gamma_Y(\Delta_1, \Delta_2) := \text{diag}(\Delta_1 \Delta_2^T) Y. \quad (4.9)$$

To see this, we begin with the condition that $Y(t)$ remains on $T_{n,d}$,

$$\text{diag}(Y Y^T) = I_n. \quad (4.10)$$

Differentiating this equation twice, we obtain,

$$\text{diag}(\ddot{Y} Y^T + 2\dot{Y} \dot{Y}^T + Y \ddot{Y}^T) = 0. \quad (4.11)$$

In order for $Y(\cdot)$ to be a geodesic, $\ddot{Y}(t)$ must be in the normal space at $Y(t)$, i.e.,

$$\ddot{Y}(t) = D(t) Y(t) \quad (4.12)$$

for some diagonal matrix $D(t)$. To obtain an expression for D , substitute (4.12) into (4.11), which yields (4.9).

The function Γ_Y is the matrix notation of the Christoffel symbols, Γ_{ij}^k , with respect to E_1, \dots, E_{nd} , the standard basis vectors of $\mathbb{R}^{n \times d}$. More precisely, $\nabla_{E_i} E_j = \sum_{k=1}^{nd} \Gamma_{ij}^k E_k$ with Γ_{ij}^k defined by $\langle \Gamma_Y(X_1, X_2), E_k \rangle = \sum_{i,j=1}^{nd} \Gamma_{ij}^k (X_1)_i (X_2)_j$.

The geodesic at $Y(0) \in T_{n,d}$ in the direction $\Delta \in T_{Y(0)} T_{n,d}$ is given by,

$$Y_i(t) = \cos(\|\Delta_i\|t) Y_i(0) + \frac{1}{\|\Delta_i\|} \sin(\|\Delta_i\|t) \Delta_i. \quad (4.13)$$

for $i = 1, \dots, n$, per component on the sphere. By differentiating, we obtain an expression for the evolution of the tangent along the geodesic:

$$\dot{Y}_i(t) = -\|\Delta_i\| \sin(\|\Delta_i\|t) Y_i(0) + \cos(\|\Delta_i\|t) \Delta_i. \quad (4.14)$$

Since $\text{Chol}_{n,d}$ is a Riemannian sub-manifold of $T_{n,d}$ it has the same geodesics.

4.4. Parallel transport along a geodesic. We consider this problem per component on the sphere. If $\Delta^{(2)} \in T_{Y^{(1)}} T_{n,d}$ is parallel transported along a geodesic starting from $Y^{(1)}$ in the direction of $\Delta^{(1)} \in T_{Y^{(1)}} T_{n,d}$, then decompose $\Delta^{(2)}$ in terms of $\Delta^{(1)}$,

$$\Delta_i^{(2)}(t) = \langle \Delta_i^{(1)}(0), \Delta_i^{(2)}(0) \rangle \Delta_i^{(1)}(t) + R_i, \quad R_i \perp \Delta_i^{(1)}(0).$$

Then $\Delta_i^{(1)}(t)$ changes according to (4.14) and R_i remains unchanged. Parallel transport from $Y^{(1)}$ to $Y^{(2)}$ defines an isometry $\tau(Y^{(1)}, Y^{(2)}) : T_{Y^{(1)}} T_{n,d} \rightarrow T_{Y^{(2)}} T_{n,d}$. When it is clear in between which two points is transported, then parallel transport is denoted simply by τ . Since $\text{Chol}_{n,d}$ is a Riemannian sub-manifold of $T_{n,d}$ it has the same equations for parallel transport.

4.5. The gradient. Since $\text{Chol}_{n,d}$ is a sub-manifold of $\mathbb{R}^{n \times d}$ we can use coordinates of $\mathbb{R}^{n \times d}$ to express the differential F_Y of F at the point Y , namely $(F_Y)_{ij} = \frac{\partial F}{\partial Y_{ij}}$. The gradient $\text{grad} F$ of a function F on $\text{Chol}_{n,d}$ can be determined by (4.2). It follows that,

$$\text{grad} F = \pi_{T_Y \text{Chol}_{n,d}}(F_Y) = \zeta(F_Y - \text{diag}(F_Y Y^T) Y). \quad (4.15)$$

4.6. Hessian. The Hessian $\text{Hess} F$ of a function F is a second covariant derivative of F . More precisely, let Δ_1, Δ_2 be two vector fields, then

$$\text{Hess} F(\Delta_1, \Delta_2) = \langle \nabla_{\Delta_1} \text{grad} F, \Delta_2 \rangle$$

In local coordinates of $\mathbb{R}^{n \times d}$

$$\text{Hess} F(\Delta_1, \Delta_2) = F_{YY}(\Delta_1, \Delta_2) - \langle F_Y, \Gamma_Y(\Delta_1, \Delta_2) \rangle, \quad (4.16)$$

where

$$F_{YY}(\Delta_1, \Delta_2) = \frac{d}{dt} \frac{d}{ds} \Big|_{t=s=0} F(Y(t, s)), \text{ with } \frac{d}{dt} \Big|_{t=0} Y = \Delta_1, \frac{d}{ds} \Big|_{s=0} Y = \Delta_2.$$

Newton's method requires inverting the Hessian at minus the gradient, therefore we need to find the tangent Δ to $\text{Chol}_{n,d}$ such that

$$\text{Hess}F(\Delta, X) = \langle -\text{grad}F, X \rangle, \text{ for all tangents } X \text{ to } \text{Chol}_{n,d}. \quad (4.17)$$

To solve (4.17), it is convenient to calculate the unique tangent vector $\mathcal{H} = \mathcal{H}(\Delta)$ satisfying

$$\text{Hess}F(\Delta, X) = \langle \mathcal{H}, X \rangle, \text{ for all tangents } X \text{ to } \text{Chol}_{n,d},$$

since then the Newton Equation (4.17) becomes $\mathcal{H}(\Delta) = -\text{grad}F$. From (4.9) and (4.16), we obtain

$$\mathcal{H}(\Delta) = \pi_{T_Y \text{Chol}_{n,d}}(F_{YY}(\Delta)) - \text{diag}(F_Y Y^T) \Delta, \quad (4.18)$$

where the notation $F_{YY}(\Delta)$ means the tangent vector satisfying

$$F_{YY}(\Delta) = \frac{d}{dt} \Big|_{t=0} F_Y(Y(t)), \quad \dot{Y}(0) = \Delta.$$

4.7. Algorithms. We are now in a position to state the conjugate gradient algorithm, given as Algorithm 1, and the Newton algorithm, given as Algorithm 2, for optimisation over the Cholesky manifold. These algorithms are instances of the geometric programs presented in Smith (1993), for the particular case of the Cholesky manifold.

5. Discussion of convergence properties. In this section, we discuss convergence properties of the geometric programs: global convergence and the local rate of convergence.

5.1. Global convergence. First, we discuss global convergence for the Riemannian-Newton algorithm. It is well known that the Newton algorithm, as displayed in Algorithm 2, is not globally convergent to a local minimum. Moreover, the steps in Algorithm 2 may even not be well defined, because the Hessian mapping could be singular. The standard way to resolve these issues, is to introduce jointly a steepest descent algorithm. So Algorithm 2 is adjusted in the following way. When the new search direction $\Delta^{(k)}$ has been calculated, then we also consider the steepest descent search direction $\Delta_{\text{Steep}}^{(k)} = -\text{grad}F(Y^{(k)})$. Subsequently, a line minimization of the objective value is performed in both directions, $\Delta^{(k)}$ and $\Delta_{\text{Steep}}^{(k)}$. We then take as the next point of the algorithm whichever search direction finds the point with lowest objective value. Such a steepest descent method with line minimization is well known to have guaranteed convergence to a local minimum.

Second, we discuss global convergence for conjugate gradient algorithms. For the Riemannian case, we have not seen any global convergence results for conjugate gradient algorithms in the literature. Therefore we focus on the results obtained for the flat-Euclidean case. Zoutendijk (1970) and Al-Baali (1985) establish global convergence of the Fletcher & Reeves (1964) conjugate gradient method with line minimization. Gilbert & Nocedal (1992) establish alternative line search minimizations that guarantee global convergence of the Polak & Ribière (1969) conjugate gradient method.

5.2. Local rate of convergence. Local rates of convergence for geometric optimisation algorithms are established in Smith (1993), Edelman, Arias & Smith (1998) and Dedieu, Priouret & Malajovich (2003).

In Theorem 3.3 of Smith (1993), the following result is established for the Riemannian-Newton method. If \hat{Y} is a non-degenerate stationary point, then there exists an open set \mathcal{U} containing \hat{Y} , such that starting from any $Y^{(0)}$ in \mathcal{U} , the sequence of points produced by Algorithm 2 converges quadratically to \hat{Y} .

In Theorem 4.3 of Smith (1993), the following result is stated for the Riemannian Fletcher & Reeves (1964) and Polak & Ribière (1969) conjugate gradient methods. Suppose \hat{Y} is a non-degenerate stationary point such that the Hessian at \hat{Y} is positive definite. Suppose $\{Y^{(j)}\}_{j=0}^{\infty}$ is a sequence of points, generated by Algorithm 1, converging to \hat{Y} . Then, for sufficiently large j , the sequence $\{Y^{(j)}\}_{j=0}^{\infty}$ has $\dim(\text{Chol}_{n,d})$ -steps quadratic convergence to \hat{Y} .

As a numerical illustration, convergence runs have been displayed in Figure 5.1, for reducing a 10×10 correlation matrix to rank 3. The following algorithms are compared:

Algorithm 1 Conjugate gradient for minimizing $F(Y)$ on $\text{Chol}_{n,d}$

Input: $Y^{(0)}, F(\cdot)$.

Require: $Y^{(0)}$ such that $Y^{(0)}(Y^{(0)})^T = I_n$.

Compute $G^{(0)} = \text{grad}F(Y^{(0)}) = \zeta(F_Y - \text{diag}(F_Y(Y^{(0)})^T)Y^{(0)})$ and set $H^{(0)} = -G^{(0)}$.

for $k = 0, 1, 2, \dots$ **do**

Minimize $F(Y^{(k)}(t))$ over t where $Y^{(k)}(t)$ is a geodesic on $\text{Chol}_{n,d}$ starting from $Y^{(k)}$ in the direction of $H^{(k)}$.

Set $t_k = t_{\min}$ and $Y^{(k+1)} = Y^{(k)}(t_k)$.

Compute $G^{(k+1)} = \text{grad}F(Y^{(k+1)}) = \zeta(F_Y - \text{diag}(F_Y(Y^{(k+1)})^T)Y^{(k+1)})$.

Parallel transport tangent vectors $H^{(k)}$ and $G^{(k)}$ to the point $Y^{(k+1)}$.

Compute the new search direction

$$H^{(k+1)} = -G^{(k+1)} + \gamma_k \tau H^{(k)} \quad \text{where} \quad \begin{cases} \gamma_k = \frac{\langle G^{(k+1)} - \tau G^{(k)}, G^{(k+1)} \rangle}{\langle G^{(k)}, G^{(k)} \rangle}, & \text{Polak \& Ribière (1969),} \\ \gamma_k = \frac{\|G^{(k+1)}\|^2}{\|G^{(k)}\|^2}, & \text{Fletcher \& Reeves (1964).} \end{cases}$$

Reset $H^{(k+1)} = -G^{(k+1)}$ if $k + 1 \equiv 0 \pmod{n(d-1) - \frac{1}{2}d(d-1)}$.

end for

Algorithm 2 Newton's method for minimizing $F(Y)$ on $\text{Chol}_{n,d}$.

Input: $Y^{(0)}, F(\cdot)$.

Require: $Y^{(0)}$ such that $\text{diag}(Y^{(0)}(Y^{(0)})^T) = I_n$.

for $k = 0, 1, 2, \dots$ **do**

Compute $G^{(k)} = \text{grad}F(Y^{(k)}) = \zeta(F_Y - \text{diag}(F_Y Y^T)Y)$.

Compute $\Delta^{(k)} = -\mathcal{H}^{-1}G^{(k)}$, i.e. $\Delta^{(k)} \in T_Y \text{Chol}_{n,d}$ and

$$\zeta\left(F_{YY}(\Delta^{(k)}) - \text{diag}\left(F_{YY}(\Delta^{(k)})(Y^{(k)})^T\right)Y^{(k)}\right) - \text{diag}(F_Y(Y^{(k)})^T)\Delta^{(k)} = -G^{(k)}.$$

Move from $Y^{(k)}$ in direction $\Delta^{(k)}$ to $Y^{(k)}(1)$ along the geodesic.

Set $Y^{(k+1)} = Y^{(k)}(1)$.

end for

1. Steepest descent, for which the search direction $H^{(k+1)}$ in Algorithm 1 is equal to $-G^{(k+1)}$, i.e., to minus the gradient. The steepest descent method has a linear local rate of convergence, see Smith (1993, Theorem 2.3).
2. PRCG, Polak-Ribière conjugate gradient.
3. FRCG, Fletcher-Reeves conjugate gradient.
4. Newton.
5. Lev.-Mar., the Levenberg (1944) & Marquardt (1963) method, which is a Newton-type method.

The code that is used for this test is the package 'LRCM min', to be discussed in Section 7. This package also contains the correlation matrix used for the convergence run test. Figure 5.1 clearly illustrates the convergence properties of the various geometric programs. The efficiency of the algorithms is studied in Section 7 below.

6. A special case: Distance minimization. In this section, the primary concern of this paper to minimize the objective function of (1.2) is studied. The outline of this section is as follows. First, some particular choices for n and d are examined. Second, the differential and Hessian of F are calculated. Third, the connection with Lagrange multipliers is stated; in particular, this will lead to an identification method of whether a local minimum is a global minimum. Fourth, we discuss the PCA with re-scaling method for obtaining an initial feasible point.

6.1. The case of $d = n$. The case that C is a symmetric matrix and the closest positive semidefinite matrix X is to be found allows a successive projection solution, which was shown by Higham (2002).




FIG. 5.1. Convergence runs for various geometric programs: log relative residual $\ln(\|\text{grad}F(Y^{(i)})\|/\|\text{grad}F(Y^{(0)})\|)$ versus the iterate i .

6.2. The case of $d = 2$, $n = 3$. A 3×3 symmetric matrix with ones on the diagonals is denoted by

$$\begin{pmatrix} 1 & x & y \\ x & 1 & z \\ y & z & 1 \end{pmatrix}.$$

Its determinant is given by

$$\det = -\{x^2 + y^2 + z^2\} + 2xyz + 1.$$

By straightforward calculations it can be shown that $\det = 0$ implies that all eigenvalues are nonnegative. The set of 3 by 3 correlation matrices of rank 2 may thus be represented by the set $\{\det = 0\}$. To get an intuitive understanding of the complexity of the problem, the feasible region has been displayed in Figure 1.1.

6.3. Formula for the differential of F . Consider the specific case of the weighted Hadamard semi-norm of (1.2). This semi-norm can be represented by a Frobenius norm by introducing the Hadamard product \circ . The Hadamard product denotes entry-by-entry multiplication. Formally, for two matrices A and B of equal dimensions, the Hadamard product $A \circ B$ is defined by $(A \circ B)_{ij} = A_{ij}B_{ij}$. The objective function (1.2) can then be written as

$$F(Y) = \frac{1}{2} \sum_{i < j} W_{ij} (C_{ij} - Y_i Y_j^T)^2 = \frac{1}{2} \|W^{\circ 1/2} \circ \psi\|_F^2 = \frac{1}{2} \langle W^{\circ 1/2} \circ \psi, W^{\circ 1/2} \circ \psi \rangle,$$

with $\psi := YY^T - C$ and with $(W^{\circ 1/2})_{ij} = \sqrt{W_{ij}}$. Then

$$\begin{aligned} \frac{d}{dt} F(Y(t)) &= \langle W^{\circ 1/2} \circ \dot{\psi}, W^{\circ 1/2} \circ \psi \rangle = \langle \dot{\psi}, W \circ \psi \rangle \\ &= \langle \Delta Y^T + Y \Delta^T, W \circ \psi \rangle \\ &= \langle \Delta Y^T, W \circ \psi \rangle + \langle Y \Delta^T, W \circ \psi \rangle \\ &= \langle \Delta, 2(W \circ \psi)Y \rangle = \langle \Delta, F_Y \rangle, \quad \forall \Delta. \end{aligned}$$

Thus from (4.1) we have

$$F_Y = 2(W \circ \psi)Y. \tag{6.1}$$

Similarly, we may compute the second derivative

$$F_{YY}(\Delta) = \left. \frac{d}{dt} \right|_{t=0} F_Y(Y(t)) = 2 \left((W \circ \psi) \Delta + (W \circ (\Delta Y^T + Y \Delta^T)) Y \right),$$

with $Y(\cdot)$ any curve starting from Y in the direction of Δ .

6.4. Connection normal with Lagrange multipliers. The following lemma provides the basis for the connection of the normal vector at Y versus the Lagrange multipliers of the algorithm of Zhang & Wu (2003) and Wu (2003). The result is novel since previously only an expression was known for the matrix Y given the Lagrange multipliers. The result below establishes the reverse direction. This Lagrange result will allow us to identify whether a local minimum is also a global minimum. That we are able to efficiently determine whether a local minimum is a global minimum, is a very rare phenomenon in non-convex optimisation, and makes the rank reduction problem (non-convex for $d < n$) all the more interesting.

Note that the Lagrange theory is based on an efficient expression of the low-rank projection by an eigenvalue decomposition. Therefore the theory below can be extended efficiently only for the Hadamard norm with equal weights and for the weighted Frobenius norm, see also the discussion in Section 2.1. The proof of the following lemma has been deferred to Appendix F.

LEMMA 6.1. *Let $Y \in T_{n,d}$ be such that $\text{grad}F(Y) = 0$. Here, $\text{grad}F$ is the gradient of F on $T_{n,d}$, $\text{grad}F(Y) = \pi_{T_Y T_{n,d}}(F_Y) = F_Y - \text{diag}(F_Y Y^T) Y$, with F_Y in (6.1). Define*

$$\lambda := \frac{1}{2} \text{diag} (F_Y Y^T)$$

and define $C(\lambda) := C + \lambda$. Then there exist a joint eigenvalue decomposition

$$C(\lambda) = Q D Q^T, \quad Y Y^T = Q D^* Q^T$$

where D^* can be obtained by selecting at most d nonnegative entries from D (here if an entry is selected it retains the corresponding position in the matrix).

The characterization of the global minimum for Problem (1.1) was first achieved in Zhang & Wu (2003) and Wu (2003), which we repeat here: Denote by $\{X\}_d$ a matrix obtained by eigenvalue decomposition of X together with leaving in only the d largest eigenvalues (in absolute value). Denote for $\lambda \in \mathbb{R}^n$: $C(\lambda) = C + \text{diag}(\lambda)$. The proof of the following theorem has been repeated for clarity in Appendix G.

THEOREM 6.2. (Characterization of the global minimum of Problem (1.1), see Zhang & Wu (2003) and Wu (2003)) *Let C be a symmetric matrix. Let λ^* be such that there exists $\{C + \text{diag}(\lambda^*)\}_d \in C_{n,d}$ with*

$$\text{diag} (\{C + \text{diag}(\lambda^*)\}_d) = \text{diag}(C). \quad (6.2)$$

Then $\{C + \text{diag}(\lambda^*)\}_d$ is a global minimizer of Problem (1.1).

This brings us in a position to identify whether a local minimum is a global minimum:

THEOREM 6.3. *Let $Y \in T_{n,d}$ be such that $\text{grad}F(Y) = 0$ on $T_{n,d}$. Let λ and $C(\lambda)$ be defined as in Lemma 6.1. If $Y Y^T$ has the d largest eigenvalues from $C(\lambda)$ (in absolute value) then $Y Y^T$ is a global minimizer to the Problem (1.1).*

Proof: Apply Lemma 6.1 and Theorem 6.2. □

6.5. Initial feasible point. To obtain an initial feasible point $Y \in T_{n,d}$ we use a method of Flury (1988). We first perform an eigenvalue decomposition

$$C = Q \Lambda Q^T, \quad \Lambda_{11} \geq \dots \geq \Lambda_{nn}. \quad (6.3)$$

Then we define Y by assigning to each row

$$Y_i = \frac{Z}{\|Z\|_2}, \quad Z = (Q_d \Lambda_{d+}^{1/2})(i, :), \quad i = 1, \dots, n,$$

where Q_d consists of the first d columns of Q and where Λ_{d+} is the principal sub-matrix of Λ of degree d , filled only with the non-negative elements from Λ . The scaling is to ensure that each row of Y is of

TABLE 7.1
Excerpt of Table 3 in De Jong et al. (2004).

	γ_1	γ_2	γ_3	γ_4
estimate	0.000	0.480	1.511	0.186
standard error	-	0.099	0.289	0.127

unit length. If row i is a priori of zero length, then we choose Y_i to be an arbitrary vector in \mathbb{R}^d of unit length. Finally, to obtain an initial feasible point in $\text{Chol}_{n,d}$, we perform a Cholesky decomposition as in the proof of Theorem 3.6.

Note that the condition of decreasing norm in (6.3) is thus key to ensure that the initial point is close to the global minimum, see the result of Theorem 6.3.

7. Numerical results. There are many different algorithms available in the literature, as detailed in Section 2. Some of these have an efficient implementation, i.e., the cost of a single iteration is low. Some algorithms have fast convergence, for example, the Newton method has quadratic convergence. Algorithms with fast convergence usually require less iterations to attain a predefined convergence criterion. Thus, the real-world performance of an algorithm is a trade-off between cost-per-iterate and number of iterations required. A priori, it is not clear which algorithm will perform best. Therefore, in this section, the numerical performance of geometric optimisation is compared to other methods available in the literature.

7.1. Acknowledgement. Our implementation of geometric optimisation over low-rank correlation matrices ‘LRCM min’⁵ is an adoption of the ‘SG min’ template of Edelman & Lippert (2000) (written in MATLAB) for optimisation over the Stiefel and Grassmann manifolds. This template contains four distinct well-known non-linear optimisation algorithms adapted for geometric optimisation over Riemannian manifolds: Newton algorithm; dogleg step or Levenberg (1944) and Marquardt (1963) algorithm; Polak & Ribière (1969) conjugate gradient; and Fletcher & Reeves (1964) conjugate gradient.

7.2. Numerical comparison. The performances of the following seven algorithms, all of these described in Sections 4 and 2, except for item 7 (`fmincon`), are compared:

1. Geometric optimisation, Newton (Newton).
2. Geometric optimisation, Fletcher-Reeves conjugate gradient (FRCG).
3. Majorization, e.g., Pietersz & Groenen (2004b) (Major.).
4. Parametrization, e.g., Rebonato (1999b) (Param.).
5. Alternating projections without normal vector correction, e.g., Grubišić (2002) (Alt. Proj.).
6. Lagrange multipliers, e.g., Zhang & Wu (2003) (Lagrange).
7. `fmincon`, a MATLAB built-in medium-scale constrained nonlinear program (`fmincon`).

Note that the first two algorithms in this list have been developed in this paper. The algorithms are tested on a large number of randomly generated correlation matrices. The benefit of testing on many correlation matrices is, that the overall and generic performance of the algorithms may be assessed. The correlation matrices are randomly generated as follows. A parametric form for (primarily interest rate) correlation matrices is posed in De Jong, Driessen & Pelsser (2004, Equation (8)). We repeat the parametric form here for completeness.

$$\rho(T_i, T_j) = \exp \left\{ -\gamma_1 |T_i - T_j| - \frac{\gamma_2 |T_i - T_j|}{\max(T_i, T_j)^{\gamma_3}} - \gamma_4 |\sqrt{T_i} - \sqrt{T_j}| \right\},$$

with $\gamma_1, \gamma_2, \gamma_4 > 0$ and with T_i denoting the expiry time of rate i . (Our particular choice is $T_i = i$, $i = 1, 2, \dots$) This model was then subsequently estimated with USD historical interest rate data. In Table 3 of De Jong et al. (2004) the estimated γ -parameters are listed, along with their standard error. An excerpt of this table has been displayed in Table 7.1. The random correlation matrix that we use is obtained by randomizing the γ -parameters. We assume the γ -parameters distributed normally with mean and standard errors given by Table 7.1, with $\gamma_1, \gamma_2, \gamma_4$ capped at zero.

As the benchmark criterion for the performance of an algorithm, we take its obtained accuracy of fit given a fixed amount of computational time. Such a criterion corresponds to financial practice, since decisions based on derivative valuation calculations often need to be made within seconds. To display the

⁵LRCM min can be downloaded from www.few.eur.nl/few/people/pietersz.




FIG. 7.1. Performance profile with $n = 30$, $d = 3$, 2 seconds of computational time, Hadamard norm with equal weights. A rule of thumb is, that the higher the graph of an algorithm, the better its performance.

comparison results, we use the state-of-the-art and convenient *performance profiles*; see Dolan & Moré (2002). The reader is referred there for details, but the idea is briefly described here. There are 100 test correlation matrices $p = 1, \dots, 100$, and seven algorithms $s = 1, \dots, 7$. As *performance measure* we take the obtained function value $F^{(p,s)}$ of algorithm s on problem p given the limited computational time. The *performance ratio* $\rho^{(p,s)}$ is defined to be the ratio of the performance measure of the algorithm over the best obtained performance measure for all seven algorithms,

$$\rho^{(p,s)} = \frac{F^{(p,s)}}{\min\{F^{(p,s)} : s = 1, \dots, 7\}}.$$

The cumulative distribution function $\Omega^{(s)}$ of the performance ratio for algorithm s , viewed as a random variable $p \rightarrow \rho^{(p,s)}$, is then the *performance profile* of algorithm s ,

$$\Omega^{(s)}(\xi) = \frac{1}{100} \#\{p : \rho^{(p,s)} \leq \xi, p = 1, \dots, 100\}.$$

A rule of thumb is, that the higher the profile of an algorithm, the better its performance. The performance profiles have been displayed in Figures 7.1–7.4, for various choices of n , d , and computational times. Each performance profile represents a benchmark on 100 different test interest rate correlation matrices. For Figures 7.1–7.3, an objective function with equal weights is used. For Figure 7.4, we use a Hadamard semi-norm with non-constant weights. These weights are chosen so as to reflect the importance of the correlation entries for a specific trigger swap, as outlined in, e.g., Rebonato (2004b, Section 20.4.3). For this specific trigger swap, the first three rows and columns are important. Therefore the weights matrix W takes the form

$$W_{ij} = \begin{cases} 1 & \text{if } i \leq 3 \text{ or } j \leq 3, \\ 0 & \text{otherwise.} \end{cases}$$

From Figures 7.1–7.4 it becomes clear that geometric optimisation compares favourably to the other methods available in the literature, with respect to obtaining the best fit to the original correlation matrix within a limited computational time.




FIG. 7.2. Performance profile with $n = 50$, $d = 4$, 1 second of computational time, Hadamard norm with equal weights.

8. Conclusions. We applied geometric optimisation tools for finding the nearest low-rank correlation matrix. The differential geometric machinery provided us with an algorithm more efficient than any existing algorithm in the literature, at least for the numerical cases considered. The geometric approach also allows for insight and more intuition into the problem. We established a technique that allows one to straightforwardly identify whether a local minimum is a global minimum.

REFERENCES

- Abraham, R., Marsden, J. E. & Ratiu, T. (1988), *Manifolds, Tensor Analysis, and Applications*, Springer-Verlag, Berlin.
- Al-Baali, M. (1985), ‘Descent property and global convergence of the Fletcher-Reeves method with inexact line search’, *IMA Journal of Numerical Analysis* **5**, 121–124.
- Brace, A., Gątarek, D. & Musiela, M. (1997), ‘The market model of interest rate dynamics’, *Mathematical Finance* **7**(2), 127–155.
- Brigo, D. (2002), A note on correlation and rank reduction, www.damianobrigo.it.
- Brigo, D. & Mercurio, F. (2001), *Interest Rate Models: Theory and Practice*, Springer-Verlag, Berlin.
- do Carmo, M. P. (1992), *Riemannian Geometry.*, 12 edn, Birkhäuser, Boston, MA.
- Dai, Q. & Singleton, K. (2003), ‘Term structure dynamics in theory and reality’, *Review of Financial Studies* **16**(3), 631–678.
- De Jong, F., Driessen, J. & Pelsser, A. A. J. (2004), ‘On the information in the interest rate term structure and option prices’, *Review of Derivatives Research* **7**(2), 99–127.
- Dedieu, J.-P., Priouret, P. & Malajovich, G. (2003), ‘Newton’s method on Riemannian manifolds: Covariant alpha-theory’, *IMA Journal of Numerical Analysis* **23**(3), 395–419.
- Depczynski, U. & Stöckler, J. (1998), A differential geometric approach to equidistributed knots on Riemannian manifolds, in C. K. Chui & L. L. Schumaker, eds, ‘Approximation Theory IX, Theoretical Aspects’, Vol. 1, Vanderbilt University Press, Nashville, TN, pp. 99–106.
- Dolan, E. D. & Moré, J. J. (2002), ‘Benchmarking optimization software with performance profiles’, *Mathematical Programming, Series A* **91**(2), 201–213.
- Duistermaat, J. J. & Kolk, J. A. C. (2000), *Lie Groups*, Springer-Verlag, Berlin.
- Dykstra, R. L. (1983), ‘An algorithm for restricted least squares regression’, *Journal of the American Statistical Association* **87**(384), 837–842.
- Edelman, A. & Lippert, R. (2000), Nonlinear eigenvalue problems with orthogonality constraints (section 8.3), in Z. Bai, J. Demmel, J. Dongarra, A. Ruhe & H. van der Vorst, eds, ‘Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide’, SIAM, Philadelphia, NJ.
- Edelman, A., Arias, T. A. & Smith, S. T. (1998), ‘The geometry of algorithms with orthogonality constraints’, *SIAM Journal*




FIG. 7.3. Performance profile with $n = 60$, $d = 5$, 3 seconds of computational time, Hadamard norm with equal weights.

of *Matrix Analysis and its Applications* **20**(2), 303–353.

- Fletcher, R. & Reeves, C. M. (1964), ‘Function minimization by conjugate gradients’, *Computer Journal* **7**(2), 149–154.
- Flury, B. (1988), *Common Principal Components and Related Multivariate Models*, J. Wiley & Sons, New York.
- Gilbert, J.-C. & Nocedal, J. (1992), ‘Global convergence properties of conjugate gradient methods for optimization’, *SIAM Journal on Optimization* **2**(1), 21–42.
- Golub, G. H. & van Loan, C. F. (1996), *Matrix Computations*, 3 edn, John Hopkins University Press, Baltimore, MD.
- Grubišić, I. (2002), Interest rate theory: BGM model, Master’s thesis, Leiden University. www.math.uu.nl/people/grubisic.
- Han, S.-P. (1988), ‘A successive projection method’, *Mathematical Programming* **40**, 1–14.
- Higham, N. J. (2002), ‘Computing the nearest correlation matrix—a problem from finance’, *IMA Journal of Numerical Analysis* **22**(3), 329–343.
- Horn, R. A. & Johnson, C. R. (1990), *Matrix Analysis*, Cambridge University Press, Cambridge.
- Hull, J. C. & White, A. (2000), ‘Forward rate volatilities, swap rate volatilities, and implementation of the LIBOR market model’, *Journal of Fixed Income* **10**(2), 46–62.
- Jamshidian, F. (1997), ‘LIBOR and swap market models and measures’, *Finance and Stochastics* **1**(4), 293–330.
- Joshi, M. S. (2003), ‘Rapid computation of drifts in a reduced factor LIBOR market model’, *Wilmott Magazine* **5**, 84–85.
- Levenberg, K. (1944), ‘A method for the solution of certain non-linear problems in least squares’, *Quarterly of Applied Mathematics* **2**, 164–168.
- Marquardt, D. W. (1963), ‘An algorithm for least-squares estimation of nonlinear parameters’, *Journal of the Society for Industrial and Applied Mathematics* **11**(2), 431–441.
- Miltersen, K. R., Sandmann, K. & Sondermann, D. (1997), ‘Closed form solutions for term structure derivatives with log-normal interest rates’, *Journal of Finance* **52**(1), 409–430.
- Morini, M. & Webber, N. (2004), An EZI method to reduce the rank of a correlation matrix, www.cass.city.ac.uk/facfin/facultypages/nwebber/.
- Munkres, J. R. (1975), *Topology*, Prentice-Hall, London.
- Musiela, M. & Rutkowski, M. (1997), ‘Continuous-time term structure models: Forward measure approach’, *Finance and Stochastics* **1**(4), 261–291.
- Pietersz, R. & Groenen, P. J. F. (2004a), ‘A major LIBOR fit’, *Risk Magazine*. December issue.
- Pietersz, R. & Groenen, P. J. F. (2004b), ‘Rank reduction of correlation matrices by majorization’, *Quantitative Finance*. Forthcoming.
- Polak, E. & Ribière, G. (1969), ‘Note sur la convergence de méthodes de directions conjuguées’, *Revue Française d’Informatique et de Recherche Opérationnelle* **16**, 35–43.
- Rapisarda, F., Brigo, D. & Mercurio, F. (2002), Parametrizing correlations: A geometric interpretation, Banca IMI Working Paper, www.fabiomercurio.it.
- Rebonato, R. (1999a), ‘Calibrating the BGM model’, *Risk Magazine*, pp. 74–79. March.
- Rebonato, R. (1999b), ‘On the simultaneous calibration of multifactor lognormal interest rate models to Black volatilities




FIG. 7.4. Performance profile with $n = 15$, $d = 3$, 1 second of computational time, trigger swap Hadamard semi-norm.

and to the correlation matrix', *Journal of Computational Finance* **2**(4), 5–27.

Rebonato, R. (1999c), *Volatility and Correlation in the Pricing of Equity, FX and Interest-Rate Options*, J. Wiley & Sons, Chichester.

Rebonato, R. (2002), *Modern Pricing of Interest-Rate Derivatives*, Princeton University Press, New Jersey.

Rebonato, R. (2004a), 'Interest-rate term-structure pricing models: a review', *Proceedings of the Royal Society London* **460**, 667–728. Series A.

Rebonato, R. (2004b), *Volatility and Correlation: The Perfect Hedger and the Fox*, 2 edn, J. Wiley & Sons, Chichester, UK.

Sidenius, J. (2000), 'LIBOR market models in practice', *Journal of Computational Finance* **3**(3), 5–26.

Smith, S. T. (1993), *Geometric Optimization Methods for Adaptive Filtering*, PhD thesis, Harvard University, Cambridge, MA.

Weigel, P. (2004), 'Optimal calibration of LIBOR market models to correlations', *Journal of Derivatives* **12**(3).

Wu, L. (2003), 'Fast at-the-money calibration of the LIBOR market model using Lagrange multipliers', *Journal of Computational Finance* **6**(2), 39–77.

Zhang, Z. & Wu, L. (2003), 'Optimal low-rank approximation to a correlation matrix', *Linear Algebra and its Applications* **364**, 161–187.

Zoutendijk, G. (1970), Nonlinear programming, computational methods, in J. Abadie, ed., 'Integer and nonlinear programming', pp. 37–86.

Appendix A. The proof of Theorem 3.3.

Proof of (i). The maps Ψ and Φ are well defined: To show that Ψ is well defined, we need to show that if $Y_2 \in [Y_1]$, then $Y_2 Y_2^T = Y_1 Y_1^T$. From the assumption, we have that $\exists Q \in O_d : Y_2 = Y_1 Q$. It follows that

$$Y_2 Y_2^T = (Y_1 Q)(Y_1 Q)^T = Y_1 Q Q^T Y_1^T = Y_1 Y_1^T,$$

which was to be shown.

To show that Φ is well defined, we need to show:

(A) If $X \in C_{n,d}$ then there exists $Y \in T_{n,d}$ such that $X = Y Y^T$.

(B) If $Y, Z \in T_{n,d}$, with $Y Y^T = Z Z^T =: X$ then there exists $Q \in O_d$ such that $Y = Z Q$.

Ad (A): Let

$$X = Q \Lambda Q^T, \quad Q \in O_n, \quad \Lambda = \text{diag}(\Lambda),$$

be an eigenvalue decomposition with $\Lambda_{ii} = 0$ for $i = d + 1, \dots, n$. Note that such a decomposition of the specified form is possible because of the restriction $X \in C_{n,d}$. Then note that

$$Q \sqrt{\Lambda} = ((Q \sqrt{\Lambda})(:, 1:d) \mid 0).$$

Thus if we set $Y = (Q \sqrt{\Lambda})(:, 1:d)$ then $Y Y^T = X$ and $Y \in T_{n,d}$, which was to be shown.

Ad (B): Let $\text{rank}(Y) = \text{rank}(Z) = \text{rank}(X) = k \leq d$. Without loss of generality, we may assume that the first k rows of Y and Z are independent. We extend the set of k row vectors $\{Y_1, \dots, Y_k\}$ to a set of d row vectors $\{Y_1, \dots, Y_k, \tilde{Y}_{k+1}, \dots, \tilde{Y}_d\}$, such that the latter forms a basis of \mathbb{R}^d . Similarly, we obtain a basis $\{Z_1, \dots, Z_k, \tilde{Z}_{k+1}, \dots, \tilde{Z}_d\}$ of \mathbb{R}^d . It follows that there exists an orthogonal rotation Q , $Q Q^T = I_d$, such that $Q Y_i = Z_i$ ($i = 1, \dots, k$), $Q \tilde{Y}_i = \tilde{Z}_i$ ($i = k + 1, \dots, d$). Note that then also $Q Y_i = Z_i$ for $i = k + 1, \dots, n$, by linearity of Q and since the last $n - k$ row vectors are linearly dependent on the first k row vectors by assumption. It follows that $Y Q = Z$, which was to be shown. \square

Proof (ii). Diagram (3.3) is commutative: To show that $\Psi \circ \pi = s$: Let $Y \in T_{n,d}$, then $\pi(Y) = [Y]$ and $\Psi([Y]) = Y Y^T$ and also $s(Y) = Y Y^T$. To show that $\Phi \circ s = \pi$: Let $Y \in T_{n,d}$, then $s(Y) = Y Y^T$ and $\Phi(Y Y^T) = [Y]$ and also $\pi(Y) = [Y]$. \square

Proof of (iii). The map Ψ is a homeomorphism with inverse Φ : It is straightforward to verify that $\Phi \circ \Psi$ and $\Psi \circ \Phi$ are both the identity maps. The map Ψ is thus bijective with inverse Φ . To show that Ψ is continuous, note that for quotient spaces we have: The map Ψ is continuous if and only if $\Psi \circ \pi$ is continuous (see for example Abraham, Marsden & Ratiu (1988), Proposition 1.4.8). In our case, $\Psi \circ \pi = s$ with $s(Y) = Y Y^T$ is continuous. The proof now follows from a well-known lemma from topology: A continuous bijection from a compact space into a Hausdorff space is a homeomorphism (see for example Munkres (1975), Theorem 5.6). \square

Appendix B. Proof of Proposition 3.4.

1. It is sufficient to show that $\{Y \in \mathbb{R}^{n \times d} : \text{rank}(Y) = d\}$ is open in $\mathbb{R}^{n \times d}$, since $T_{n,d}^*$ is open in $T_{n,d}$ if and only if $\{Y \in \mathbb{R}^{n \times d} : \text{rank}(Y) = d\}$ is open in $\mathbb{R}^{n \times d}$. Since the rank of a symmetric matrix is a locally constant function, it follows that $\{Y \in \mathbb{R}^{n \times d} : \text{rank}(Y) = d\} = s^{-1}(\text{rank}^{-1}(d))$ is an open subset of $\mathbb{R}^{n \times d}$, with $s(Y) = Y Y^T$ as in Definition 3.2. \square
2. This part is a corollary of Theorem 1.11.4 of Duistermaat & Kolk (2000). This theorem essentially states that for a smooth action of a Lie group on a manifold the quotient is a manifold if the action is proper and free. First, we show that the action of O_d on $T_{n,d}^*$ is proper⁶. Let

$$\phi : T_{n,d}^* \times O_d \rightarrow T_{n,d}^* \times T_{n,d}^*, \quad (Y, Q) \mapsto (Y Q, Y)$$

and K a compact subset of $T_{n,d}^* \times T_{n,d}^*$. We have to show that $\phi^{-1}(K)$ is compact. By continuity of ϕ , $\phi^{-1}(K)$ is closed in $T_{n,d}^* \times O_d$. Because $T_{n,d}^* \times O_d$ is bounded it follows that $\phi^{-1}(K)$ is compact.

Second, we show that the O_d -action on $T_{n,d}^*$ is free. Let $Y \in T_{n,d}^*$ and $Q \in O_d$ such that $Y Q = Y$. Since $\text{rank}(Y) = d$, it follows from the proof of Theorem 3.3 (i) that there exists precisely one $Q \in O_d$ such that $Y Q = Y$. Thus, this Q must be the identity matrix.

The dimension of $M_{n,d}^* = \dim(T_{n,d}^*) - \dim(O_d) = n(d - 1) - \frac{1}{2}d(d - 1)$. \square

⁶For a definition, see Duistermaat & Kolk (2000, page 53).

Appendix C. Proof of Proposition 3.5.

This part is a corollary of Theorem 1.11.4 of Duistermaat & Kolk (2000). This theorem states that there is only one differentiable structure on the orbit space which satisfies the following: Suppose that, for every $[Y] \in M_{n,d}^*$, we have an open neighbourhood $\mathcal{U} \subseteq M_{n,d}^*$ and a bijective map:

$$\tau : \pi^{-1}(\mathcal{U}) \rightarrow \mathcal{U} \times O_d, \quad Y \mapsto (\pi(Y), \chi(Y)),$$

such that, for every $Y \in \pi^{-1}(\mathcal{U})$, $Q \in O_d$, $\tau(YQ) = (\pi(Y), \chi(Y)Q)$. The differentiable structure on $M_{n,d}^*$ is the one which makes τ into a diffeomorphism. The topology of $M_{n,d}^*$ obtained in this manner is equal to the quotient topology.

Let $Y \in T_{n,d}^*$ and σ_Y be a section over \mathcal{U}_Y defined in (3.4). We define $\tau_Y : \pi^{-1}(\mathcal{U}_Y) \rightarrow \mathcal{U}_Y \times O_d$ as follows. For $Z \in \pi^{-1}([Z])$, $[Z] \in \mathcal{U}_Y$, there is a unique element $Q_Z \in O_d$ such that $Z = \sigma_Y([Z])Q_Z$. Then we define τ_Y by $\tau_Y(Z) = ([Z], Q_Z)$. By definition, we have that $\tau_Y^{-1}([Z], Q) = \sigma_Y([Z])Q$. Since $\tau_Y^{-1} : \pi^{-1}(\mathcal{U}_Y) \rightarrow \mathcal{U}_Y \times O_d$ is a bijective map, we have that τ_Y is bijective, too. It can be easily verified that τ_Y satisfies the condition $\tau_Y(YQ) = ([Y], Q_Y Q)$ of Theorem 1.11.4 of Duistermaat & Kolk (2000) stated above. It follows that τ_Y is a diffeomorphism if and only if $\sigma_Y : \mathcal{U} \rightarrow \sigma_Y(\mathcal{U})$ is a diffeomorphism. Thus, the differentiable structure on $M_{n,d}^*$ is the one which makes $\sigma_Y : \mathcal{U}_Y \rightarrow \sigma(\mathcal{U}_Y)$ into a diffeomorphism. \square

Appendix D. Proof of Theorem 3.6.

Let $X \in C_{n,d}$ and suppose that $\text{rank}(X) = k \leq d$. Then there is a $Y \in T_{n,k}$ such that $YY^T = X$, by Theorem 3.3. Apply to Y the procedure⁷ outlined in Section 3.3, to obtain a lower-triangular matrix $Y^* \in T_{n,k}$, such that $Y^*(Y^*)^T = X$. A lower-triangular matrix $\bar{Y} \in \text{Chol}_{n,d}$ that satisfies $\bar{Y}\bar{Y}^T = X$ can now easily be obtained by setting

$$\bar{Y} = \begin{pmatrix} Y^* & \underbrace{0}_{n \times (d-k)} \end{pmatrix},$$

which was to be shown. \square

Appendix E. Proof of Theorem 3.7.

First, we prove the ‘only if’ part. Note that it is sufficient to show that the map $s : \text{Chol}_{n,d} \rightarrow C_{n,d}$ is *open*. For then if Y attains a local minimum of F on the open neighbourhood $\mathcal{U} \subset \text{Chol}_{n,d}$, then $s(Y) = YY^T$ attains a local minimum of \tilde{F} on the open neighbourhood $s(\mathcal{U})$ of YY^T , since for any $X' = Y'Y'^T \in s(\mathcal{U})$, $\tilde{F}(X') = \tilde{F}(Y'Y'^T) = F(Y') \geq F(Y) = \tilde{F}(YY^T)$.

To show that $s : \text{Chol}_{n,d} \rightarrow C_{n,d}$ is open, note that it is sufficient to show that $\pi : \text{Chol}_{n,d} \rightarrow M_{n,d}$ is open, since $\Psi : M_{n,d} \rightarrow C_{n,d}$ is a homeomorphism (see Proposition 3.4, item 3) and $s = \Psi \circ \pi$.

Suppose, then, that \mathcal{U} is open in $\text{Chol}_{n,d}$. We have to show that $\pi^{-1}(\pi(\mathcal{U}))$ is open in $T_{n,d}$, by definition of the quotient topology of $M_{n,d}$. We have

$$\pi^{-1}(\pi(\mathcal{U})) = \{ YQ : Y \in \mathcal{U}, Q \in O_d \}.$$

It is sufficient to show that the complement $(\pi^{-1}(\pi(\mathcal{U})))^c$ is closed. Let $\{Y^{(i)}\}$ be a sequence in $(\pi^{-1}(\pi(\mathcal{U})))^c$ converging to Y , i.e. $\lim_{i \rightarrow \infty} \|Y^{(i)} - Y\| = 0$. We can write $Y^{(i)} = Z^{(i)}Q^{(i)}$ with $Z^{(i)} \in \mathcal{U}^c$ and $Q^{(i)} \in O_d$. Then,

$$\lim_{i \rightarrow \infty} \|Y^{(i)} - Y\| = \lim_{i \rightarrow \infty} \|Z^{(i)}Q^{(i)} - Y\| = \lim_{i \rightarrow \infty} \|Z^{(i)} - Y(Q^{(i)})^T\| = 0. \quad (\text{E.1})$$

Since $\mathcal{U}^c \times O_d$ is compact, there exists a convergent subsequence $\{Z^{(i_j)}, Q^{(i_j)}\}$, with $Z^{(i_j)} \rightarrow Z^* \in \mathcal{U}^c$ and $Q^{(i_j)} \rightarrow Q^*$, say. From (E.1) it follows that $Z^* = Y(Q^*)^T \in \mathcal{U}^c$, which implies $Y \in (\pi^{-1}(\pi(\mathcal{U})))^c$.

The reverse direction is obvious since the map $s : \text{Chol}_{n,d} \rightarrow C_{n,d}$ is continuous. \square

Appendix F. Proof of Lemma 6.1.

It is recalled from matrix analysis that X_1 and X_2 admit a joint eigenvalue decomposition if and only if their Lie bracket $[X_1, X_2] = X_1X_2 - X_2X_1$ equals zero. Define $\bar{C}(\lambda) := -\psi + \lambda$. Note that $2\lambda Y$ is the projection $\pi_{N_Y T_{n,d}}(F_Y)$ of F_Y onto the normal space at Y . Note also that

$$YY^T + \bar{C}(\lambda) = C(\lambda). \quad (\text{F.1})$$

⁷The procedure in Section 3.3 is stated in terms of d , but k should be read there in this case.

We calculate

$$\bar{C}(\lambda)Y = \{-\psi + \lambda\}Y = -\frac{1}{2}F_Y + \frac{1}{2}\pi_{N_Y T_{n,d}}(F_Y) = 0. \quad (\text{F.2})$$

The last equality follows from the assumption that $\text{grad}F(Y) = 0$, i.e. the differential F_Y is normal at Y . (Here, $\text{grad}F(Y)$ denotes the gradient on $T_{n,d}$.) It follows from (F.2) and from the symmetry of $\bar{C}(\lambda)$ that

- (i) $YY^T\bar{C}(\lambda) = 0$ and also,
- (ii) $[YY^T, \bar{C}(\lambda)] = 0$.

From (ii), YY^T and $\bar{C}(\lambda)$ admit a joint eigenvalue decomposition, but then also jointly with $C(\lambda)$ because of (F.1). Suppose $\bar{C}(\lambda) = Q\bar{D}Q^T$. From (i) we then have that D_{ii}^* and \bar{D}_{ii} cannot both be non-zero. The result now follows since YY^T is positive semidefinite and has rank less than or equal to d . \square

Appendix G. Proof of Theorem 6.2.

Define the Lagrangian

$$\mathcal{L}(X, \lambda) := -\|C - X\|_F^2 - 2\lambda^T \text{diag}(C - X), \quad \text{and}$$

$$V(\lambda) := \min \{ \mathcal{L}(X, \lambda) : \text{rank}(X) = d \}. \quad (\text{G.1})$$

Note that the minimization problem in Equation (G.1) is attained by any $\{C(\lambda)\}_d$ (see e.g., Equation (30) of Wu (2003)). For any $X \in C_{n,d}$,

$$\|C - X\|_F^2 \stackrel{(a)}{=} -\mathcal{L}(X, \lambda^*) \stackrel{(b)}{\geq} -V(\lambda^*) \stackrel{(c)}{=} \|C - \{C(\lambda)\}_d\|_F^2.$$

(This is the equation at the end of the proof of Theorem 4.4 of Zhang & Wu (2003).) Here (in-)equality

- (a) is obtained from the property that $X \in C_{n,d}$,
- (b) is by definition of V , and
- (c) is by assumption of (6.2). \square