Efficient Rank Reduction of Correlation Matrices[1]

Igor Grubišić[2], Raoul Pietersz[3]

First version: 29 August 2002, this version: 25 November 2003

Abstract. Geometric optimization algorithms are developed that efficiently find the nearest low-rank correlation matrix. We show numerically that our methods outperform the existing methods in the literature. The algorithm is shown theoretically to be globally convergent to a local minimum, with a quadratic local rate of convergence. The connection with the Lagrange multiplier method is established, along with an identification of wether a local minimum is a global minimum. The additional benefits of the geometric approach are: (i) any weighted norm can be applied whereas previous methods only allowed for a weighted Frobenius norm, and (ii) neighborhood search can straightforwardly be applied.

Key words: geometric optimization, correlation matrix, structured rank reduction, Thomson problem

JEL Classification: G13

[1]We are grateful for the comments of Eric Balder, Pieter Eendebak, Andrea Gaccobe, Raymond Lacey, Antoon Pelsser, Jordan Rizov and seminar participants at Leiden University and Utrecht University. All errors are our own.

[2]Mathematical Institute, University of Utrecht, P.O. Box 80010, 3508 TA Utrecht, The Netherlands (email: grubisic@math.uu.nl, tel: +31 30 2531529, fax: +31 30 2518394)

[3]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, fax: +31 10 4089136) and Market Risk – Modelling and Product Analysis (HQ1059), ABN AMRO Bank, P.O. Box 283, 1000 EA Amsterdam, The Netherlands

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:

$$
\begin{array}{rl}
\textbf{Find} & X \in \mathbb{S}_n \\
\textbf{to minimize} & \frac{1}{2}\|C - X\|^2 \\
\textbf{subject to} & \operatorname{rank}(X) \le d \\
 & X_{ii} = 1,\ i = 1 : n \\
 & X \succeq 0
\end{array}
\tag{1}
$$

Or in other words: Find the low-rank correlation matrix X nearest to the given $n \times n$ matrix C. The choice of the norm will reflect what is meant by nearness of the two matrices. In our setting C has the interpretation of measured correlation. Due to the measurement error C will in general not satisfy the constraints. It can however be the case that some specific entries of the matrix C are better measured than others[4]. This dictates that we have to put weights on individual entries of matrices. This consideration provides the natural norm to our problem. In the literature this is a well known norm called the Hadamard norm and it is denoted by $\|\cdot\|_H$.

The importance of this problem in finance has been recognized by several researchers such as [Reb99], [Hig02] and [ZhW03]. One of the first articles to address this problem in finance is [Reb99], in which the set of rank d correlation matrices is parameterized. The minimization technique was however not addressed. Instead only the principal components analysis (PCA) method was proposed to obtain a feasible point. This simple method is due to Flury [Flu88] and is the main technique used by practitioners today. Zhang and Wu recognized the shortcomings of PCA and in their recent work [ZhW03] they established the most significant result up to now in this area. They extended the PCA technique (call it the ZW algorithm) such that if C obeys some additional conditions and if ZW converges then it produces the globally nearest point. This method is however still often plagued by non-convergence as shown numerically in this paper. Moreover the Zhang-Wu algorithm is only applicable for the weighted Frobenius norm.

The other important contribution is due to Higham [Hig02]. In this paper Higham develops an algorithm that solves the problem globally if $d = n$

[4] Typically this happens in a finance setting.

without any additional conditions on C. This method is also applicable only for the weighted Frobenius norm.

We propose a novel technique that can solve the problem locally without any additional assumptions and which incorporates the Hadamard norm. It will be shown that our method numerically outperforms the work of Zhang & Wu and Higham. 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.

In our paper we formulate the problem in terms of Riemannian geometry. This approach allows us to use gradient methods which are numerically stable and efficient, in particular the Riemann-Newton method is applied.

The paper is organized as follows. In section 2, the constraints of the problem are formulated in terms of differential geometry. We identify the set of correlation matrices of rank at most d with a set of equivalence classes of n products of the $d-1$-sphere. This is the canonical space for the optimization of the arbitrary smooth function subject to the same constraints. Section 3 describes the Riemannian structure on n products of the $d-1$-sphere and the quotient space. Formulas are given for parallel transport, geodesics and the horizontal space. The minimization algorithms are made explicit. In section 4 we study theoretically the convergence of the algorithms. We establish global convergence and the local convergence rates. The application of the algorithms to the problem of finding the nearest low-rank correlation matrix is worked out in detail in section 5. In section 6 we present numerical results with randomly generated correlation matrices. These are compared with the algorithm of Zhang and Wu. The application to Thomson's problem is given in section 7. Finally in section 8 we conclude the paper.

2 Solution methodology

The problem with general F. Note that problem (1) is a special case of the following more general problem:

(2)

Find	$X \in \mathbb{S}_n$
to minimize	$F(X)$
subject to	$\text{rank}(X) \leq d$
	$X_{ii} = 1,\ i = 1 : n$
	$X \succeq 0$

Here we have used the following notation:

- $\mathbb{S}_n$ is the set of n by n real symmetric matrices.

- $F : \mathbb{S}_n \to \mathbb{R}$ is the objective function.

- $\succeq$ denotes positive semidefiniteness.

In this paper methods will be developed to solve problem (2) for the case when F is smooth. Note that such methods will thus also solve problem (1), which is the primary goal of this paper.

Basic idea. The idea for solving problem (2) is to equip the constraint set with a differentiable structure and subsequently utilize the recent advances of geometric optimization over manifolds. These advances include [EAS99], [Smi93] and [DPM03]. The advantage of these methods is that the coordinates are chosen judiciously, generally with a higher number of parameters than necessary, however it is this choice that leads to the simplest form for the gradient and geodesics, which in turn leads to an efficient implementation. We identify the constraint set with 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 $\mathbb{R}^d$ with the corresponding inner products. This representation is independent of the choice of basis of $\mathbb{R}^d$. In other words, any orthogonal rotation of the configuration does not alter the associated correlation matrix. This idea is developed more rigorously below.

A homeomorphism. In this section the set of $n \times n$ correlation matrices of rank d is equipped with a topology. We subsequently establish a homeomorphism between the latter topological space with the quotient space of n products of the $d-1$ sphere over the group of orthogonal transformations of $\mathbb{R}^d$.

Definition 1 *The set of rank d symmetric $n \times n$ correlation matrices is defined by*

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

Here I_n denotes the $n \times n$ identity matrix and diag *denotes the map $\mathbb{R}^{n \times n} \to \mathbb{R}^{n \times n}$*

$$\operatorname{diag}(X) = \begin{pmatrix} X_{11} & & 0 \\ & \ddots & \\ 0 & & X_{nn} \end{pmatrix} \text{ for } X \in \mathbb{R}^{n \times n}.$$

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, elements of n products of the $d-1$ sphere are denoted as a matrix $Y \in \mathbb{R}^{n\times d}$, with each row vector Y_i of unit length. Elements of the orthogonal group acting on $\mathbb{R}^d$ are denoted by a $d\times d$ orthogonal matrix Q.

Definition 2 *The product of n spheres S^{d-1} is denoted by $T_{n,d}$. Denote by O_d the group of orthogonal transformation of d-space. We define the following right O_d-action*[5] *on $T_{n,d}$:*

$$\begin{aligned} T_{n,d} \times O_d &\to T_{n,d} \\ (Y, Q) &\mapsto YQ. \end{aligned} \tag{3}$$

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 or space of orbits $T_{n,d}/O_d$ is denoted by $M_{n,d}$. The canonical projection $T_{n,d} \to T_{n,d}/O_d = M_{n,d}$ is denoted by π. Define the map[6]

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

$$\Psi\big(\,[Y]\,\big) = YY^T.$$

Consider a map[7] *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 shown in theorem 3. Then set

$$\Phi(X) = [Y].$$

Finally, define the map $s : T_{n,d} \to C_{n,d}$,

$$s(Y) = YY^T.$$

The following theorem relates the spaces $C_{n,d}$ and $M_{n,d}$; the proof has been deferred to appendix 1.

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

[6]It will be shown in theorem 3 that this map is well defined.

[7]It will be shown in theorem 3 that this map is well defined.

Theorem 3 *Consider the following diagram*

(4)

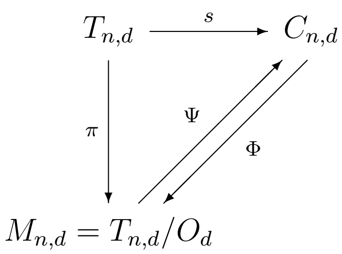

with the objects and maps as in definitions 1 and 2. We have the following:

(i) The maps Ψ and Φ are well defined.

(ii) The diagram is commutative.

(iii) The map Ψ is a homeomorphism with inverse Φ.

The space $M_{n,d}$ equipped with a differentiable structure. To be able to apply techniques from geometric optimization, the topological space $M_{n,d}$ requires a differentiable structure, preferably as a manifold. It turns out that $M_{n,d}$ is not a manifold, but a so-called *stratified space*. This follows from the fact that $T_{n,d}$ with the smooth action of the compact Lie group O_d forms (by definition) an O_d-manifold. For details the reader is referred to [DuK]. However there is a subspace of $M_{n,d}$ that is a manifold. This turns out to be sufficient for establishing convergence results for the geometric optimization techniques, see section 4.

Proposition 4 *Let $T^*_{n,d} \subset T_{n,d}$ be the subspace defined by*

$$T^*_{n,d} =: \left\{ \, Y \in T_{n,d} \; : \; \mathrm{rank}(Y) = d \, \right\}.$$

Then we have the following:

1. *$T^*_{n,d}$ is a submanifold of $T^*_{n,d}$*
2. *Denote by $M^*_{n,d}$ the quotient space $T^*_{n,d}/O_d$. Then $M^*_{n,d}$ is manifold.*
3. *$T^*_{n,d}$ is dense in $T_{n,d}$ and Mnd^* is dense in $M_{n,d}$. is a manifold.*

Proof:

1. It is enough to show that $T^*_{n,d}$ is open in $T_{n,d}$. Let $Y \in T^*_{n,d}$ and (U, Σ, V) be

the SVD of Y. Let $X = U\Sigma_d U^T$ where $\Sigma_d = \mathrm{diag}(\Sigma_{11}, \cdots, \Sigma_{d-1,d-1}, 0, \cdots, 0)$. Then, by Eckart-Young theorem we have

$$\min_{Z \in T_{n,d}; \mathrm{rank}(Z)<d} \|Y - Z\|^2 \geq \min_{Z \in \mathbb{R}^{n \times d}; \mathrm{rank}(Z)<d} \|Y - Z\|^2 = \|Y - X\|^2 = \Sigma_{d,d}^2.$$

Choose $\varepsilon \in (0, \Sigma_{d,d}^2)$. Then

$$\mathcal{U} := \{Z \in T_{n,d} \mid \|Y - Z\|^2 < \varepsilon\}$$

is open in $T_{n,d}$ and $\mathcal{U} \subseteq T_{n,d}^*$ and because Y was arbitrarily we have that $T_{n,d}^*$ is open in $T_{n,d}$.

2. This part is a corollary of theorem (17) by taking $M = T_{n,d}^*$ and $G = O_d$. Thus it is enough to show that conditions of the theorem are satisfied. We first show that action of O_d on $T_{n,d}^*$ is proper[8]. Let

$$\begin{array}{rcl} \Phi : T_{n,d}^* \times O_d & \rightarrow & T_{n,d}^* \times T_{n,d}^* \\ (m, g) & \mapsto & (mg^{-1}, m) \end{array}$$

and K a compact subset of $T_{n,d}^* \times T_{n,d}^*$. Then, by continuity Φ, $\Phi^{-1}(K)$ is closed in $T_{n,d}^* \times O_d$. Because $T_{n,d}^* \times O_d$ is relatively compact it follows that $\Phi^{-1}(K)$ is compact.

Finally, we show that O_d acts free on $T_{n,d}^*$. Let $Y \in T_{n,d}^*$ and $Q \in O_d$ such that $YQ^T = Y$. If we denote the rows of Y with y_i, we can rewrite this as follows

$$Qy_i = y_i, \qquad i = 1 : n.$$

This means that y_i are eigenvectors with eigenvalue 1. Because the $\mathrm{rank}(Y) = d$ we can find the sequence $\{y_1, .., y_d\}$ of d-independent rows of Y. These are independent eigenvectors of Q. Only square matrix with all eigenvalues equal to 1 is the identity matrix.

3. Suppose $Y \in T_{n,d}$ with $\mathrm{rank}(YY^T) = k$ where $1 < k < d$. Then we can choose a set of k rows vectors $y_1, .., y_k$ of Y such that they span k-dimensional linear subspace of $\mathbb{R}^d$. We denote this subspace with H_k.

For $i = k : d - 1$ we have the following:

Choose a linear subspace H_{i+1} of $\mathbb{R}^d$ such that $H_i \subset H_{i+1}$ and $\dim H_{i+1} = i + 1$. Then there exists a sequence $\{x_{i+1}^n\}_n \in H_{i+1} \backslash H_i$ such that $\|y_{i+1} - x_{i+1}^n\| \rightarrow 0$ as $n \rightarrow \infty$. Define new sequence $\{y_{i+1}^n\}_n = \{x_{i+1}^n / \|x_{i+1}^n\|\}_n$. Then we have the following:

[8] for definition see (16)

1. $\text{span}\{y_1, \cdots, y_i, y_{i+1}^n\} = H_{i+1}$ for all n,
2. $\{y_{i+1}^n\}_n \in S^{d-1}$,
3. $\|y_{i+1} - y_{i+1}^n\| \to 0$ as $n \to \infty$.

Define $\{Y_n\}_n \in \mathbb{R}^{n\times d}$ as follows:

$$Y_n := \begin{pmatrix} y_1 \\ \vdots \\ y_k \\ y_{k+1}^n \\ \vdots \\ y_d^n \\ y_{d+1} \\ \vdots \\ y_n \end{pmatrix}.$$

Then by construction $\{Y_n\} \in T_{n,d}^*$ and $\|Y - Y_n\| \to 0$ as $n \to \infty$.

Finally we show that $M_{n,d}^*$ is dense in $M_{n,d}$. Let $[Y] \in M_{n,d}$ and take a representant $Y \in [Y]$. Then $Y \in T_{n,d}$ and by previous result there is a sequence $\{Y_n\}_n \in T_{n,d}^*$ such that $\|Y - Y_n\| \to 0$ as $n \to \infty$. From continuity of the canonical projection π we have that $\lim_{n\to\infty} \pi(Y_n) = \pi(Y)$. Thus $M_{n,d}^*$ is dense in $M_{n,d}$. □

Choice of representation. In principle we could elect another manifold M and another Lie group G with M/G homeomorphic to the constraint set. 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 $M = T_{n,d}$ and $G = O_d$ then convenient expressions for geodesics etc. are obtained.

In the next section the geometric optimization tools are developed for $M_{n,d}^*$.

3 Optimization over manifolds

For the development of minimization algorithms on a manifold, certain objects of the manifold need to calculated explicitly, such as geodesics, parallel transport etc. In this section these objects are introduced and made explicit for $T_{n,d}$ and $M_{n,d}^*$.

From a theoretical view point 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 $T_{n,d}$ is viewed as subset of $\mathbb{R}^{n \times d}$. A point $Y \in T_{n,d}$ will be represented as a $n \times d$ matrix. If we think of $T_{n,d}$ as n products of S^{d-1} then it is convenient to represent Y as

$$Y = \left\{ \begin{pmatrix} Y_1 \\ \vdots \\ Y_n \end{pmatrix} : Y_i \in S^{d-1}, \right\}.$$

This representation reveals that to calculate geodesics and parallel transport on $T_{n,d}$ it is sufficient to calculate these on a single sphere.

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

Normal Space. The normal space is defined to be the orthogonal complement of the tangent space. Orthogonality depends on the metric chosen. We choose to embed the manifold in Euclidean space and then the inner product for two tangents Δ_1, Δ_2 is defined as

$$\langle \Delta_1, \Delta_2 \rangle = \mathrm{tr} \Delta_1 \Delta_2^T,$$

which is the Frobenius inner product for $n \times d$ matrices. It is then straightforward to verify that the normal space $N_Y T_{n,d}$ at $Y \in T_{n,d}$ is given by

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

It follows that the normal space is n dimensional. The projections π_N and π_T onto the normal and tangent space are given by

$$\pi_N(\Delta) = \mathrm{diag}(\Delta Y^T) Y \;\; \text{and} \;\; \pi_T(\Delta) = Y - \mathrm{diag}(\Delta Y^T) Y,$$

respectively.

Geodesics. Geodesics are also well known for the sphere. The geodesics on $T_{n,d}$ are then the products of geodesics on the sphere. The geodesic at $Y_i(0)$ in the direction Δ_i is given by

$$Y_i(t) = \cos(\; \|\Delta_i\| t \;) \; Y_i(0) + \frac{1}{\|\Delta_i\|} \sin \left(\; \|\Delta_i\| t \; \right) \; \Delta_i. \tag{5}$$

By differentiating, we obtain an expression for the evolution of the tangent along the geodesic:

$$\dot{Y}_i(t) = -\|\Delta_i\| \sin\left(\|\Delta_i\| t \right) Y_i(0) + \cos\left(\|\Delta_i\| t \right) \Delta_i. \tag{6}$$

Parallel transport. Again, we consider this problem per component on the sphere. If Δ_2 is parallel transported along a geodesic starting from Y in the direction of Δ_1, then decompose Δ_2 in terms of Δ_1,

$$\Delta_2 = \langle \Delta_1, \Delta_2 \rangle \Delta_1 + R, \quad R \perp \Delta_1.$$

Then Δ_1 changes according to equation (6) and R remains unchanged. Parallel transport from Y_1 to Y_2 defines a map $\tau(Y_1, Y_2) : T_{Y_1} T_{n,d} \to T_{Y_2} T_{n,d}$. When it is clear in between which two points is transported, then parallel transport is denoted simply by τ.

Geometry of the quotient space. We now study the regular part of the quotient space $M_{n,d}$. The regular part consists of all correlation matrices of rank d and is denoted by $M^*_{n,d}$. From the theory of G-manifolds it follows that $M^*_{n,d}$ is an open submanifold of $M_{n,d}$. As the geometry of the sphere is well known, we try to extend the results obtained on the sphere to $M^*_{n,d}$. The key techniques here are the horizontal and vertical spaces. At a point Y in some manifold M the vertical space consists of those tangents in $T_Y M$ that are tangent to the set $[Y] \subset M$. The vertical space at $Y \in M$ will be denoted by $V_Y M$, as long as it is clear what the equivalence relation is. The horizontal space is the orthogonal complement of the vertical space. The geometric idea is that when a point is moved along a tangent in the vertical set, then we remain in the equivalence class of the point. If a point is moved along a tangent in the horizontal set, then we actually move over equivalence classes in the quotient space.

In the developments below, it will be shown that a geodesic along a tangent in the horizontal space will always remain in the horizontal space, *for the case of the quotient of $T_{n,d}$ with the orthogonal group.* Thus from a representative of an equivalence class and a tangent in the horizontal space a geodesic on $M^*_{n,d}$ may be constructed. First we proceed by characterizing the vertical space.

Vertical space. The vertical space associated with $T_{n,d}$ and the orthogonal group is characterized by the following lemma.

Lemma 5 *The vertical space at a point $Y \in T_{n,d}$ and associated with the equivalence classes generated by the orthogonal group O_d is given by*

$$V = \left\{ \; \Delta \in T_Y T(n,d) \; ; \; \Delta = YZ, \;\; Z \in \mathbb{R}^{d \times d}, \;\; Z^T = -Z \; \right\}.$$

Proof:

- *$V \subset$ vertical space.* If $\Delta \in V$ then infinitesimally

$$Y + \varepsilon\Delta = Y \; \underbrace{\left(\; I_d + \varepsilon Z \; \right)}_{=:Q(\varepsilon)}$$

  Because the tangent space of the orthogonal group is given by the skew-symmetric matrices, it follows that $Q(\varepsilon)$ remains infinitesimally in the orthogonal group; thus $Y + \varepsilon\Delta$ remains infinitesimally in the equivalence class of Y, which was to be shown.

- *Vertical space $\subset V$.* If v is an element of the vertical space then infinitesimally

$$Y + \varepsilon v \in [Y],$$

  thus there exists $Q(\varepsilon) \in O_d$ such that

$$Y + \varepsilon v = YQ(\varepsilon).$$

  Applying a Taylor expansion to Q yields $Q(\varepsilon) = I_d + \varepsilon Z$, with Z skew-symmetric (because $Q(\varepsilon)$ remains in O_d). It follows that $v = YZ$, which was to be shown. □

The next lemma shows that there is a more simple way to represent the vertical space.

Lemma 6 *The vertical space is given by*

$$V_Y T_{n,d} = \left\{ \; \Delta = YZ \; ; \; Z \in \mathbb{R}^{d \times d}, \;\; Z^T = -Z \; \right\} \subset T_Y T_{n,d}.$$

In other words, all matrices of the form YZ with Z skew-symmetric are in the tangent space at Y.

Proof:

$$\begin{aligned} \langle Y, YZ \rangle &= \operatorname{tr}\left(\; Y(YZ)^T \; \right) \\ &= \operatorname{tr}\left(\; YZ^T Y^T \; \right) \\ &= -\operatorname{tr}\left(\; YZY^T \; \right) \\ &= -\operatorname{tr}\left(\; (YZ)Y^T \; \right) \\ &= -\langle YZ, Y \rangle \; = \; -\langle Y, YZ \rangle. \end{aligned}$$

Thus $\langle Y, YZ\rangle = 0$ and YZ is tangent at Y. $\square$

Geodesics on $M^*_{n,d}$. In this section we show: If a point is moved along a geodesic in the direction of a tangent in the horizontal space, then the tangent along the geodesic remains in the horizontal space. Geodesics in $M^*_{n,d}$ may thus be represented as a select set of restricted geodesics on $T_{n,d}$. We proceed by characterizing the horizontal space in an alternative fashion. The proofs of the following lemma and subsequent theorem have been deferred to appendix 2.

Lemma 7 *The tangent Δ is in the horizontal space if and only if*

$$\sum_{i=1}^{n} \langle \underbrace{\Delta_i^T}_{d\times 1} \underbrace{Y_i}_{1\times d}, Z\rangle = 0, \quad \forall Z \in \mathbb{R}^{d\times d}, \quad Z \textit{ skew-symmetric.}$$

Theorem 8 *Consider the manifold $T_{n,d}$ and the horizontal and vertical spaces associated with the action given in equation (3). If a point is moved along a geodesic in the direction of a tangent in the horizontal space, then the tangent along the geodesic remains in the horizontal space.*

It follows that if $Y(\cdot)$ is a geodesic in $T^*_{n,d}$ with initial tangent in the horizontal space, then $[Y(t)]$ is a geodesic in $M^*_{n,d}$.

Now we return to the geometry of the quotient space $M_{n,d}$. Note that $M^*_{n,d}$ is an open manifold of dimension $n(d-1) - d(d-1)/2$. We can decompose the quotient as follows:

$$M^*_{n,d} \cup \partial_{n,d}, \text{ with } \partial_{n,d} = M_{n,d} \backslash M^*_{n,d}.$$

Because $M_{n,d}$ is not a manifold we can not define geodesics on the whole space. However on $M^*_{n,d}$ we can. The geodesics on $T_{n,d}$ define a curve in $M_{n,d}$. This curve is defined for all t and for all $[Y_1]$ and $[Y_2]$ in $M_{n,d}$ there is such a curve that joins them. Any such curve restricted to $M^*_{n,d}$ is a geodesic.

Remark 9 (*Dimension counting argument*) If $n > d$ then the dimension of the irregular part of $M_{n,d}$ corresponding to the matrices of rank $d-1$ is at least 2 lower than the dimension of the regular part $M^*_{n,d}$.

Proof: Note that $\dim(M^*_{n,d}) = n(d-1) - \frac{1}{2}d(d-1)$. Then a straightforward calculation yields

$$\dim(M^*_{n,d}) - \dim(M^*_{n,d-1}) = n - d + 1 \geq 2.$$

Here the inequality follows from the assumption $n > d$. $\square$

The gradient. The *differential* F_Y of a function F is defined by

$$\langle F_Y, \Delta \rangle = \frac{\mathrm{d}}{\mathrm{d}t} F(Y(t))\big|_{t=0}, \quad \forall \Delta, \tag{7}$$

with $Y(\cdot)$ any curve starting from Y in the direction of Δ, i.e., $\dot{Y}(0) = \Delta$. The *gradient* ∇F of a function F is defined as the projection of the differential onto the tangent space. In our case,

$$\nabla F = \pi_T(F_Y) = F_Y - \mathrm{diag}(F_Y Y^T)Y.$$

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

$$F_{YY}(\Delta) = \frac{\mathrm{d}}{\mathrm{d}t} F_Y(Y(t))\big|_{t=0}.$$

Algorithms. We are now in a position to state algorithms 1 and 2 for the optimization over $M_{n,d}$. These are modified versions of those presented by [EAS99].

4 Convergence analysis

In this section we show that the geometric optimization algorithms of the previous section converge globally to a local minimum and we establish their rate of convergence.

Global convergence. Zangwill [Zan69] derived generic sufficient conditions when an iterative algorithm converges. The result is repeated in the developments below in a form adapted to the case of minimizing F over $M_{n,d}$.

Let M be a compact set. We specify a subset $\Omega \subset M$ called the *solution set*. Any point $Y \in \Omega$ is called a *solution*. An *(autonomous) iterative algorithm* is a map $A : M \to M \cup \{\mathrm{stop}\}$ with $A^{-1}(\{\mathrm{stop}\}) = \Omega$. The proof of the following theorem is adapted from the proof of theorem 1 in [Zan69].

Theorem 10 (Global convergence) *Consider the problem of minimizing the function F of problem (2) over $M_{n,d}$ by use of algorithms 1 and 2. Suppose given a fixed tolerance level ε. A point Y is called a solution if $\|\nabla F(Y)\| < \varepsilon$. The algorithms 1 and 2 are extended with the following:*

- *Stop whenever $\|\nabla F(Y^{(i)})\| < \varepsilon$.*

Algorithm 1 Newton's Method for Minimizing $F(Y)$ on $M_{n,d}$

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

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

1: **for** $k = 0, 1, 2, \ldots$ **do**

2: Compute $G^{(k)} = \nabla F(Y^{(k)}) = F_Y - \mathrm{diag}(F_Y Y^T)Y$.

3: Compute $\Delta^{(k)} = -\mathrm{Hess}^{-1} G^{(k)}$ such that $\mathrm{diag}(\Delta^{(k)}(Y^{(k)})^T) = 0$ and

$$F_{YY}(\Delta^{(k)}) - \mathrm{diag}\big(\, F_{YY}(\Delta^{(k)})(Y^{(k)})^T \,\big)Y^{(k)} = -G^{(k)}.$$

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

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

6: **end for**

Algorithm 2 Conjugate Gradient for Minimizing $F(Y)$ on $T_{n,d}$

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

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

1: Compute $G^{(0)} = \nabla F(Y^{(0)})$ and set $H^{(0)} = -G^{(0)}$.

2: **for** $k = 0, 1, 2, \ldots$ **do**

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

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

5: Compute $G^{(k+1)} = \nabla F(Y^{(k+1)})$.

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

7: Compute the new search direction

$$H^{(k+1)} = -G^{(k+1)} + \gamma_k \tau H^{(k)} \quad \text{where} \quad \gamma_k = \frac{\big\langle G^{(k+1)} - \tau G^{(k)}, G^{(k+1)} \big\rangle}{\big\langle G^{(k)}, G^{(k)} \big\rangle}$$

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

9: **end for**

- *If we end up in the irregular part or if the Newton step fails, perform a steepest descent step as if we were minimizing the function F over the manifold $T_{n,d}$.*

- *Given a search direction in $M_{n,d}$, continue along a geodesic until there is no further decrease in the gradient ∇F.*

If the function F is twice continuously differentiable then for the augmented algorithms we have: From any starting point $Y^{(0)} \in M_{n,d}$ the algorithm either stops at a solution or produces an infinite sequence of points none of which are solutions, for which the limit of any convergent subsequence is a solution point.

Proof: Without loss of generality we may assume that the procedure generates an infinite sequence of points $\{Y^{(i)}\}$ none of which are solutions. Since $M_{n,d}$ is compact any subsequence must contain a convergent subsequence. It remains to prove that the limit of any convergent subsequence must be a solution.

First, note that the algorithm $A(\cdot)$ is continuous in virtue of the twice continuous differentiability of $F(\cdot)$. Namely the next point $A(Y)$ is entirely defined in terms of components up to the second derivative – all continuous by assumption. Second, note that if $Y^{(i)}$ is not a solution then

$$\|\nabla F(Y^{(i+1)})\| = \|\nabla F(A(Y^{(i)}))\| < \|\nabla F(Y^{(i)})\|. \tag{8}$$

Namely if $Y^{(i)}$ is not a solution then its gradient is non-negligible, thus at least by steepest descent on $T_{n,d}$ we will find a point $Y^{(i+1)} := A(Y^{(i)})$ with strictly smaller gradient norm. Third, note that the sequence $\{\|\nabla F(Y^{(i)})\|\}_{i=1}^{\infty}$ has a limit since it is monotonically decreasing and bounded from below by 0.

Let $\{Y^{(i_j)}\}_{j=1}^{\infty}$ be a subsequence that converges to Y^*, say. It must be proven that Y^* is a solution. Assume Y^* is not a solution. By continuity of the iterative procedure $A(Y^{(i_j)}) \to A(Y^*)$. Thus by the continuity of $\|\nabla F(\cdot)\|$, we have

$$\|\nabla F(A(Y^{(i_j)}))\| \downarrow \|\nabla F(A(Y^*))\| < \|\nabla F(Y^*)\|,$$

which is in contradiction with $\|\nabla F(A(Y^{(i_j)}))\| \to \|\nabla F(Y^*)\|$. □

Local rate of convergence. The proof of the following proposition may be found in for example (in whole or in part) [Smi93], [EAS99], [DPM03] or [AMS02].

Proposition 11 (Local rate of convergence) *For the conjugate gradient algorithms we have locally in the regular part $M^*_{n,d}$ a quadratic rate of convergence. For the Riemannian-Newton algorithm we have locally in the regular part $M^*_{n,d}$ a quadratic rate of convergence if the function F is non-degenerate at the local minimum.*

We conclude this section by remarking that the algorithms are stable numerically due to explicit and analytical knowledge of the geodesics, see also [Yan99], page 890.

5 Case distance minimization

In this section the primary concern of this paper to minimize $F(X) = \frac{1}{2}\|X - C\|_F^2$ is studied. The outline of this section is as follows. First, some particular choices for n and d are examined. Second, the differential of F is calculated. Third, the connection with the Zhang-Wu algorithm is stated; in particular this will lead to (i) an identification method of wether a local minimum is a global minimum, and (ii) an appropriate choice for the starting point of the algorithm. Fourth, an objective function F with weights on the individual elements of the matrix is studied (Hadamard norm).

Some particular choices for n and d.

Case $d = n$. The case that C is a (not necessary positive definite) symmetric matrix and the closest positive definite matrix X is to be found allows a successive projection solution. This was shown by [Hig02].

Case 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 the other 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.

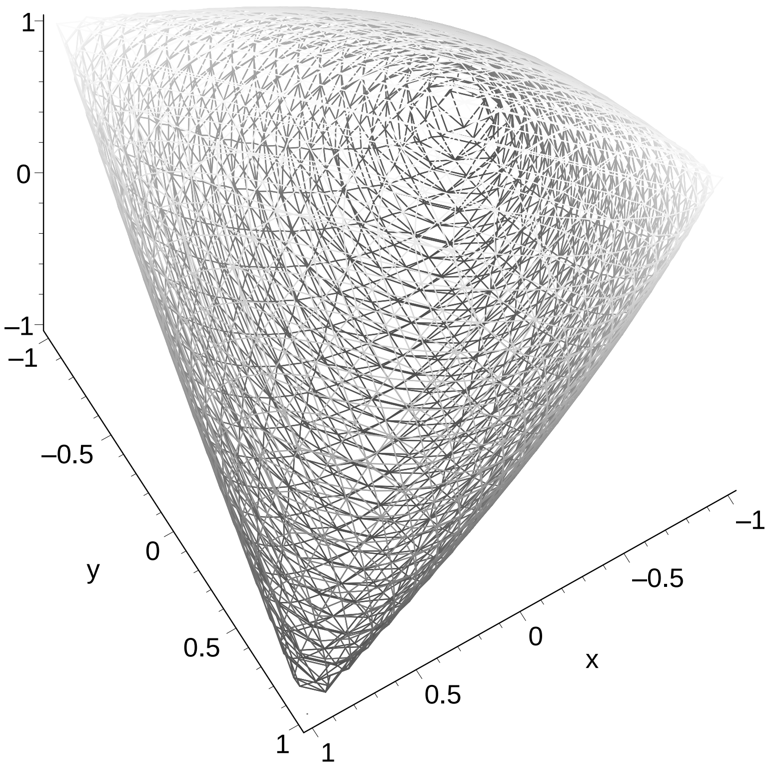

Figure 1: The shell represents the set of 3×3 correlation matrices of rank 2 or less.

Formula for the differential of F. Consider the specific case

$$F(Y) = \frac{1}{2}\|YY^T - C\|^2 = \frac{1}{2}\langle\psi, \psi\rangle,$$

with $\psi := YY^T - C$. Then

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

Thus from equation (7) we have

$$F_Y = 2\psi Y. \tag{9}$$

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 Zhang-Wu algorithm.

Lemma 12 *Let $Y \in T_{n,d}$ be such that $\nabla F(Y) = 0$. Here F is associated with the symmetric matrix C. Define*

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

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

$$C(\lambda) = QDQ^T, \ \ YY^T = QD^*Q^T$$

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).

Proof: 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 $\pi_N(F_Y)$ of F_Y onto the normal space at Y. Note also that

$$YY^T + \bar{C}(\lambda) = C(\lambda). \tag{10}$$

We calculate

$$\bar{C}(\lambda)Y = \{ \ -\psi + \lambda \ \}Y = -\frac{1}{2}F_Y + \frac{1}{2}\pi_N\big(\ F_Y \ \big) = 0. \tag{11}$$

The last equality follows from the assumption that the gradient of F at Y is zero $\nabla F(Y) = 0$, i.e., the differential F_Y is normal at Y. It follows from equation (11) 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 equation (10). Suppose $\bar{C}(\lambda) = Q\bar{D}Q^T$. From (i) we then have that $D^*_{ii} = 0$ if and only if $\bar{D}_{ii} \neq 0$. The result now follows since YY^T is positive semidefinite and has rank less than or equal to d.□

Lemma 12 will allow us to identify whether a local minimum is also a global minimum. The characterization of the global minimum for problem (1) was first achieved in [ZhW03] and [Wu03], 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 norm). Denote for $\lambda \in \mathbb{R}^n$: $C(\lambda) = C + \text{diag}(\lambda)$.

Theorem 13 (Characterization of the global minimum of problem (1), see [ZhW03] and [Wu03]) *Let C be a symmetric matrix. Let λ^* be such that there exists $\{C + \text{diag}(\lambda^*)\}_d \in C_{n,d}$ with*

$$\text{diag}\big(\ \{C + \text{diag}(\lambda^*)\}_d \ \big) = \text{diag}(C). \tag{12}$$

Then $\{C + \text{diag}(\lambda^*)\}_d$ is a global minimizer of problem (1).

Proof (Repeated here for clarity): Define the Lagrangian

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

$$V(\lambda) := \min \big\{ \ \mathcal{L}(X, \lambda) \ : \ \text{rank}(X) = d \ \big\}. \tag{13}$$

Note that the minimization problem in equation (13) is attained by any $\{C(\lambda)\}_d$ (see e.g., equation (30) of [Wu03]). For any $X \in C_{n,d}$,

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

(This is the equation at the end of the proof of theorem 4.4 of [ZhW03].) Here equation

(a) is obtained from the property that $X \in C_{n,d}$,

(b) is by definition of V, and

(c) is by assumption of equation (12). □

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

Theorem 14 *Let $Y \in T_{n,d}$ be such that $\nabla F(Y) = 0$. Let λ and $C(\lambda)$ be defined as in lemma 12. If YY^T has the d largest eigenvalues from $C(\lambda)$ (in the norm) then YY^T is a global minimizer to the problem (1).*

Proof: Apply lemma 12 and theorem 13. □

Initial feasible point [Flu88]. To obtain an initial feasible point $[Y] \in M_{n,d}$ close to the global minimum, we first perform an eigenvalue decomposition

$$C = Q\Lambda Q^T, \quad |\Lambda_{11}| \geq \cdots \geq |\Lambda_{nn}|. \tag{14}$$

Then we define Y by assigning to each row

$$Y_i = \frac{1}{\|\{\ Q_d\Lambda_d\ \}_i\|}\{\ Q_d\Lambda_d\ \}_i$$

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. The scaling is to ensure that each row of Y is of unit length. If row i is a priori of zero length, then we choose Y_i to be an arbitrary vector in $\mathbb{R}^d$.

Note that the condition of decreasing norm in equation (14) is thus key to ensure that the initial point is close to the global minimum, c.f. the result of theorem 14.

Weighted norms. The Frobenius norm in the objective function F can be replaced by (i) a weighted Frobenius norm or (ii) a norm with arbitrary weights per element of the matrix (this is a so-called Hadamard norm). The developed optimization theory can be applied to both cases. For the case (i) of the weighted Frobenius norm we obtain an analytical expression for the differential in the same way as we did in equation (9). For the case (ii) of the Hadamard norm however such an analytical expression is not known to us. Nevertheless, the differential can be approximated by means of finite differences.

6 Numerical results case distance minimization

The outline of this section is as follows. First, the numerical performance is compared of geometric optimization, Zhang and Wu's algorithm and parameterized optimization. Second, various comparative features between the algorithms (including also Dykstra's algorithm) are discussed.

Acknowledgement. Our implementation of geometric optimization over low-rank correlation matrices is an adoption of the 'SG min' template of [EdL00] (written in MATLAB) for optimization over the Stiefel and Grassmann manifolds. This template contains four distinct algorithms for geometric optimization: Newton, dog-leg step algorithm, Polak-Ribière conjugate gradient and Fletcher-Reeves conjugate gradient. For a description see section 8.3.2 of [EdL00].

Numerical comparison. The performance of the following algorithms is compared:

1. The geometric optimization algorithms developed in this paper.
2. The Zhang-Wu algorithm, see [ZhW03] and [Wu03].
3. The parametrization methods of [Reb99], [Bri02] and [BMR02]. Here the set of low-rank correlation matrices is parameterized by trigonometric functions. In essence, this approach is the same as geometric optimization, bar the key difference that our approach yields simple expressions for the gradient, etc. This leads to an efficient implementation as shown by the numerical results in this section.

The two types of random correlation matrices are described as follows:

- *Random Davies-Higham correlation matrices.* The key reference here is [DaH00]. This random correlation matrix generator is available in the MATLAB gallery function 'randcorr'.
- *Random 'financial' or 'DJDP' correlation matrices.* A parametric form for (primarily financial) correlation matrices is posed in [DJDP02], equation (8). We repeat here the parametric form for completeness.

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

  with $\gamma_1, \gamma_2, \gamma_4 > 0$ and with T_i denoting the expiry time of rate i. (Our particular choice is $T_i = i$, $i = 1, 2, \ldots$) This model was then

Table 1: Excerpt of table 3 in [DJDP02].

	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

subsequently estimated with USD historical interest rate data. In table 3 of [DJDP02] the estimated γ parameters are listed, along with their standard error. An excerpt of this table has been displayed in table 1. The random DJDP matrix that we used is obtained by randomizing the γ-parameters. We assumed the γ-parameters distributed normally with mean and standard errors given by table 1, with $\gamma_1, \gamma_2, \gamma_4$ capped at zero.

Each of the Lagrange (Zhang-Wu), Newton, dog, Polak-Ribière conjugate gradient (prcg), Fletcher-Reeves conjugate gradient (frcg), and parameterized algorithms were applied to reduce the matrices to the desired rank, given the available CPU time. For each algorithm, this produces a distance δ to the original matrix

$$\delta_{\text{algo}} = \|X_{\text{algo}} - C\|_F^2.$$

Here X_{algo} denotes the rank reduced matrix found by algorithm 'algo' and C denotes the original matrix.

The worst of geometric optimization and Zhang-Wu versus the parametrization method. In the left panel of figure 2 we plot the empirical distribution of the data

$$\frac{\delta^{(i)}_{\text{param}} - \delta^{(i)}_{\text{worst}}}{\delta^{(i)}_{\text{param}}}, \text{ with } \delta^{(i)}_{\text{worst}} := \max\{\delta^{(i)}_{\text{lagrange}}, \delta^{(i)}_{\text{newton}}, \delta^{(i)}_{\text{dog}}, \delta^{(i)}_{\text{prcg}}, \delta^{(i)}_{\text{frcg}}\}. \tag{15}$$

The allowed CPU time for either Lagrange or geometric optimization is 2 seconds. The parametrization method however was allowed 4000 iterations; this led to CPU times of 4-11 seconds.

The worst of geometric optimization versus the parametrization method. In the right panel of figure 2 we plot (for a different set of random matrices) the empirical distribution of the data in equation (15) without comparison with the Lagrange algorithm:

$$\delta^{(i)}_{\text{worst}} := \max\{\delta^{(i)}_{\text{newton}}, \delta^{(i)}_{\text{dog}}, \delta^{(i)}_{\text{prcg}}, \delta^{(i)}_{\text{frcg}}\}.$$

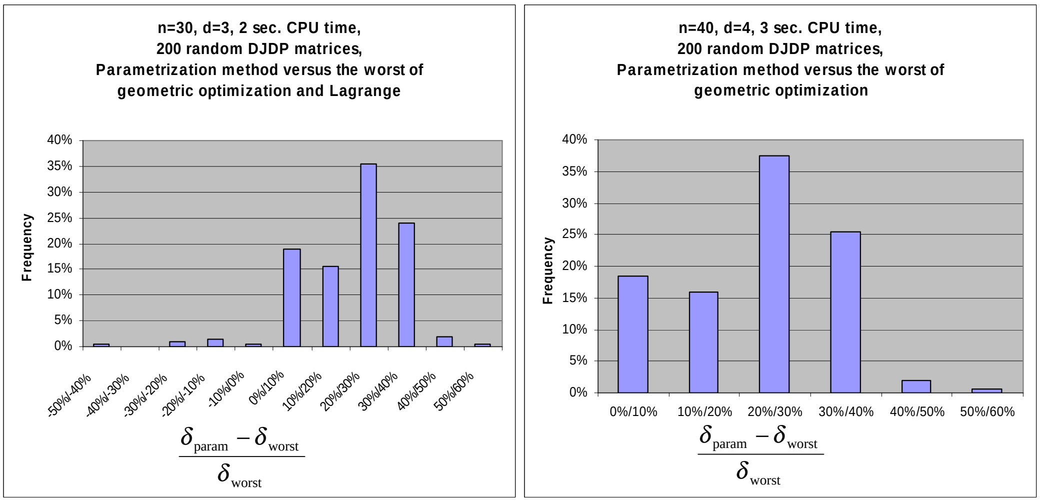

Figure 2: Results versus the parametrization method.

The allowed CPU time for geometric optimization is 3 seconds. The parametrization method however was allowed 2000 iterations; this led to CPU times of 6-17 seconds.

Geometric optimization versus the Zhang-Wu algorithm. To jointly compare distances produced by the algorithms over the various matrices, all lengths are normalized by the nearest found matrix. We obtain the data

$$\left(\frac{\delta^{(i)}_{\text{lagrange}}}{\delta^{(i)}_{\text{min}}} \quad \frac{\delta^{(i)}_{\text{newton}}}{\delta^{(i)}_{\text{min}}} \quad \frac{\delta^{(i)}_{\text{dog}}}{\delta^{(i)}_{\text{min}}} \quad \frac{\delta^{(i)}_{\text{prcg}}}{\delta^{(i)}_{\text{min}}} \quad \frac{\delta^{(i)}_{\text{frcg}}}{\delta^{(i)}_{\text{min}}} \right)$$

with

$$\delta^{(i)}_{\text{min}} = \min \left\{ \delta^{(i)}_{\text{lagrange}},\ \delta^{(i)}_{\text{newton}},\ \delta^{(i)}_{\text{dog}},\ \delta^{(i)}_{\text{prcg}},\ \delta^{(i)}_{\text{frcg}} \right\}.$$

The distribution of this data for the two sets of matrices has been displayed in figures 3 and 4.

Discussion of the results. From figures 3 and 4 it becomes clear that geometric optimization is by far superior over the Lagrange multiplier algorithm of Zhang and Wu, for the particular numerical cases studied. In turn, the Zhang-Wu algorithm is by far superior over the parametrization method, for the numerical case investigated in the left panel of figure 2.

Comparative features of the algorithms. Various comparative features of the algorithms for finding the nearest correlation matrix are listed

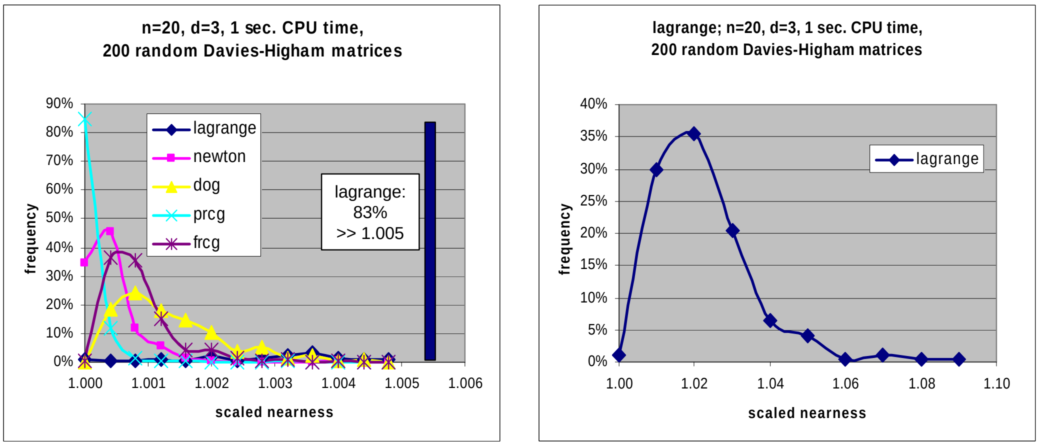

Figure 3: Results on random Davies-Higham matrices.

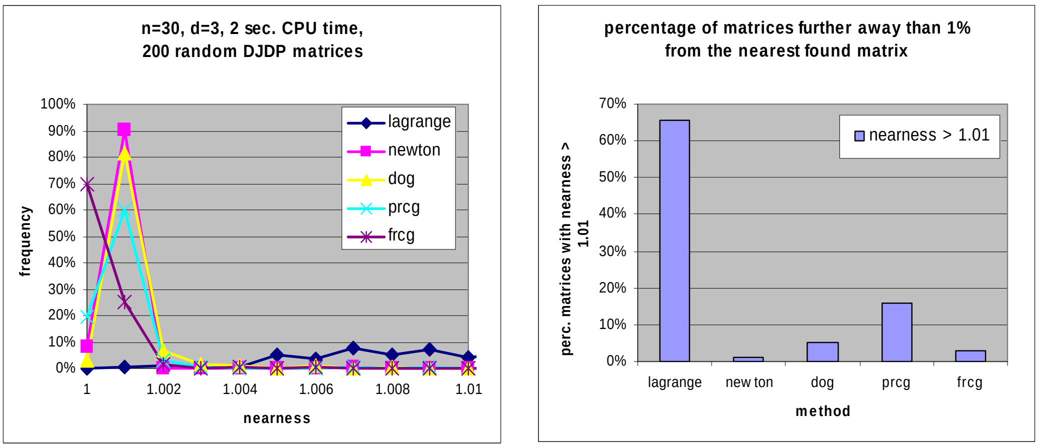

Figure 4: Results on random DJDP matrices.

Table 2: Comparative features of the algorithms.

feature	**Zhang-Wu**	**Dykstra, Han, Higham**	**geometric optimization**
convergence	not guaranteed	global	global
if converged	global minimum	global minimum	local minimum
rate of convergence	at best linear	at best linear	locally quadratic
freedom of norm	weighted Frobenius	weighted Frobenius	any norm
restriction on target matrix	positive definiteness	none	none
neighbourhood search	no	no	yes
rank conditions	yes	no	yes

in table 2. Here the algorithm of Dykstra, Han and Higham is the successive projection algorithm ([Dyk83], [Han88]) applied in [Hig02] to the problem of finding the nearest (possibly full-rank) correlation matrix. Some remarks explaining the statements in the table are given below:

- *Global convergence.* The Zhang-Wu algorithm is guaranteed to have at least one accumulation point by proposition 4.1 of [ZhW03]. The same proposition then states that: 'if the accumulation point has distinct d^{th} and $d+1^{\text{th}}$ eigenvalues, then that point is a global minimizer'. As the distinct eigenvalue condition however does not hold in general (we have seen counterexamples in our numerical experiments), it follows that global convergence is not guaranteed. The Dykstra, Han and Higham and geometric optimization algorithms are guaranteed to converge globally. For details the reader is referred to [Hig02] and proposition 10, respectively.

- *Rate of convergence.* Methods such as the Zhang-Wu and Higham methods can be applied to linear sets; for linear sets the rate of convergence of the alternating projections method is linear, see [Deu83] and [DeH97]. Therefore we can expect at best linear convergence for the Higham and Zhang-Wu algorithms. See also the remark after algorithm 3.3 in [Hig02]. For the convergence rate of the geometric optimization algorithms, see proposition 11.

- *Restriction on the target matrix.* In the definition of the Zhang-Wu algorithm matrices are formed that take the d largest (in norm) eigenvalues of a given matrix (also of the target matrix). If the target matrix is not positive semidefinite, some of these eigenvalues at the optimal Lagrange multiplier may be negative. It follows that the algorithm would produce matrices that are not positive semidefinite.

For explanations of properties of the algorithms that have not been explained in the above remarks the reader is referred to [ZhW03], [Hig02] and the current text, respectively. Note that the entries for the parametrization method in table 2 would be identical to the entries of the geometric optimization method. The distinctive difference (not listed in the table) is that for geometric optimization, we obtain simple expressions for the gradient, Hessian etc. Theoretically, that leads to a more efficient implementation. This has also been shown numerically by the tests reported in figure 2.

7 The Thomson problem

Table 3: Energy minimization results for various n after 20 iterations. The column [HSS94] are conjecturally minimal configurations.

n	this paper: newton	dog	prcg	frcg	[DeS98]	[HSS94]
5	6.4752	6.4758	6.4789	6.4771		6.4747
10	32.7210	32.7204	32.7210	32.8403		32.7169
25	243.8983	243.8495	243.8544	244.4211		243.8128
50	1055.668	1055.781	1055.695	1056.737	1055.5128	1055.1823

The purpose of this section is to show that the optimization tools developed in this paper can be applied to other areas.

The Thomson problem is concerned with minimizing the potential energy of n charged particles on the 2-sphere $S^2 \subset \mathbb{R}^3$ ($d = 3$). The associated potential is given by

$$F(Y) = \sum_{i=1}^{n} \sum_{j=i+1}^{n} \frac{1}{\|Y_i - Y_j\|_2}.$$

Note that the potential energy is invariant under the action of the orthogonal group; the minimization problem is thus over the set $M_{n,d}$. This problem has a long history, see for example `www.ogre.nu/sphere.htm`.

Geometric optimization techniques have previously been applied to the Thomson problem in the literature in [DeS98], but these authors have only considered conjugate gradient techniques on $T_{n,d}$. In comparison, *we stress here that* our approach considers the quotient space $M_{n,d}$, which allows for Newton's algorithm (the latter not developed in [DeS98]). For numerical illustration, the results for various dimensions n have been included in table 3.

8 Conclusions

We applied geometric optimization tools for finding the nearest low-rank correlation matrix. Despite the involved differential geometric machinery, it is interesting to see that the approach results in an algorithm more efficient than any existing algorithm in the literature (for the numerical cases considered). The geometric approach also allows for insight and more intuition into the problem.

References

[AMR88] Abraham, R., Marsden, J.E., Ratiu, T.: Manifolds, Tensor Analysis, and Applications. *New York: Springer* (1988)

[AMS02] Absil, P.-A., Mahoney, R., Sepulchre, R.: Riemannian geometry of Grassmann manifolds with a view on algorithmic computation. *Submitted to Acta Applicandae Mathematicae* (2002)

[BMR02] Brigo, D., Mercurio, F., Rapisarda, F.: Parametrizing correlations: A geometric interpretation. *Milan: Banca IMI Working Paper* (2002) `www.fabiomercurio.it`

[Bri02] Brigo, D.: A Note on Correlation and Rank Reduction. *Milan: Banca IMI Working Paper* (2002) `www.damianobrigo.it`

[DaH00] Davies, P.I., Higham, N.J.: Numerically Stable Generation of Correlation Matrices and Their Factors. *BIT* **40** 640-651 (2000)

[DeH97] Deutsch, F., Hundal, H.: The rate of convergence for the method of alternating projections, II. *Journal of Mathematical Analysis and its Applications* **205** 381-405 (1997)

[DeS98] Depczynski, U., Stöckler, J.: A Differential Geometric Approach to Equidistributed Knots on Riemannian Manifolds. *Approximation Theory IX, Volume I: Theoretical Aspects, Nashville: Vanderbilt University Press* (1998)

[Deu83] Deutsch, F.: Von Neumann's alternating method: The rate of convergence. *Approximation Theory IV (C.K. Chui, L.L. Schumaker & J.D. Ward, eds.) New York: Academic Press* 427-434 (1983)

[DJDP02] De Jong, F., Driessen, J., Pelsser, A.A.J.: On the Information in the Interest Rate Term Structure and Option Prices. *Working paper* (2002)

[DPM03] Dedieu, J.-P., Priouret, P., Malajovich, G.: Newton's Method on Riemannian Manifolds: Covariant Alpha-Theory. *Working paper* (2003)

[DuK] Duistermaat J.J., Kolk J.A.C.: Lie Groups *Springer-Verlag*2000

[Dyk83] Dykstra, R.L.: An algorithm for restricted least squares regression. *Journal of the American Statistical Association* **87** (384) 837-842 (1983)

[EAS99] Edelman, A., Arias, T.A., Smith, S.T.: The geometry of algorithms with orthogonality constraints. *SIAM Journal of Matrix Analysis and its Applications* **20** (2) 303-353 (1999)

[EdL00] Edelman, A., Lippert, R.: Nonlinear Eigenvalue Problems with Orthogonality Constraints (Section 8.3). In Bai, Z., Demmel, J., Dongarra, J., Ruhe, A. and Vorst, H. van der, editors: Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide. *Philidelphia: SIAM* (2000)

[Flu88] Flury, B.: Common Principal Components and Related Multivariate Models. *New-York: Wileys*

[Han88] Han, S.-P.: A successive projection method. *Mathematical Programming* **40** 1-14 (1988)

[Hig02] Higham, N.J.: Computing the nearest correlation matrix–a problem from finance. *IMA Journal of Numerical Analysis* **22** 329-343 (2002)

[HSS94] Hardin, R.H., Sloan, N.J.A., Smith, W.D.: Minimal energy arrangements of points on a sphere. `www.research.att.com/~njas/electrons/index.html` (1994)

[Mun75] Munkres, J.R.: Topology. *Prentice-Hall* (1975)

[Reb99] Rebonato, R.: On the simultaneous calibration of multifactor lognormal interest rate models to Black volatilities and to the correlation matrix. *Journal of computational finance* (1999)

[Smi93] Smith, S.T.: Geometric Optimization Methods for Adaptive Filtering. *PhD thesis, Harvard University, Cambridge, Massachusetts* (1993)

[Wu03] Wu, L.: Fast At-The-Money Calibration of the Libor Market Model Using Lagrange Multipliers. *Journal of Computational Finance* **6** (2) 39-77 (2003)

[Yan99] Yang, Y.: Optimization on Riemannian Manifold. *Proceeding of the* 38^{th} *Conference on Decision and Control, Phoenix, Arizona* (1999)

[Zan69] Zangwill, W.I.: Convergence Conditions for Nonlinear Programming Algorithms. *Management Science-A* **16** (1) 1-13 (1969)

[ZhW03] Zhang, Z., Wu, L.: Optimal Low-Rank Approximation to a Correlation Matrix. *Linear Algebra and its Applications* **364** 161187 (2003)

Appendix 1: Proof of theorem 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 $\exists Q \in O_d :\ Y_2 = Y_1 Q$. If 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 $X = YY^T$.

(B) If $Y_i \in T_{n,d}$, $i = 1, 2$ with $Y_1 Y_1^T = Y_2 Y_2^T =: X$ then there exists $Q \in O_d$ such that $Y_1 = Y_2 Q$.

Ad (A): Let

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

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

$$Q\sqrt{\Lambda} = \left(\; (Q\sqrt{\Lambda})(:, 1 : d) \; \middle| \; 0 \; \right).$$

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

Ad (B): Apply Gram-Schmidt jointly to the sets of vectors

$$\left\{ \; Y_1^{(i)}, \ldots Y_n^{(i)} \; \right\} \subset \mathbb{R}^d, \;\; i = 1, 2.$$

We find two orthogonal sets of vectors

$$\left\{ \; Z_1^{(i)}, \ldots Z_{e^{(i)}}^{(i)} \; \right\} \subset \mathbb{R}^d, \;\; i = 1, 2,$$

with $e^{(i)} \leq d$, such that

$$\mathrm{span}\big(\; Y_1^{(i)}, \ldots Y_n^{(i)} \; \big) = \mathrm{span}\big(\; Z_1^{(i)}, \ldots Z_{e^{(i)}}^{(i)} \; \big).$$

The key idea in the proof of item (B) is that the Gram-Schmidt procedure is defined in terms of the vector products of the Y vectors; and these are equal for both the $Y^{(1)}$ and $Y^{(2)}$ vectors and determined by the vector product matrix X. The first result that follows from this idea is that the spaces spanned by $Z^{(1)}$ and $Z^{(2)}$ have equal dimensions,

$$e := e^{(1)} = e^{(2)}.$$

Note that there exists an orthogonal transformation $Q \in O_d$ such that[9]

$$Z_k^{(1)} Q = Z_k^{(2)}, \quad k = 1, \ldots, e.$$

Claim. We have

$$Y_j^{(1)} Q = Y_j^{(2)}, \quad j = 1, \ldots, n.$$

Note that the proof of item (B) is complete when it is shown that this claim holds.

Proof of claim. The proof follows by induction on j.

- For $j = 1$, note that $Z_1^{(i)} = Y_1^{(i)}$ since $\|Y_1^{(i)}\| = 1$. Thus

$$Y_1^{(1)} Q = Z_1^{(1)} Q = Z_1^{(2)} = Y_1^{(2)}.$$

- Suppose the claim holds for $j' = 1, \ldots, j-1$. Then write

$$\begin{aligned} Y_j^{(i)} &= P_j^{(i)} + R_j^{(i)}, \quad \text{with} \\ P_j^{(i)} &= \sum_{j'=1}^{j-1} \langle Y_{j'}^{(i)}, Y_j^{(i)} \rangle Y_{j'}^{(i)}, \quad \text{and} \\ R_j^{(i)} &= Y_j^{(i)} - P_j^{(i)}. \end{aligned} \tag{16}$$

[Projection P and remainder R.] Note that we have $P_j^{(1)} Q = P_j^{(2)}$ since

$$\begin{aligned} P_j^{(1)} Q &= \sum_{j'=1}^{j-1} \langle Y_{j'}^{(1)}, Y_j^{(1)} \rangle Y_{j'}^{(1)} Q \\ &\overset{(a)}{=} \sum_{j'=1}^{j-1} \langle Y_{j'}^{(1)}, Y_j^{(1)} \rangle Y_{j'}^{(2)} \\ &\overset{(b)}{=} \sum_{j'=1}^{j-1} \langle Y_{j'}^{(2)}, Y_j^{(2)} \rangle Y_{j'}^{(2)} = P_j^{(2)}. \end{aligned} \tag{17}$$

[9] If $e = d$ then such Q is determined uniquely, else if $e < d$ then there exists more than one such Q. In the latter case, the Z may be extended to a full orthogonal basis of $\mathbb{R}^d$. Given (extended) Z matrices then Q is uniquely determined by $Q = (Z^{(1)})^T Z^{(2)}$.

Equality (a) follows by the induction hypothesis and equality (b) by the joint property $Y^{(i)}(Y^{(i)})^T = X$.

A formula for the length of $R_j^{(i)}$ is given by

$$\begin{aligned}\|R_j^{(i)}\|^2 &= \langle Y_j^{(i)} - P_j^{(i)}, Y_j^{(i)} - P_j^{(i)}\rangle \\ &= \|Y_j^{(i)}\|^2 - 2\sum_{j'=1}^{j-1}\langle Y_{j'}^{(i)}, Y_j^{(i)}\rangle^2 \\ &\quad + \sum_{j_1=1}^{j-1}\sum_{j_2=1}^{j-1}\langle Y_{j_1}^{(i)}, Y_j^{(i)}\rangle\langle Y_{j_2}^{(i)}, Y_j^{(i)}\rangle\langle Y_{j_1}^{(i)}, Y_{j_2}^{(i)}\rangle \\ &= X_{jj} - 2\sum_{j'=1}^{j-1} X_{j'j} + \sum_{j_1=1}^{j-1}\sum_{j_2=1}^{j-1} X_{j_1,j}X_{j_2,j}X_{j_1,j_2}.\end{aligned}$$

Thus $\|R_j\| := \|R_j^{(1)}\| = \|R_j^{(2)}\|$.

If $R_j^{(1)} = 0$ then $R_j^{(2)} = 0$ and from equations (16) and (17) it follows

$$Y_j^{(1)}Q = P_j^{(1)}Q = P_j^{(2)} = Y_j^{(2)}$$

and the induction step holds. Without loss of generality we may thus assume $R_j \neq 0$. Denote by $k(j)$ the minimum integer such that

$$\mathrm{span}\big(\, Z_1^{(i)}, \ldots, Z_{k(j)}^{(i)} \,\big) = \mathrm{span}\big(\, Y_1^{(i)}, \ldots, Y_j^{(i)} \,\big).$$

Note that the definition of $k(\cdot)$ is independent of i, since the Gram-Schmidt process is defined in terms of the inner products that are common for both $i = 1, 2$ and determined by X. Note that the new Gram-Schmidt vector is given by the normalized remainder term R:

$$Z_{k(j)}^{(i)} = \frac{1}{\|R_j\|} R_j^{(i)}$$

and

$$R_j^{(1)}Q = \|R_j\| Z_{k(j)}^{(1)}Q = \|R_j\| Z_{k(j)}^{(2)} = R_j^{(2)}.$$

Thus $Y_j^{(1)} = P_j^{(1)} + R_j^{(1)}$ and both $P_j^{(1)}$ and $R_j^{(1)}$ map to their respective counterparts under Q. By linearity, $Y_j^{(1)}$ then is mapped to its respective counterpart $Y_j^{(2)}$, which was to be shown.

Proof of (ii). Diagram (4) is commutative: To show that $\Psi \circ \pi = s$: Let $Y \in T_{n,d}$, then $\pi(Y) = [Y]$ and $\Psi([Y]) = YY^T$ and also $s(Y) = YY^T$. To show that $\Phi \circ s = \pi$: Let $Y \in T_{n,d}$, then $s(Y) = YY^T$ and $\Phi(YY^T) = [Y]$ and also $\pi(Y) = [Y]$.

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 [AMR88], proposition 1.4.8). In our case, $\Psi \circ \pi = s$ with $s(Y) = YY^T$ 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 [Mun75], theorem 5.6). □

Appendix 2: The proofs of section 3

First we prove lemma 7 and second we prove theorem 8.

Proof of lemma 7: For $\Delta \in T_Y T_{n,d}$,

$$\begin{aligned} \Delta \in \mathrm{H_Y T_{n,d}} &\Leftrightarrow \langle \Delta, X \rangle = 0, \ \ \forall X \in \mathrm{V_Y T_{n,d}} \\ &\Leftrightarrow \langle \Delta, YZ \rangle = 0, \ \ \forall Z \in \mathbb{R}^{d\times d}, Z \text{ skew-symmetric.} \end{aligned} \tag{18}$$

Now $\langle \Delta, YZ \rangle = \mathrm{tr}(\Delta Z^T Y^T)$. We calculate $\Delta Z^T Y^T$:

$$\begin{aligned} \Delta Z^T Y^T &= \left(\sum_{k=1}^{d} \Delta_{ik} Z_{jk} \right)_{i=1,j=1}^{i=n,j=d} Y^T \\ &= \left(\sum_{k_1=1}^{d} \sum_{k_2=1}^{d} \Delta_{ik_1} Z_{k_2 k_1} Y_{jk_2} \right)_{i,j=1}^{n}. \end{aligned}$$

Taking the trace, we find

$$\begin{aligned} \langle \Delta, YZ \rangle &= \sum_{i=1}^{n} \sum_{k_1=1}^{d} \sum_{k_2=1}^{d} \Delta_{ik_1} Y_{ik_2} Z_{k_2 k_1} \\ &= \sum_{i=1}^{n} \langle \Delta_i^T Y_i, Z \rangle. \end{aligned} \tag{19}$$

From equations (18) and (19) the claim then follows. □

Proof of theorem 8: Let $Y(\cdot)$ be a geodesic starting at $Y(0)$ along a tangent Δ in the horizontal space. By lemma 7 it suffices to show that

$$\sum_{i=1}^{n} \langle \dot{Y}_i^T(t) Y_i(t), Z \rangle = 0, \ \ \forall Z \in \mathbb{R}^{d \times d}, Z \text{ skew-symmetric.} \tag{20}$$

So let such Z be given. First of all note that

$$\langle X_1^T X_2, Z \rangle = -\langle X_2^T X_1, Z \rangle \tag{21}$$

for any X_1 and X_2 since Z is skew-symmetric. In particular

$$\langle X^T X, Z \rangle = 0, \ \ \text{for any } X. \tag{22}$$

Then we have from the geodesic formulas (5):

$$\begin{aligned} \dot{Y}_i^T(t) Y_i(t) &= \Big\{ -\|\Delta_i\| \sin\big(\|\Delta_i\| t\big)\ Y_i(0) + \cos\big(\|\Delta_i\| t\big)\ \Delta_i \Big\} \\ &\qquad \Big\{ \cos(\|\Delta_i\| t)\ Y_i^T(0) + \frac{1}{\|\Delta_i\|} \sin\big(\|\Delta_i\| t\big)\ \Delta_i^T \Big\} \\ &= \Big\{ A + B \Big\} \Big\{ C + D \Big\}, \text{ say.} \end{aligned}$$

When we take the inner product with Z then the terms $A \times C$ and $B \times D$ cancel because of the property of equation (22). We are left with:

$$\begin{aligned} \langle \dot{Y}_i^T(t) Y_i(t), Z \rangle &= -\sin^2\big(\|\Delta_i\| t\big) \langle Y_i(0) \Delta_i^T, Z \rangle + \cos^2\big(\|\Delta_i\| t\big) \langle \Delta_i Y_i^T(0), Z \rangle \\ &= \langle \Delta_i Y_i^T(0), Z \rangle, \end{aligned}$$

where in the second equality we changed the sign and transposed $Y_i(0)\Delta_i^T$ in the first term by virtue of equation (21). Subsequently the familiar equation $\sin^2 + \cos^2 = 1$ was applied. Since Δ is horizontal at $Y(0)$ by assumption, equation (20) follows, which was to be shown. □

Appendix 3: Quotient of manifolds

Definition 15 *Let M be a smooth manifold, and G a Lie group. A smooth action from the right of G on M is a smooth map $A : M \times G \to M$ denoted by $A(m, g) = mg^{-1}$ that satisfy following conditions*

- *If e is identity element of G, then for any $m \in M$*

$$me = e$$

- *If $g_1, g_2 \in G$, then for any $m \in M$ we have*

$$m(g_1 g_2)^{-1} = (m g_2^{-1}) g_1^{-1}$$

.

Definition 16 *Let Φ be a continuous mapping from a topological space U to topological space V. Then Φ is called* proper *if $\Phi^{-1}(K)$ is compact in U for every compact subset of V.*

A continuous action of a topological group G on a topological space M is said to be a proper action *if*

$$(23) \qquad (m, g) \mapsto (mg^{-1}, m) \quad \textit{is a proper mapping}: \quad M \times G \to M \times M$$

Theorem 17 *Let G be a Lie group, M a smooth manifold and A a smooth action of G on M that is proper and free. Then orbit space M/G has a unique structure of a smooth manifold of dimension equal to $\dim M - \dim G$ with following properties. If $\pi : M \to M/G$ is the canonical projection $m \mapsto mG$, then for every $b \in M/G$ there is an open neighborhood $S \subseteq M/G$ and a smooth diffeomorphism:*

$$\tau : m \mapsto (s(m), \chi(m)) : \pi^{-1}(S) \to S \times G,$$

such that, for every $m \in \pi^{-1}(S)$, $g \in G$:

$$\pi(m) = s(m) \qquad \textit{and} \qquad \tau(mg^{-1}) = (s(m), g\chi(m)).$$

The topology of M/G is equal to quotient topology