

Geometric means of structured matrices ^{*}

Dario A. Bini[†] Bruno Iannazzo[‡] Ben Jeuris[§]
Raf Vandebril[§]

December 20, 2012

Abstract

The geometric mean of positive definite matrices is usually identified with the Karcher mean which verifies all properties –inspired by the scalar case– a geometric mean is expected to satisfy. Unfortunately, the Karcher mean is typically not structure preserving, and destroys, e.g., Toeplitz and band structures, which emerge in many applications. For this reason, the Karcher mean is not always recommended for modeling averages of structured matrices. In this article a new definition of a geometric mean for structured matrices is introduced, its properties are outlined, algorithms for its computation, and numerical experiments are provided. In the Toeplitz case an existing mean based on the Kähler metric is analyzed for comparison.

1 Introduction

The wish to generalize the concept of a geometric mean to positive definite (*positive* for short) matrices and, on the other hand, the need to average quantities expressed by positive matrices in certain applications have led to the definition and the study of the Karcher mean [6, 7, 27].

Without describing all facets, one can consider the set of $n \times n$ positive matrices, denoted by \mathcal{P}_n , as a manifold [1], in particular, there is a smooth map from \mathcal{P}_n to \mathbb{R}^{n^2} . In each point of $X \in \mathcal{P}_n$ one can define the tangent space $T_X \mathcal{P}_n$ which can be identified with the space of Hermitian matrices. The Karcher mean can now be defined in terms of a Riemannian geometry defined on \mathcal{P}_n and induced by the inner product

$$g_A(X, Y) := \operatorname{tr}(A^{-1}XA^{-1}Y), \quad X, Y \in T_A \mathcal{P}_n \quad (1.1)$$

^{*}This work was partially supported by MIUR grant number 2002014121; by the Research Council KU Leuven, projects OT/11/055 (Spectral Properties of Perturbed Normal Matrices and their Applications), CoE EF/05/006 Optimization in Engineering (OPTEC); by the Fund for Scientific Research–Flanders (Belgium) project G034212N (Reestablishing Smoothness for Matrix Manifold Optimization via Resolution of Singularities); and by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office, Belgian Network DYSCO (Dynamical Systems, Control, and Optimization).

[†]Dipartimento di Matematica, Università di Pisa, Largo Bruno Pontecorvo 5, 56125 Pisa, Italy bini@dm.unipi.it

[‡]Dipartimento di Matematica e Informatica, Università di Perugia, Via Vanvitelli 1, 06123 Perugia, Italy bruno.iannazzo@dm.unipg.it

[§]Departement Computerwetenschappen, KU Leuven, Celestijnenlaan 200A, 3000 Leuven, Belgium {ben.jeuris,raf.vandebril}@cs.kuleuven.be

on the tangent space $T_A\mathcal{P}_n$. This inner product g_A makes \mathcal{P}_n a complete Riemannian manifold with non-positive curvature and yields the following distance between two matrices $X, Y \in \mathcal{P}_n$:

$$\delta(X, Y) = \left(\sum_{k=1}^n \log^2 \lambda_k \right)^{1/2}, \quad (1.2)$$

where $\lambda_1, \dots, \lambda_n$, are the eigenvalues of $X^{-1}Y$, which are positive numbers (for all the proofs see [6, Ch. 6]). The Karcher mean of a set of m positive matrices, $A_1, \dots, A_m \in \mathcal{P}_n$, is defined as the unique positive minimizer $G(A_1, \dots, A_m)$ of the function

$$f(X; A_1, \dots, A_m) := \sum_{j=1}^m \delta^2(X, A_j). \quad (1.3)$$

Since this mean minimizes the sum of squared intrinsic distances to each of the matrices A_j it is a barycenter of these matrices with respect to the aforementioned metric.

An important quality of the Karcher mean is that it verifies all the properties desired by a geometric mean, like the ten Ando-Li-Mathias (ALM) axioms [2]. For this reason, it is a viable tool in applications requiring some of these properties [5, 28]. A geometric mean should for instance be: permutation invariant, monotone, joint concave, and should satisfy the arithmetic-geometric-harmonic inequality (see [2] for the precise statements of the properties). In particular, one of the most characterizing properties of a geometric mean is its invariance under inversion:

$$G(A_1^{-1}, \dots, A_m^{-1}) = G(A_1, \dots, A_m)^{-1}. \quad (1.4)$$

Prior to having the proofs of all the properties of the Karcher mean, some of which very elusive [25], other definitions of a matrix geometric mean had been proposed [2, 8, 11, 29], even if nowadays there is large agreement in considering the Karcher mean as the “right” matrix geometric mean.

In certain applications, however, besides the positive definiteness, the data matrices have some further structure in the sense that they belong to some special subset \mathcal{S} , say a linear space. For instance, in the design and analysis of certain radar systems, the matrices to be averaged are correlation matrices, which are positive Toeplitz matrices [24]. In these cases, one would like the geometric mean to belong to the same class \mathcal{S} as the data. Unfortunately, the Karcher mean does not preserve many structures, in particular the Karcher mean of Toeplitz and/or band matrices is typically not of Toeplitz and/or band form anymore, as illustrated by the following simple example.

Example 1.1. Let \mathcal{S} be the set of tridiagonal Toeplitz matrices and choose $A_1, A_2 \in \mathcal{S}$ where $A_1 = I$, and $A_2 = \text{tridiag}(1, 2, 1)$ being the matrix with 2’s on the main, and 1’s appearing on sub- and superdiagonals. We have $A_1 A_2 = A_2 A_1$, thus the Karcher mean equals $(A_1 A_2)^{1/2}$. For $n = 3$ we get

$$(A_1 A_2)^{1/2} = \frac{\sqrt{2}}{4} \begin{bmatrix} \sqrt{2 + \sqrt{2}} + 2 & \sqrt{2}\sqrt{2 - \sqrt{2}} & \sqrt{2 + \sqrt{2}} - 2 \\ \sqrt{2}\sqrt{2 - \sqrt{2}} & \sqrt{2 + \sqrt{2}} & \sqrt{2}\sqrt{2 - \sqrt{2}} \\ \sqrt{2 + \sqrt{2}} - 2 & \sqrt{2}\sqrt{2 - \sqrt{2}} & \sqrt{2 + \sqrt{2}} + 2 \end{bmatrix} \quad (1.5)$$

which is neither tridiagonal nor Toeplitz.

In this paper we introduce the concept of a structured geometric mean of positive matrices in such a way that if $A_1, \dots, A_m \in \mathcal{S}$ also their mean belongs to \mathcal{S} . Given a subset \mathcal{S} of \mathcal{P}_n and matrices $A_1, \dots, A_m \in \mathcal{S}$, we say that $G \in \mathcal{S}$ is a *structured geometric mean* with respect to \mathcal{S} of A_1, \dots, A_m if the function $f(X; A_1, \dots, A_m)$ of (1.3) takes its minimum value over \mathcal{S} at G . The set of the structured geometric means of A_1, \dots, A_m with respect to \mathcal{S} is denoted by $G_{\mathcal{S}} = G_{\mathcal{S}}(A_1, \dots, A_m)$.

We show that if \mathcal{S} is closed then $G_{\mathcal{S}}$ is nonempty and the matrices $G \in G_{\mathcal{S}}$ satisfy most of the ALM axioms, in a suitably adjusted form, for instance, the invariance under inversion property (1.4) turns into

$$G_{\mathcal{S}}(A_1, \dots, A_m) = G_{\mathcal{S}^{-1}}(A_1^{-1}, \dots, A_m^{-1})^{-1},$$

where for a set $\mathcal{U} \subseteq \mathcal{P}_n$ we denote $\mathcal{U}^{-1} = \{X^{-1} : X \in \mathcal{U}\}$. That is, the inverse of any structured geometric mean of the matrices $A_1, \dots, A_m \in \mathcal{S}$ with respect to \mathcal{S} coincides with a structured mean of the inverses $A_1^{-1}, \dots, A_m^{-1}$ with respect to the set \mathcal{S}^{-1} where these inverses reside.

Moreover, we show that, in many interesting cases, structured geometric means can be characterized in terms of the positive solutions of a suitable vector equation and provide algorithms for their computation.

In the Toeplitz case we also consider a different approach, where the mean is defined as a barycenter for a suitable metric on the manifold [4]. We analyze this barycenter and its properties in detail, obtaining an explicit expression in the real case and a quick algorithm in the complex case.

The article is organized as follows. In Section 2 the cost function (1.3) is examined with special focus on the existence of the minimizer over a closed set. The structured matrix mean itself is the subject of study in Section 3, where the theoretical properties it should satisfy are examined. Section 4 proposes two algorithms for computing a structured mean G in a linear space together with their convergence analysis. For one algorithm, it is shown that the convergence speed is independent of the condition number of the mean and is faster when the condition numbers of the matrices $A_i^{-1/2} G A_i^{-1/2}$ are smaller, for $i = 1, \dots, n$. Because of its nature and its convergence properties, this algorithm can be viewed as the natural extension to the structured case of the Richardson-like algorithm introduced and analyzed in [10] for the computation of the Karcher mean of unstructured matrices. In Section 5, for Toeplitz matrices, a different structured matrix mean [4] as a barycenter is considered, and an algorithm for computing it is developed. Section 6 shows numerical experiments related to accuracy and speed for computing the structured matrix mean.

Here we recall some basic notation and properties that will be used in the rest of the paper. Given a matrix A , we define $\sigma(A)$ the spectrum of A , that is, the set of all the eigenvalues of A , and $\rho(A) = \max_{\lambda \in \sigma(A)} |\lambda|$ the spectral radius of A . Moreover we denote by $\|A\|_F := (\text{trace}(A^* A))^{1/2} = (\sum_{i,j} |a_{ij}|^2)^{1/2}$ the Euclidean (Frobenius) norm of A , and $\|A\|_2 = \rho(A^* A)^{1/2}$ the spectral norm. By A^* we denote the transposed conjugate of A . Recall that for a positive matrix A there exists a unique positive solution to the equation $X^2 = A$. This solution, denoted by $A^{1/2}$, is called the square root of A [6]. Given a matrix $A \in \mathbb{C}^{n \times n}$, we use the vec-operator to build $\text{vec}(A) \in \mathbb{C}^{n^2}$, a long vector obtained by stacking the columns of A . We will use the Kronecker product \otimes such that $A \otimes B$ is the

block matrix whose (i, j) th block is defined as $a_{ij}B$. The vec operator and the Kronecker product interplay in the following way [20]

$$\text{vec}(ABC) = (C^T \otimes A) \text{vec}(B). \quad (1.6)$$

Finally, we recall a natural partial order in \mathcal{P}_n that will be used in the following: let A and B be positive, we write $A \geq B$ if the matrix $A - B$ is semidefinite positive.

2 Existence of structured geometric means

In this section the existence of a structured geometric mean and its relation to the classical Karcher mean is studied. First some necessities are repeated.

2.1 Uniqueness of the Karcher mean for positive matrices

The Riemannian geometry on \mathcal{P}_n given by the inner product (1.1) turns out to be complete and a parametrization of the geodesic joining two positive matrices A and B is known to be [6]

$$A\#_t B = A^{1/2}(A^{-1/2}BA^{-1/2})^t A^{1/2} = A(A^{-1}B)^t, \quad t \in [0, 1], \quad (2.1)$$

where the midpoint $A\#_{1/2}B$ coincides with the geometric mean of two matrices [21].

Given a set of matrices $A_1, \dots, A_m \in \mathcal{P}_n$, the function $f(X) = f(X; A_1, \dots, A_m)$ in (1.3) is strictly geodesically convex, which means that for any two different matrices $X, Y \in \mathcal{P}_n$, we have

$$f(X\#_t Y) < (1-t)f(X) + tf(Y) \quad 0 < t < 1. \quad (2.2)$$

This property follows from [6, Exercise 6.1.13], where it is stated that for $m = 1$ the function $f(X)$ is strictly geodesically convex. The case $m > 1$ follows by summing up the m inequalities obtained by applying (2.2) to the functions $f(X) = f(X; A_i)$, for $i = 1, \dots, m$, respectively.

Geodesical convexity is a key ingredient for the proof of the existence of a unique minimizer of f over \mathcal{P}_n given in [6, Ch. 6]. A different proof is obtained using the fact that \mathcal{P}_n , with the inner product (1.1), forms a Cartan-Hadamard manifold [14, 26], which is a Riemannian manifold, complete, simply connected and has non-positive sectional curvature everywhere. On such a Cartan-Hadamard manifolds the Karcher mean (the so-called center-of-mass) is unique [15, 23].

The notion of geodesical convexity in \mathcal{P}_n is different from the customary convexity in the Euclidean space where one requires that

$$f((1-t)X + tY) \leq (1-t)f(X) + tf(Y), \quad t \in [0, 1].$$

In fact, the function f is not convex in the traditional sense as the following example shows.

Example 2.1. Consider the set made of the unique matrix $A = 1$, and $x, y \in \mathbb{R}_+^* = \mathcal{P}_1$. We have $f(x) = \delta^2(x, A) = \log^2(x)$ which is not convex. On the

other hand the function $\log^2(x)$ is strictly geodesically convex and this can be shown by an elementary argument: in fact, it is continuous and

$$\begin{aligned}\delta^2(\sqrt{xy}, 1) &= \log^2(\sqrt{xy}) = \frac{1}{4} (\log^2 x + \log^2 y + 2 \log x \log y) \\ &= \frac{1}{2} (\log^2 x + \log^2 y) - \frac{1}{4} (\log x - \log y)^2 \\ &< \frac{1}{2} (\log^2 x + \log^2 y) = \frac{1}{2} (\delta^2(x, 1) + \delta^2(y, 1)).\end{aligned}$$

Iterative selection of midpoints, by using midpoints and a continuity argument completes the proof.

Since f is strictly geodesically convex, it can be proved that it has a unique minimizer over any geodesically convex subset \mathcal{S} of \mathcal{P}_n , where we say that a subset $\mathcal{S} \subseteq \mathcal{P}_n$ is geodesically convex if for any $X, Y \in \mathcal{S}$, the entire geodesic $X \#_t Y$, $t \in [0, 1]$ belongs to \mathcal{S} . In fact, by absurd, if X_1 and X_2 were two different matrices in \mathcal{S} where f takes its minimum, then from (2.2) it would follow that $f(X_1 \#_t X_2) < f(X_1) = f(X_2)$ for any $0 < t < 1$ which contradicts the assumption.

2.2 Existence of structured geometric means on a closed set

For a generic closed subset \mathcal{U} of \mathcal{P}_n which is not necessarily geodesically convex, we can prove the existence of a minimum point by using the fact that $f(X)$ is continuous. In order to prove this, we first give a couple of preliminary results.

Lemma 2.2. *Let $A, X, Y \in \mathcal{P}_n$ be such that $Y = A^{-1/2} X A^{-1/2}$. Then for any operator norm it holds that*

$$\|Y\| \geq \|X\| / \|A^{1/2}\|^2, \quad \|Y^{-1}\| \geq \|X^{-1}\| / \|A^{-1/2}\|^2.$$

Proof. The condition $Y = A^{-1/2} X A^{-1/2}$ can be rewritten as $X = A^{1/2} Y A^{1/2}$. Taking norms yields $\|X\| \leq \|A^{1/2}\|^2 \|Y\|$ from which the first inequality follows. The second inequality holds similarly starting from $Y^{-1} = A^{1/2} X^{-1} A^{1/2}$. \square

Lemma 2.3. *For the function $\delta^2(X, A)$ it holds that*

$$\delta^2(X, A) \geq \log^2 s$$

where $s = \max \{ \rho(A^{-1/2} X A^{-1/2}), \rho(A^{1/2} X^{-1} A^{1/2}) \}$.

Proof. It follows from the equation

$$\delta^2(X, A) = \sum_i \log^2 \lambda_i \left(A^{-1/2} X A^{-1/2} \right).$$

Since all terms are positive, $\sum_i \log^2 \lambda_i (A^{-1/2} X A^{-1/2})$ is greater than any single term in the summation, in particular those given by the extreme eigenvalues of $A^{-1/2} X A^{-1/2}$, that is, $\rho(A^{-1/2} X A^{-1/2})$ and $1/\rho(A^{1/2} X^{-1} A^{1/2})$. \square

We are now ready to prove the main result of this section.

Theorem 2.4. *Let $\mathcal{U} \subseteq \mathcal{P}_n$ be a closed subset. Then for any $A_1, \dots, A_m \in \mathcal{P}_n$ the function $f(X) = f(X; A_1, \dots, A_m)$ has a minimum in \mathcal{U} .*

Proof. If \mathcal{U} is bounded, then it is compact and the continuous function $f(X)$ has a minimum in it, so we may assume that \mathcal{U} is unbounded. Let $t > 0$ and $\mathcal{A}_t = \{X \in \mathcal{P}_n : \|X\|_2 \leq t, \|X^{-1}\|_2 \leq t\}$, such that \mathcal{A}_t is closed and bounded. We claim that there exists a sufficiently large value t such that outside the set $\mathcal{U} \cap \mathcal{A}_t$ the function $f(X)$ takes values larger than $\gamma = \inf_{X \in \mathcal{U}} f(X)$. In this way, the set where we minimize the function can be restricted to $\mathcal{U} \cap \mathcal{A}_t$ which is compact and hence $f(X)$ takes its minimum over it. For simplicity, we prove the existence of t for $m = 1$. The case $m > 1$ can be obtained by using the same arguments.

Combining Lemma 2.2 with $\|\cdot\| = \|\cdot\|_2$, Lemma 2.3, and using the properties of the spectral norm, we find that there exist positive constants α and β (depending on A) such that

$$\delta^2(X, A) \geq \max \{ \log^2(\alpha \|X\|_2), \log^2(\beta \|X^{-1}\|_2) \} \quad (2.3)$$

for any $X \in \mathcal{P}_n$. Choosing t sufficiently large in such a way that $\log^2(\alpha t), \log^2(\beta t) > \gamma$, it follows from (2.3) that $\delta^2(X, A) > \gamma$ for any X having $\|X\|_2 > t$ or $\|X^{-1}\|_2 > t$. This completes the proof of the existence of a minimum of $f(X, A)$. Considering the summation in (1.3) this generalizes to an arbitrary $f(X)$. \square

In general, uniqueness of the point where $f(X)$ takes its minimum cannot be guaranteed. For instance, if both A and A^{-1} belong to \mathcal{U} while $I = A \#_{1/2} A^{-1}$ does not, then the function $f_1(X) := \delta^2(X, A) + \delta^2(X, A^{-1})$ reaches its minimum at a point $G \neq I \in \mathcal{U}$. Clearly, $f_1(G^{-1}) = f_1(G)$ and if $G^{-1} \neq G$ belongs to \mathcal{U} , then we have at least two distinct points of minimum. A more concrete example is the following.

Example 2.5. Consider the 2×2 matrices $A = I$ and $B = \begin{bmatrix} a & 0 \\ 0 & a^{-1} \end{bmatrix}$, where $a > 1$. Define the segment $\mathcal{U} = \{G(t) = A + t(B - A), t \in [0, 1]\}$ which is closed and convex, but not geodesically convex. The function $f(t) = \delta^2(G(t), A) + \delta^2(G(t), B)$ takes the form $f(t) = \log^2((1-t)/a + t) + \log^2(a(1-t) + t) + \log^2((1-t) + t/a) + \log^2((1-t) + at)$ and is symmetric with respect to $t = 1/2$. For $a = 200$ the function has the graph shown in Figure 1 with a local maximum at $t = 1/2$ and two global minima close to the edges of the segment.

3 A theoretical exploration of the structured geometric mean

In this section we discuss the relation between the structured and generic geometric mean, together with the adaptation of the generic properties to the structured setting. We will discuss just the real case, so in this section, the set \mathcal{P}_n stands for the manifold of real positive definite matrices whose tangent space is the set of real and symmetric matrices.




Figure 1: Graph of $f(t) = \delta^2(G(t), A) + \delta^2(G(t), B)$ for $G(t) = A + t(B - A)$ with $A = I$ and $B = \text{diag}(200, 1/200)$.

3.1 The geometric and structured geometric mean relation

The properties shown in Section 2 imply that a structured geometric mean with respect to \mathcal{U} , as defined in the introduction, always exists for any closed subset \mathcal{U} of \mathcal{P}_n . In particular, this holds in the cases where $\mathcal{U} = \mathcal{S} \cap \mathcal{P}_n$ for any linear space \mathcal{S} of matrices \mathcal{S} and also for $\mathcal{U}^{-1} := \mathcal{S}^{-1} \cap \mathcal{P}_n$ where $\mathcal{S}^{-1} = \{A^{-1} : A \in \mathcal{S}, \det A \neq 0\}$. This captures a wide class of interesting structures emerging in applications, e.g., Toeplitz and band matrices, as well as their inverses. For simplicity we will restrict our analysis in the remainder of the article to the real case.

More general structures are given in terms of a parametrization $\sigma(t) : \mathcal{V} \rightarrow \mathbb{R}^{n \times n}$, with σ a differentiable function defined in the open subset \mathcal{V} of \mathbb{R}^q , which we will call the *parameter space*. The set $\mathcal{T} = \sigma(\mathbb{R}^q)$ is the structure determined by σ . If σ is linear and $\mathcal{V} = \mathbb{R}^q$, then \mathcal{T} is a linear space. Examples of sets \mathcal{T} of interest which generally do not form a linear space are the set of matrices with a given displacement rank [9], the set of semiseparable [33], and quasiseparable matrices [17]. For an $n \times n$ symmetric Toeplitz matrix, a possible parametrization is given by

$$\sigma(t) = \sigma([t_0, t_1, \dots, t_{n-1}]) = \begin{bmatrix} t_0 & t_1 & \dots & t_{n-1} \\ t_1 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & t_1 \\ t_{n-1} & \dots & t_1 & t_0 \end{bmatrix}. \quad (3.1)$$

For a band matrix, one can, e.g., just store the nonzero-elements in a long vector and map them onto their exact locations. In the following, given a closed set \mathcal{T} we let $\mathcal{U} = \mathcal{T} \cap \mathcal{P}_n$.

In Example 2.5 we illustrated that the minimum of the cost function re-

stricted to a closed subset $\mathcal{U} \subseteq \mathcal{P}_n$ is not necessarily unique. For this reason, we consider the structured geometric mean $G_{\mathcal{U}} = G_{\mathcal{U}}(A_1, \dots, A_m)$ of $A_1, \dots, A_m \in \mathcal{U}$, as the set of matrices in \mathcal{U} where the function $f(X)$ attains its minimum. Formally speaking, for $A_1, \dots, A_m \in \mathcal{U}$, let $g \in \mathbb{R}^q$ be such that $\widehat{G} = \sigma(g) \in G_{\mathcal{U}}(A_1, \dots, A_m)$, then

$$f(\sigma(g); A_1, \dots, A_m) = \min_{t \in \mathbb{R}^q} f(\sigma(t); A_1, \dots, A_m).$$

Since $\mathcal{U} \subseteq \mathcal{P}_n$, the minimum over \mathcal{P}_n is less than or equal to the minimum over \mathcal{U} . In general it will often happen that $\widehat{G} \neq G(A_1, \dots, A_m)$ like in (1.5).

3.2 Properties of the geometric mean conveyed to the structured mean setting

Some desired properties from a matrix geometric mean were stated by Ando, Li and Mathias in [2], of which the most noticeable are enlisted here.

Consistency with scalars If A_1, \dots, A_m commute, then

$$G(A_1, \dots, A_m) = (A_1 \cdots A_m)^{1/m}.$$

Permutation invariance For any permutation π of $\{1, \dots, k\}$, it holds that

$$G(A_1, \dots, A_m) = G(A_{\pi(1)}, \dots, A_{\pi(m)}).$$

Joint homogeneity

$$G(\alpha_1 A_1, \alpha_2 A_2, \dots, \alpha_m A_m) = (\alpha_1 \cdots \alpha_m)^{1/m} G(A_1, \dots, A_m).$$

Monotonicity If $A_i \geq A'_i$, for $i = 1, \dots, k$, then

$$G(A_1, \dots, A_m) \geq G(A'_1, \dots, A'_m).$$

Invariance under congruence For any nonsingular M ,

$$G(M^* A_1 M, \dots, M^* A_m M) = M^* G(A_1, \dots, A_m) M.$$

Invariance under inversion

$$G(A_1, \dots, A_m)^{-1} = G(A_1^{-1}, \dots, A_m^{-1}).$$

Arithmetic-geometric-harmonic mean inequality

$$\frac{1}{m}(A_1 + \cdots + A_m) \geq G(A_1, \dots, A_m) \geq m(A_1^{-1} + \cdots + A_m^{-1})^{-1}.$$

Yet another property naturally desired from a geometric mean, but not required in the list of Ando, Li and Mathias, is the *repetition invariance*, that is, for any set of positive matrices $A_1, \dots, A_m \in \mathcal{P}_n$, it holds that

$$G(A_1, \dots, A_m, A_1, \dots, A_m) = G(A_1, \dots, A_m). \quad (3.2)$$

Now, we consider the properties of the structured geometric mean. Some properties such as the permutation invariance trivially hold, others should be restated. In fact, in the generic case the structures we consider, are not invariant under inversion nor under congruence. That is because if $A \in \mathcal{U}$ then it need not necessarily hold that $A^{-1} \in \mathcal{U}$ nor $M^*AM \in \mathcal{U}$.

We start with the invariance under inversion as this is one of the most characterizing properties of the geometric mean. To this end we consider the set $\mathcal{T}^{-1} = \{T^{-1} : T \in \mathcal{T}, \det T \neq 0\}$ parametrized with the function $\sigma_{-1}(t) := \sigma(t)^{-1}$. Clearly, the intersection \mathcal{U} of \mathcal{T} with \mathcal{P}_n yields always invertible matrices, so that $\mathcal{T}^{-1} \cap \mathcal{P}_n = \mathcal{U}^{-1}$.

According to our definition, the structured geometric mean of $A_1^{-1}, \dots, A_m^{-1} \in \mathcal{U}^{-1}$ is given by the set $G_{\mathcal{U}^{-1}}(A_1^{-1}, \dots, A_m^{-1})$. For any $\tilde{G} \in G_{\mathcal{U}^{-1}}$, we have $\tilde{G} = \sigma(\tilde{g})^{-1}$ such that

$$f(\sigma(\tilde{g})^{-1}; A_1^{-1}, \dots, A_m^{-1}) = \min_{t \in \mathbb{R}^q} f(\sigma(t)^{-1}; A_1^{-1}, \dots, A_m^{-1}).$$

Since $\delta(A, B) = \delta(A^{-1}, B^{-1})$, one gets $f(X; A_1, \dots, A_m) = f(X^{-1}; A_1^{-1}, \dots, A_m^{-1})$ so that

$$f(\sigma(\tilde{g}); A_1, \dots, A_m) = \min_{t \in \mathbb{R}^q} f(\sigma(t); A_1, \dots, A_m)$$

and thus $\tilde{G}^{-1} \in G_{\mathcal{U}}(A_1, \dots, A_m)$. Since \tilde{G} was chosen arbitrarily, and since \mathcal{U} can be interchanged with \mathcal{U}^{-1} , we have the analogous of the invariance under inversion for the structured geometric mean:

$$G_{\mathcal{U}}(A_1, \dots, A_m)^{-1} = G_{\mathcal{U}^{-1}}(A_1^{-1}, \dots, A_m^{-1}). \quad (3.3)$$

In a similar manner we can restate the invariance under congruence in a structured style by defining, for any nonsingular M the set $\mathcal{U}_M := M^*\mathcal{U}M = \{M^*TM : T \in \mathcal{U}\}$. The invariance under congruence is then understood as

$$G_{\mathcal{U}_M}(M^*A_1M, \dots, M^*A_mM) = M^*G_{\mathcal{U}}(A_1, \dots, A_m)M$$

which holds true.

Joint homogeneity, in order to be defined, requires that the set \mathcal{T} satisfies the following property:

$$A \in \mathcal{T} \Rightarrow \alpha A \in \mathcal{T}$$

for any scalar $\alpha > 0$. This property clearly holds if \mathcal{T} is a linear space or the set formed by the inverses of the nonsingular matrices of a linear space. For these sets, the joint homogeneity holds.

Repetition invariance holds true as well by (1.3), since

$$f(X; A_1, \dots, A_m, A_1, \dots, A_m) = 2f(X; A_1, \dots, A_m),$$

so the minimizers (over a subset) of the functions $f(X; A_1, \dots, A_m, A_1, \dots, A_m)$ and $f(X; A_1, \dots, A_m)$ are the same.

Regarding the remaining properties, we observe that the consistency with scalars is not verified, as Example 1.1 shows. Nevertheless, weaker consistency properties hold, as idempotency, namely $G_{\mathcal{U}}(A, A, \dots, A) = A$ for each structure \mathcal{U} and $A \in \mathcal{U}$. Moreover, if the set \mathcal{U} is geodesically convex then

$$G_{\mathcal{U}}(A_1, \dots, A_m) = G(A_1, \dots, A_m),$$

so the geometric and structured geometric mean coincide. An interesting case of geodesically convex set is given by $\mathcal{U} = \mathcal{T} \cap \mathcal{P}_n$, when \mathcal{T} is an algebra, i.e., a linear space closed under multiplication and inversion.

Finally, the properties related to the ordering of positive matrices such as monotonicity are not true as shown by the following numerical example.

Example 3.1. We consider the four Toeplitz matrices

$$T_1 = \begin{bmatrix} 1 & 1/2 & 1/2 \\ 1/2 & 1 & 1/2 \\ 1/2 & 1/2 & 1 \end{bmatrix}, \quad T_2 = T_1, \quad T_3 = \begin{bmatrix} 3/4 & 1/2 & 0 \\ 1/2 & 3/4 & 1/2 \\ 0 & 1/2 & 3/4 \end{bmatrix}, \quad S = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix},$$

and, using the algorithms presented in the next sections, we compute a structured geometric mean G_ε of the three matrices T_1, T_2 and $T_3 + \varepsilon S$ for various $\varepsilon \geq 0$. The norm of $G_\varepsilon - G_0$ becomes small as ε tends to 0 and we observe that $G_\varepsilon - G_0$ is not positive (semi)definite, while $T_3 + \varepsilon S \geq T_3$. This gives numerical evidence of the lack of monotonicity of a structured geometric mean. On the other hand, computing the arithmetic mean A of T_1, T_2 and T_3 , one observes also that the expected inequality $A \geq G_0$ does not hold in this case.

3.3 The structured mean as solution(s) of a vector equation

We start from the Karcher mean which is obtained as the unique solution in \mathcal{P}_n of the matrix equation

$$\sum_{i=1}^m \log(XA_i^{-1}) = 0. \quad (3.4)$$

Equation (3.4) is obtained using the fact that f is differentiable and has a minimum at the Karcher mean. Thus the Karcher mean verifies the condition $\nabla f_X = 0$, where $\nabla f_X = 2X^{-1} \sum_{i=1}^m \log(XA_i^{-1})$ denotes the (Euclidean) gradient of f with respect to X (see [27]). We remark already that a different metric will be studied in Section 4.3.

In the general case, the restriction of f to a structure given by $\sigma(t)$ is investigated. For any minimum g (with corresponding $\sigma(g)$) not located at the boundary of the parameter space, the gradient $\nabla(f \circ \sigma)_t$ of the function with respect to t must be zero, so we are interested in the solutions of the vector equation $\nabla(f \circ \sigma)_t = 0$.

From the chain rule of derivation, one obtains that

$$\nabla(f \circ \sigma)_t = \left(\sum_{i,j} \frac{\partial f(\sigma(t))}{\partial x_{i,j}} \frac{d\sigma_{i,j}(t)}{dt_s} \right)_{s=1,\dots,q} = 0$$

which leads to the vector equation

$$\sum_{i,j} (\Gamma(\sigma(t)))_{i,j} \frac{d\sigma_{i,j}(t)}{dt_s} = 0, \quad s = 1, \dots, q, \quad (3.5)$$

where $\Gamma(X) := \frac{1}{2} \nabla f_X$.

In the case where \mathcal{T} is a linear space, the function $\sigma(t)$ is linear and can be written in matrix form as

$$\text{vec}(\sigma(t)) = Ut, \quad U \in \mathbb{R}^{n^2 \times q},$$

so that equation (3.5) turns into

$$U^T \text{vec}(\Gamma(\sigma(t))) = 0, \quad \Gamma(X) = X^{-1} \sum_{i=1}^m \log(XA_i^{-1}). \quad (3.6)$$

If $\sigma(t)$ is chosen to be orthogonal, i.e. such that $U^T U = I$, then U^T coincides with the Moore-Penrose inverse of U .

When \mathcal{T} denotes the set of symmetric Toeplitz matrices, the parametrization (3.1) leads to a matrix U having orthogonal columns. In fact one has $U^T U = D = \text{diag}(n, 2(n-1), 2(n-2), \dots, 2)$. In particular, for $n = 3$ one has

$$U^T = \left[\begin{array}{ccc|ccc} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \end{array} \right].$$

For \mathcal{T} being the set of symmetric tridiagonal matrices the parametrization

$$\sigma(t) = \left[\begin{array}{cccccc} t_1 & t_{n+1} & & & & \\ t_{n+1} & t_2 & t_{n+2} & & & \\ & \ddots & \ddots & \ddots & & \\ & & t_{2n-2} & t_{n-1} & t_{2n-1} & \\ & & & t_{2n-1} & t_n & \end{array} \right]$$

also leads to a matrix U having orthogonal columns. Moreover, $U^T U = \text{diag}(I_n, 2I_{n-1})$. For $n = 3$, e.g., one has

$$U^T = \left[\begin{array}{ccc|ccc} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ \hline 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \end{array} \right].$$

4 Algorithms for structured geometric means in the linear case

We will give two algorithms for computing structured geometric means when they are characterized in terms of the solutions g of a vector equation, as, for instance, in the linear case.

We first provide a general definition of a class of algorithms based on preconditioned functional iteration, then we specialize to two algorithms given by two different preconditioners.

The first, provided in Section 4.2 is derived by relying on the projection of the gradient with respect to the Euclidean scalar product. The second, presented in Section 4.3, is obtained through projection with respect to the Riemannian metric of \mathcal{P}_n described in Section 1.

4.1 A preconditioned functional iteration and its convergence

Throughout this section we assume that $A_1, \dots, A_m \in \mathcal{U}$, where $\mathcal{U} = \mathcal{T} \cap \mathcal{P}_n$ and \mathcal{T} is a linear space with a parametrization $\sigma(t)$ such that $\text{vec}(\sigma(t)) = Ut$, and $D = U^T U$.

The structured geometric mean $G_{\mathcal{U}}$ is the set of minimizers of the function $f(X; A_1, \dots, A_m)$ over \mathcal{U} . These minimizers must be sought among the stationary points of the function f , that is, among the solutions to the vector equation (3.6).

Therefore, a way to design algorithms for computing structured means $G_{\mathcal{U}}$ is to apply numerical techniques to solve the vector equation (3.6). We consider a preconditioned Richardson-like iteration constructed in the spirit of [12]. Let $V(X)$ be a nonsingular and sufficiently differentiable matrix function and define

$$\begin{aligned} \varphi(t) &= t - \theta S(t), \quad S(t) = V(\sigma(t))^{-1} U^T \text{vec}(\Gamma(\sigma(t))), \\ t_{\nu+1} &= \varphi(t_{\nu}), \quad \nu = 0, 1, \dots, \end{aligned} \quad (4.1)$$

where θ is a parameter introduced to enhance convergence, $V(\sigma(t))$ is a preconditioner and t_0 is a given vector such that $\sigma(t_0)$ is positive. Observe that the fixed points of $\varphi(t)$ are the solutions of the vector equation (3.6) and if convergent, the sequence t_{ν} converges to a solution of the vector equation (3.6).

In the following, given a matrix function $f(X)$, where $X = (x_{i,j})$ and $f(X)$ are $n \times n$ matrices, we denote by $J_f(G)$ the $n^2 \times n^2$ Jacobian matrix of $\text{vec}(f(X))$ with respect to the variable $\text{vec}(X)$ computed at $X = G$, similarly we denote $J_{f \circ \sigma}(t_G)$ the $n^2 \times q$ Jacobian of the composed function $\text{vec}(f(\sigma(t)))$ with respect to the variables (t_1, \dots, t_q) at $t = t_G$. In this notation, the function in the subscript as well as the variable between parentheses specify if the derivatives are taken w.r.t. the matrix variable X or the vector variable t .

Observe that if $V(\sigma(t))$ is chosen as the Jacobian of $U^T \text{vec}(\Gamma(\sigma(t)))$, then (4.1) coincides with Newton's iteration.

If t_G is a solution of (3.6) and if t_{ν} is sufficiently near to t_G , then

$$t_{\nu+1} - t_G = J_{\varphi}(t_G)(t_{\nu} - t_G) + O(\|t_{\nu} - t_G\|^2),$$

so that in order to study the local convergence of this sequence it is sufficient to estimate the spectral radius ρ or any induced norm of $J_{\varphi}(t_G)$ and determine θ in such a way that $\rho(J_{\varphi}(t_G)) < 1$. Notice that the Jacobian of $\varphi(t)$ at $t = t_G$ is given by $I - \theta K$ where $K = J_S(t_G)$ is the Jacobian of $S(t)$ at $t = t_G$. Therefore, if we can find a preconditioner $V(t)$ such that K has real positive eigenvalues with minimum and maximum eigenvalues κ_{\min} and κ_{\max} respectively, then the choice $\theta = 2/(\kappa_{\min} + \kappa_{\max})$ insures local convergence and provides the minimum spectral radius of $J_{\varphi}(t_G)$ given by

$$\rho(J_{\varphi}(t_G)) = \frac{\kappa_{\max} - \kappa_{\min}}{\kappa_{\max} + \kappa_{\min}} = \frac{\mu - 1}{\mu + 1} < 1, \quad \mu = \kappa_{\max}/\kappa_{\min}.$$

Moreover, any values $\hat{\kappa}_{\min} \leq \hat{\kappa}_{\max}$ such that $\hat{\kappa}_{\min} \leq \kappa_{\min} \leq \kappa_{\max} \leq \hat{\kappa}_{\max}$ can be used instead of κ_{\min} and κ_{\max} to determine a value $\hat{\theta} = 2/(\hat{\kappa}_{\min} + \hat{\kappa}_{\max})$ which insures convergence. Also notice that the closer μ to 1 the faster the convergence of the iteration.

Therefore our goal is to perform a spectral analysis of K and to find an upper bound to the ratio $\mu = \kappa_{\max}/\kappa_{\min}$, assuming that all the eigenvalues of K are real positive. From the composition rule of derivatives one finds that

$$K = V(\sigma(t_G))^{-1}U^T J_\Gamma(G)U + J_{V(\sigma(t_G))^{-1}}(\sigma(t_G))U^T \text{vec}(\Gamma(\sigma(t_G)))$$

and since the $U^T \text{vec}(\Gamma(\sigma(t_G))) = 0$ then

$$K = V(\sigma(t_G))^{-1}U^T J_\Gamma(G)U. \quad (4.2)$$

To evaluate $J_\Gamma(G)$, we recall that $\Gamma(X) = \sum_{i=1}^m X^{-1} \log(XA_i^{-1})$ so that it is sufficient to determine the formal expression of $J_\psi(G)$ for $\psi(G, A) = G^{-1} \log(GA^{-1})$ for a generic A and then write $J_\Gamma(G) = \sum_{i=1}^m J_{\psi(G, A_i)}(G)$. In order to evaluate $J_\psi(G)$, we rely on the definition of Fréchet derivative of a matrix function $f(X)$ at X in the direction E

$$Df_X[E] = \lim_{t \rightarrow 0} \frac{f(X + tE) - f(X)}{t} = \left. \frac{d}{dt} \right|_{t=0} f(X + tE).$$

In fact, the $n^2 \times n^2$ Jacobian matrix $J_f(X)$ of the vector function $\text{vec} \circ f \circ \text{vec}^{-1}$ at $\text{vec}(X)$, is related to the Fréchet derivative by the following equation

$$\text{vec}(Df_X[E]) = J_f(X) \text{vec}(E). \quad (4.3)$$

We recall also the following properties of the Fréchet derivative [19] where f, g are given matrix functions and $\varphi(X) = X^{-1}$

$$\begin{aligned} D(fg)_X[E] &= Df_X[E]g(X) + f(X)Dg_X[E], & \text{product rule,} \\ D(f \circ g)_X[E] &= Df_{g(X)}[Dg_X[E]], & \text{chain rule,} \\ D\varphi_X[E] &= -X^{-1}EX^{-1}, & \text{inversion.} \end{aligned} \quad (4.4)$$

For the derivative of the exponential function we have (see [19, Eq. 10.17a])

$$J_{\exp}(Y) = (I \otimes \exp Y) \beta(Y^T \otimes I - I \otimes Y), \quad \beta(z) = (e^z - 1)/z.$$

Therefore, since $J_{\log}(X) = J_{\exp}(Y)^{-1}$ for $Y = \log X$, we find that

$$J_{\log}(X) = \gamma(\log(X^T) \otimes I - I \otimes \log X)(I \otimes X^{-1}), \quad \gamma(z) = z/(e^z - 1). \quad (4.5)$$

We are now ready to provide an explicit expression of the Fréchet derivative of the function $\psi(X, A) = X^{-1} \log(XA^{-1})$ and of the Jacobian $J_{\psi(X, A)}(X)$.

Lemma 4.1. *Let $\psi(X) = X^{-1} \log(XA^{-1})$. Assume that A, X are positive. For the matrix $J_\psi(X)$ such that $\text{vec}(D\psi_X[E]) = J_\psi(X) \text{vec}(E)$ it holds that*

$$\begin{aligned} J_\psi(X) &= -X^{-1} \log(XA^{-1}) \otimes X^{-1} + (A^{-1} \otimes X^{-1})\gamma(W)(I \otimes AX^{-1}), \\ W &= \log(XA^{-1}) \otimes I - I \otimes \log(XA^{-1}), \end{aligned}$$

with $\gamma(z) = z/(e^z - 1)$.

Proof. Since $h(X) := \log(XA^{-1})$ is the composition of $f(X) = \log(X)$ and $g(X) = XA^{-1}$, we get by (4.4)

$$Dh_X[E] = D\log_{XA^{-1}}[EA^{-1}].$$

As $\psi(X)$ is the product of $f(X) = X^{-1}$ and $h(X)$, (4.4) gives us

$$D\psi_X[E] = -X^{-1}EX^{-1}\log(XA^{-1}) + X^{-1}Dh_X[E].$$

Combining the latter two equations yields

$$D\psi_X[E] = -X^{-1}EX^{-1}\log(XA^{-1}) + X^{-1}D\log_{XA^{-1}}[EA^{-1}].$$

By using (4.3) and (1.6) we find that the matrix $J_\psi(X)$ representing $D\psi_X$ is given by

$$J_\psi(X) = -(X^{-1}\log(XA^{-1}))^T \otimes X^{-1} + (I \otimes X^{-1}) J_{\log}(XA^{-1})(A^{-T} \otimes I).$$

Replacing (4.5) in the equation above and using the fact that $A = A^T, X = X^T$ yields

$$\begin{aligned} J_\psi(X) &= -\log(A^{-1}X)X^{-1} \otimes X^{-1} \\ &\quad + (I \otimes X^{-1})\gamma(\log(A^{-1}X) \otimes I - I \otimes \log(XA^{-1}))(A^{-1} \otimes AX^{-1}). \end{aligned}$$

Using the fact that $W\log(V)W^{-1} = \log(WVW^{-1})$, the first term can be written as $-X^{-1}\log(XA^{-1}) \otimes X^{-1}$. The second term can be written as $(I \otimes X^{-1})(A^{-1} \otimes I)\gamma(\log(XA^{-1}) \otimes I - I \otimes \log(XA^{-1}))(I \otimes AX^{-1})$, which completes the proof. \square

Recall that $\Gamma(X) = \sum_{i=1}^m \psi(X, A_i)$ and $G^{-1} \sum_{i=1}^m \log(GA_i^{-1}) = 0$, for $G = \sigma(t_G)$. Then by Lemma 4.1, we obtain the following formula for the Jacobian $J_\Gamma(\sigma(t))$:

$$\begin{aligned} J_\Gamma(G) &= (I \otimes G^{-1})H(I \otimes G^{-1}), \quad H = \sum_{i=1}^m H_i \\ H_i &= (A_i^{-1} \otimes I)\gamma(\log(GA_i^{-1}) \otimes I - I \otimes \log(GA_i^{-1}))(I \otimes A_i). \end{aligned}$$

Moreover, by using the properties of the Kronecker product and the fact that $\log(GA^{-1}) = A^{1/2}\log(A^{-1/2}GA^{-1/2})A^{-1/2}$, we can write

$$\begin{aligned} H_i &= (A_i^{-1/2} \otimes A_i^{1/2})\gamma(\log M_i \otimes I - I \otimes \log M_i)(A_i^{-1/2} \otimes A_i^{1/2}), \\ M_i &= A_i^{-1/2}GA_i^{-1/2}. \end{aligned}$$

From this expression it turns out that H_i is positive, and from (4.2) we find that $J_S(t_G)$ is the product of the matrix $V(\sigma(t_G))^{-1}$ and the positive matrix $U^T(I \otimes G^{-1}) \sum_{i=1}^m H_i(I \otimes G^{-1})U$.

Thus we may conclude with the following

Theorem 4.2. *The Jacobian K of the function $S(t)$ in (4.2) at $\sigma(t_G) = G$ is given by*

$$\begin{aligned} K &= V^{-1}U^T (I \otimes G^{-1}) H (I \otimes G^{-1}) U, \\ H &= \sum_{i=1}^m H_i, \quad H_i = (A_i^{-1/2} \otimes A_i^{1/2}) \gamma(\log M_i \otimes I - I \otimes \log M_i) (A_i^{-1/2} \otimes A_i^{1/2}), \\ M_i &= A_i^{-1/2} G A_i^{-1/2}, \\ \gamma(z) &= z/(e^z - 1). \end{aligned}$$

Moreover, the eigenvalues of K are the solutions of the equation

$$\det(\kappa V - U^T (I \otimes G^{-1}) H (I \otimes G^{-1}) U) = 0.$$

4.2 An elementary preconditioner

The simplest choice for the preconditioner $V(t)$ in (4.1) is $V(t) = U^T U = D$. This corresponds to projecting the gradient of the function $f(X, A_1, \dots, A_p)$ on the set \mathcal{U} according to the Euclidean scalar product. The problem $\det(\kappa I - K) = 0$ turns into the generalized q -dimensional symmetric eigenvalue problem

$$\det(U^T (\kappa I - (I \otimes G^{-1}) H (I \otimes G^{-1})) U) = 0.$$

This problem is the projection on the space spanned by the columns of U of the problem $\det(\nu I - (I \otimes G^{-1}) H (I \otimes G^{-1})) = 0$ which has real positive solutions.

Now we recall the following result, valid for general positive matrices A, B , which relates the generalized eigenvalues of the pair (A, B) to the ones of the projected pair $(U^T A U, U^T B U)$.

Lemma 4.3. *Let A, B be $n \times n$ positive matrices and U an $n \times m$ matrix, then the generalized eigenvalues of the pair $(U^T A U, U^T B U)$ which solve the equation $\det(U^T (A - \kappa B) U) = 0$ are real positive and lie in between the maximum and minimum eigenvalues λ of the pair (A, B) , such that $\det(A - \lambda B) = 0$. Moreover, the extreme eigenvalues $\lambda_{\min}, \lambda_{\max}$ of the pair (A, B) are such that $\alpha_{\min}/\beta_{\max} \leq \lambda_{\min} \leq \lambda_{\max} \leq \alpha_{\max}/\beta_{\min}$, where $\alpha_{\min}, \alpha_{\max}, \beta_{\min}, \beta_{\max}$ are the minimum and maximum eigenvalues of the matrices A and B , respectively.*

Proof. The condition $\det(\lambda B - A) = 0$ is equivalent to $\det(\lambda I - B^{-1/2} A B^{-1/2}) = 0$ which has real positive solutions since $B^{-1/2} A B^{-1/2}$ is positive. The remaining part of the lemma follows from the fact that maximum and minimum eigenvalues of the larger and smaller problems coincide with maximum and minimum value of the Rayleigh quotient $x^T A x / x^T B x$ for $x \in \mathbb{R}^n$, and for $x \in \text{span}(U)$, respectively. \square

A first consequence of the above lemma is that the extreme eigenvalues κ_{\min} and κ_{\max} of K are in between the maximum and the minimum eigenvalue of the n^2 -dimensional symmetric matrix $Y = (I \otimes G^{-1}) H (I \otimes G^{-1})$, so that the ratio μ between the maximum and minimum eigenvalue of K is less than or equal to the condition number $\mu(Y)$ of the symmetric matrix Y . Moreover since $Y = \sum_{i=1}^m Y_i$ with,

$$Y_i = (A_i^{-1/2} \otimes A_i^{-1/2}) (I \otimes M_i^{-1}) \gamma(\log M_i \otimes I - I \otimes \log M_i) (I \otimes M_i^{-1}) (A_i^{-1/2} \otimes A_i^{-1/2})$$

one finds that $\widehat{\kappa}_{\min} := \sum_{i=1}^m \lambda_{\min}^{(i)} \leq \kappa_{\min}$ and $\widehat{\kappa}_{\max} := \sum_{i=1}^m \lambda_{\max}^{(i)} \geq \kappa_{\max}$, where $\lambda_{\min}^{(i)}$ and $\lambda_{\max}^{(i)}$ are the minimum and the maximum eigenvalues of Y_i . Moreover, from Lemma 4.3 and from the expression above for Y_i it follows that $\lambda_{\min}^{(i)} \geq \gamma_{\min}^{(i)}/(\alpha_{\max}^{(i)})^2$, $\lambda_{\max}^{(i)} \leq \gamma_{\max}^{(i)}/(\alpha_{\min}^{(i)})^2$, where $\alpha_{\min}^{(i)}$, $\alpha_{\max}^{(i)}$ are the minimum and the maximum eigenvalues of A_i , respectively, while $\gamma_{\min}^{(i)}$ and $\gamma_{\max}^{(i)}$ are the minimum and maximum eigenvalues of $(I \otimes M_i^{-1})\gamma(\log M_i \otimes I - I \otimes \log M_i)(I \otimes M_i^{-1})$, respectively.

From the properties of the matrix function $\gamma(\cdot)$ and from the properties of the Kronecker product one finds that the eigenvalues of the latter matrix can be explicitly given in terms of the eigenvalues $\nu_r^{(i)}$ of the matrix M_i . In fact, they coincide with $\frac{1}{(\nu_s^{(i)})^2}(\log t_{r,s}^{(i)})/(t_{r,s}^{(i)} - 1)$ where $t_{r,s}^{(i)} = \frac{\nu_r^{(i)}}{\nu_s^{(i)}}$.

Since the function $(\log t)/(t - 1)$ is monotone decreasing its minimum and maximum are

$$\begin{aligned}\eta_{\min}^{(i)} &= (\log \mu^{(i)})/(\mu^{(i)} - 1), \\ \eta_{\max}^{(i)} &= \log(1/\mu^{(i)})/(1/\mu^{(i)} - 1) = \mu^{(i)}(\log \mu^{(i)})/(\mu^{(i)} - 1),\end{aligned}$$

for $\mu^{(i)} = \mu(M_i)$ the spectral condition number of M_i . Additionally, taking the factor $(\nu^{(i)})^{-2}$ into consideration gives

$$\begin{aligned}\gamma_{\min}^{(i)} &\geq \eta_{\min}^{(i)} (\nu_{\max}^{(i)})^{-2}, \\ \gamma_{\max}^{(i)} &\leq \eta_{\max}^{(i)} (\nu_{\min}^{(i)})^{-2} \leq \mu^{(i)} (\nu_{\min}^{(i)})^{-2},\end{aligned}$$

where $\nu_{\min}^{(i)}$ and $\nu_{\max}^{(i)}$ represent respectively the minimum and maximum eigenvalue of M_i .

Therefore, we may conclude that the eigenvalues of K are bounded by $\widetilde{\kappa}_{\min} := \sum_{i=1}^m \eta_{\min}^{(i)}/(\nu_{\max}^{(i)} \alpha_{\max}^{(i)})^2$ and $\widetilde{\kappa}_{\max} := \sum_{i=1}^m \eta_{\max}^{(i)}/(\nu_{\min}^{(i)} \alpha_{\min}^{(i)})^2$.

Observe that this bound gets worse when either some matrix is ill-conditioned or if some matrix $A_i^{-1/2}GA_i^{-1/2}$ is ill-conditioned. The latter condition cannot occur if the matrices A_i do not differ much from G . The dependence of this bound to the conditioning of A_i makes this algorithm very inefficient as long as some A_i is ill-conditioned. This drawback is overcome in the next section where we design a more effective preconditioner.

4.3 A preconditioner based on a differential geometric viewpoint

The Karcher mean for positive matrices inherits a beautiful interpretation in terms of differential geometry. It can be considered as the center of mass for a well chosen inner product on the manifold of positive matrices. In this section and in Section 5 we consider two approaches inspired by this idea. For more information we refer to the overview in [22], and the articles [13, 18, 30–32].

When considering a manifold optimization approach, the intersection \mathcal{U} of a linear space \mathcal{T} with the manifold of positive matrices \mathcal{P}_n can be viewed as a Riemannian submanifold of \mathcal{P}_n itself, which in turn is called the *enveloping space*. This entails that the inner product from this enveloping space is induced on the submanifold. An immediate consequence is that the gradient of the cost function for the submanifold is given by the orthogonal projection (with respect

to the inner product) of the gradient for the enveloping space. Similar to the space of symmetric matrices being the tangent space to the manifold of positive matrices, the intersection \mathcal{V} of the linear space \mathcal{T} with the space of symmetric matrices is the tangent space to \mathcal{U} .

First consider the manifold of positive matrices endowed with the Euclidean inner product $g_X(A, B) = \text{tr}(AB)$, with A and B symmetric, and X a positive matrix. Note that even though this inner product g_X is independent from X , the subscript notation is kept for consistency. In this case, the orthogonal projection of a symmetric matrix A onto \mathcal{T} gives a matrix T , with

$$\text{vec}(T) = U (U^T U)^{-1} U^T \text{vec}(A),$$

or $\text{vec}(T) = Ut$, with

$$t = (U^T U)^{-1} U^T \text{vec}(A). \quad (4.6)$$

The expression for the gradient of the Karcher cost function, corresponding to the Euclidean inner product, is known for the manifold of positive matrices and is given by

$$\text{grad}_e f(X; A_1, \dots, A_m) = 2X^{-1} \sum_{i=1}^m \log(XA_i^{-1}). \quad (4.7)$$

The gradient naturally defines the direction of steepest ascent. Nevertheless, the gradient lies in the tangent space, and to build an algorithm from this a practical way is to follow the gradient and then go back to the manifold through a suitable function, called *retraction*. The precise definition of a retraction, together with general theoretical assumptions it should satisfy, can be found in [1]. Figure 2(a), graphically illustrates the concept of a retraction, where a vector ξ_X in the tangent space $T_X \mathcal{P}_n$ of the positive matrices is retracted to a point $R_X(\xi_X)$ residing on the manifold \mathcal{P}_n . On a manifold, the classical

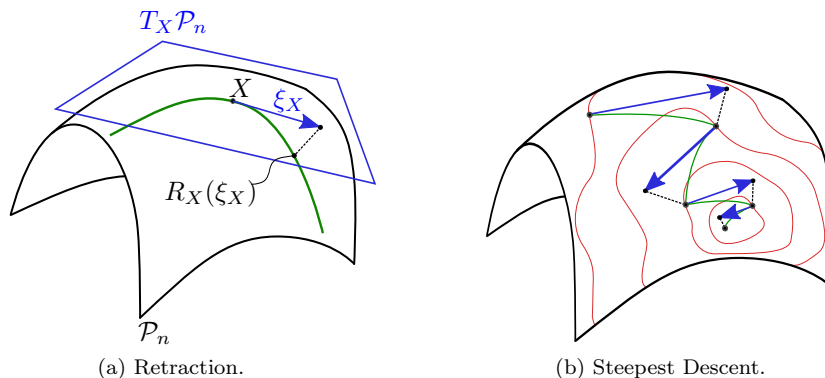


Figure 2: Graphical representation of a retraction and steepest descent flow.

steepest descent algorithm is graphically depicted in Figure 2(b). The thin red lines depict the contour lines, the blue arrows the gradients, and the green curves the retractions to the manifold.

Observe that for \mathcal{P}_n immersed in the set of symmetric matrices, the tangent space at a point is the whole set of symmetric matrices. So one can consider

the basic retraction $R_X(A) = X + A$ for a sufficiently small symmetric matrix A .

Entering now the gradient (4.7) in projection (4.6) and applying a gradient descent method with the basic retraction $R_X(A) = X + A$, we arrive exactly at the Richardson-like algorithm for finding the fixed points of function (4.1).

However, since the function f to be minimized is defined through the distance (1.2), it is more natural to consider the manifold of positive matrices endowed with the inner product $g_X(A, B) = \text{tr}(AX^{-1}BX^{-1})$, with A , B and X as before. In this case, the gradient for the enveloping space is known to be

$$\text{grad}_n f(X; A_1, \dots, A_m) = 2X \sum_{i=1}^m \log(A_i^{-1}X).$$

Note the difference with (4.7).

The orthogonal projection T onto the intersection \mathcal{V} (of \mathcal{T} and the space of symmetric matrices) of this gradient, with respect to the Riemannian scalar product, can be found as the solution of the equations

$$\begin{aligned} \text{grad}_n f(X) &= T + S, \\ g_X(S, K) &= \text{tr}(SX^{-1}KX^{-1}) = 0, \quad \text{for every } K \in \mathcal{V}. \end{aligned}$$

Writing again $\text{vec}(T) = Ut$, we find in parameter space

$$t = (U^T (X^{-1} \otimes X^{-1}) U)^{-1} U^T (X^{-1} \otimes X^{-1}) \text{vec}(\text{grad}_n f(X)). \quad (4.8)$$

The factor $U^T (X^{-1} \otimes X^{-1}) U$ is recurring and is abbreviated as D_X , where the subscript points to the intrinsic variable X . Observe that this Riemannian orthogonal projection can be seen as a Euclidean oblique projection where the two bases of the subspace are the columns of U and $(X^{-1} \otimes X^{-1})U$, respectively.

Using this expression, it is possible to define another gradient descent method where we are now searching the fixed points of the function

$$\varphi(t) = t - \theta D_{\sigma(t)}^{-1} U^T (\sigma(t)^{-1} \otimes \sigma(t)^{-1}) \text{vec}\left(\sigma(t) \sum_{i=1}^m \log(A_i^{-1}\sigma(t))\right). \quad (4.9)$$

Relying on (1.6) to incorporate the Kronecker product into the vectorization, we find that $(\sigma^{-1} \otimes \sigma^{-1}) \text{vec}(\sigma \sum_{i=1}^m \log(A_i^{-1}\sigma)) = \text{vec}(\sum_{i=1}^m \log(A_i^{-1}\sigma)\sigma^{-1})$. Applying a property of the matrix logarithm we may rewrite the latter expression as $\text{vec}(\sigma^{-1} \sum_{i=1}^m \log(\sigma A_i^{-1}))$. This way, equation (4.9) takes the form of (4.1) with

$$V = U^T (G^{-1} \otimes G^{-1}) U.$$

To analyze the convergence of (4.1) with the choice $V = U^T (G^{-1} \otimes G^{-1}) U$, we have to analyze the eigenvalues of the Jacobian $K = J_S(t_G)$ of $S(t)$ in (4.1) where the equation $\det(\kappa I - K) = 0$ takes the form of the following generalized eigenvalue problem

$$\det(U^T (\kappa(G^{-1} \otimes G^{-1}) - (I \otimes G^{-1})H(I \otimes G^{-1})) U) = 0. \quad (4.10)$$

Since the two matrices in equation (4.10) are positive, in view of Lemma 4.3, the solutions of this generalized eigenvalue problem are real positive and

are located in between the maximum and the minimum solution of the larger problem

$$\det(\lambda(G^{-1} \otimes G^{-1}) - (I \otimes G^{-1})H(I \otimes G^{-1})) = 0,$$

which in turn can be rewritten as a standard eigenvalue problem

$$\det(\lambda I - (G^{1/2} \otimes G^{1/2})(I \otimes G^{-1})H(I \otimes G^{-1})(G^{1/2} \otimes G^{1/2})) = 0.$$

Since $H = \sum_{i=1}^m H_i$, and the matrices H_i are real symmetric, the eigenvalues of this problem are located in between the sum of the minimum and the sum of the maximum eigenvalues of each subproblem

$$\det(\lambda I - (G^{1/2} \otimes G^{1/2})(I \otimes G^{-1})H_i(I \otimes G^{-1})(G^{1/2} \otimes G^{1/2})) = 0, \quad (4.11)$$

that is $\det(\lambda(G^{-1} \otimes G) - H_i) = 0$, or equivalently $\det(\lambda I - (G \otimes I)H_i(I \otimes G^{-1})) = 0$. The matrix in the latter expression is similar to $(A_i^{-1/2} \otimes A_i^{-1/2})(G \otimes I)H_i(I \otimes G^{-1})(A_i^{1/2} \otimes A_i^{1/2})$, which, using the expression of H_i provided in Theorem 4.2, can be written as

$$(M_i \otimes I)\gamma(\log M_i \otimes I - I \otimes \log M_i)(I \otimes M_i^{-1}).$$

This way, the eigenvalues of (4.11) can be explicitly given in terms of the eigenvalues $\nu_r^{(i)}$ of the matrix M_i . In fact, they coincide with the $t_{r,s}^{(i)}(\log t_{r,s}^{(i)})/(t_{r,s}^{(i)} - 1)$ where $t_{r,s}^{(i)} = \frac{\nu_r^{(i)}}{\nu_s^{(i)}}$.

Since the function $t(\log t)/(t - 1)$ is monotone, for the minimum and maximum solution to (4.11) we have

$$\begin{aligned} \eta_{\min}^{(i)} &= (1/\mu^{(i)}) \log(1/\mu^{(i)})/(1/\mu^{(i)} - 1) = (\log \mu^{(i)})/(\mu^{(i)} - 1), \\ \eta_{\max}^{(i)} &= \mu^{(i)} (\log \mu^{(i)})/(\mu^{(i)} - 1), \end{aligned}$$

respectively, for $\mu^{(i)} = \mu(M_i)$ the spectral condition number of M_i . Therefore, we may conclude that the eigenvalues of K are in between $\sum_{i=1}^m \eta_{\min}^{(i)}$ and $\sum_{i=1}^m \eta_{\max}^{(i)}$. This way, we find for the optimal value of θ and for the optimal spectral radius the estimates

$$\begin{aligned} \theta &= \frac{2}{\sum_{i=1}^m \frac{\mu^{(i)}+1}{\mu^{(i)}-1} \log \mu^{(i)}}, \\ \rho &= \frac{\sum_{i=1}^m \log \mu^{(i)}}{\sum_{i=1}^m \frac{\mu^{(i)}+1}{\mu^{(i)}-1} \log \mu^{(i)}}. \end{aligned}$$

It is interesting to point out that in this case the convergence speed is related neither to the condition number of the geometric mean G nor of the matrices A_i but is related only to the relative distances of G from each A_i measured by the quantities $\mu^{(i)} = \mu(M_i)$, $M_i = A_i^{-1/2} G A_i^{-1/2}$. The closer they are to 1 the faster the convergence. Therefore, if the matrices to average are not much far from each other so that the quantities $\mu(M_i)$ are close to 1, then the optimal value of θ is close to $1/m$ and a very fast convergence is expected. This analysis is confirmed by the numerical experiments.

4.4 The case of Toeplitz matrices

From the computational point of view, at each step of the iteration (4.1) one has to compute $U^T \text{vec}(\Gamma(\sigma(t)))$, and then solve a linear system with the matrix $V(\sigma(t))$. The former computation, based on (3.6), requires $O(mn^3)$ arithmetic operations (ops), while the cost of the latter depends on the structure of $V(\sigma(t))$.

In this section we examine the case where \mathcal{U} is the class of symmetric Toeplitz matrices and where $\sigma(t)$ associates t with the Toeplitz matrix having as first column t . We describe a way to make the algorithm of Section 4.3 more efficient.

Indeed, for the iteration analyzed in Section 4.2, V is the diagonal matrix with diagonal entries $(n, 2n - 2, \dots, 2)$ and the cost of solving a system with matrix V amounts to n divisions.

The iteration examined in Section 4.3 has a higher convergence speed but at each step an $n \times n$ system with $V = U^T(X^{-1} \otimes X^{-1})U$ must be solved, where X is a symmetric positive definite Toeplitz matrix.

We split the computation in two steps. In the first, the n^2 entries of V are computed, in the second step a standard $O(n^3)$ ops linear system solver is used. Concerning the first step we discuss two approaches.

In both approaches the inverse of the Toeplitz matrix X needs to be computed, which can be efficiently done using the Gohberg Semencul formula [9]. Here, vectors v_1, v_2, v_3, v_4 are determined such that $X^{-1} = L(v_1)L(v_2)^T - L(v_3)L(v_4)^T$, where $L(v)$ is the lower triangular Toeplitz matrix whose first column is v . From these, the n^2 entries of X^{-1} can be found. The overall cost is $O(n^2)$ ops.

1. As a first attempt, the entries of V are computed in a straightforward manner using the entries of X^{-1} :

$$V = \begin{bmatrix} \gamma_{1,1} & 2\gamma_{1,2} & \cdots & 2\gamma_{1,n} \\ 2\gamma_{1,2} & 2\gamma_{2,2} & \cdots & 2\gamma_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 2\gamma_{1,n} & 2\gamma_{2,n} & \cdots & 2\gamma_{n,n} \end{bmatrix},$$

where

$$\begin{aligned} \gamma_{1,j} &= \sum_{i=1}^n \sum_{k=1}^{n-j+1} (X^{-1})_{i,k} (X^{-1})_{i,k+j-1}, \\ \gamma_{j,p} &= \sum_{i=1}^{n-j+1} \sum_{k=1}^{n-p+1} ((X^{-1})_{i,k} (X^{-1})_{i+j-1,k+p-1} + (X^{-1})_{i,k+p-1} (X^{-1})_{i+j-1,k}). \end{aligned}$$

The cost of this approach in terms of arithmetic operations is of the order $O(n^4)$.

2. In the second approach, we show that the cost of this computation can be kept at the level of $O(n^3 \log n)$ ops by combining the Gohberg Semencul formula and the FFT. For a given i , the product vector $w_i = (X^{-1} \otimes X^{-1})Ue_i$, where e_i is the i th vector of the canonical basis, is such that $w_i = \text{vec}(X^{-1}E_iX^{-1})$, with E_i being the symmetric Toeplitz matrix whose first column is e_i . Therefore, compute first the columns of E_iX^{-1} by performing $O(n^2)$ additions, and then multiply

X^{-1} by these columns, stacking the results to obtain w_i . This computation is performed in $O(n^2 \log n)$ operations for each i by using the Gohberg Semencul formula, since the multiplication of a lower triangular Toeplitz matrix and a vector can be performed in $O(n \log n)$ operations by means of the FFT [9]. Therefore the overall computation of this stage for $i = 1, \dots, n$ is $O(n^3 \log n)$ ops. Finally, compute for any i the vector $U^T w_i$ for the cost of $O(n^2)$ additions.

The performance of these methods will be compared in Section 6.

5 Kähler metric mean for Toeplitz matrices

The Karcher mean of positive definite matrices has the specific interpretation of being the barycenter of the given matrices for the natural metric (1.2) on this manifold. Hence there are in a certain sense two possible generalizations. On the one hand, try to generalize the geometric mean concept, or, on the other hand, try to generalize the barycenter concept. Previously we focused on an extension of the geometric mean. Hereafter we focus on the positive definite Toeplitz matrix manifold itself, denoted by \mathcal{T}_n , and consider a barycenter in this case. This mean can not be called a geometric mean in the sense of satisfying all required properties, but through its intuitive definition, many desirable properties could arise.

The concept of a barycenter is not restricted to the specific metric used to define the Karcher mean. For example, when the set \mathcal{T}_n is endowed with the classical Euclidean inner product, the resulting barycenter is nothing else than the arithmetic mean. Using a probability argument, in [3, 4] a metric on \mathcal{T}_n is introduced, called the Kähler metric. This metric results in a complete, simply connected manifold with non-positive sectional curvature everywhere, or a Cartan–Hadamard manifold. Thus, by [15, 23], existence and uniqueness are guaranteed for the barycenter with respect to this metric.

We will recall some known facts about the Kähler metric, and then we will give an explicit formula for the barycenter in the real case and a numerical procedure to compute the barycenter in the complex case.

To construct the Kähler metric, a Toeplitz matrix is first transformed to an n -tuple $(P_0, \mu_1, \dots, \mu_{n-1})$ in $\mathbb{R}_+^* \times \mathbb{D}^{n-1}$, with \mathbb{R}_+^* the set of strictly positive real numbers and \mathbb{D} the set of complex numbers of modulus less than one. This transformation, denoted as $\zeta(T) = [P_T, \mu_{T,1}, \dots, \mu_{T,n-1}]^T$, is performed as follows:

$$P_T = t_0, \quad \mu_{T,j} = (-1)^j \frac{\det(S_j)}{\det(R_j)},$$

with t_0 the main diagonal element of T , R_j the principal submatrix of size j of T (the upper left $j \times j$ submatrix) and S_j obtained by shifting R_j down one row, or equivalently, by removing the first row and last column of R_{j+1} (the inverse transformation can be found in [34]). In what follows, we use this one-to-one relation between the Toeplitz matrices and the corresponding n -tuple, and when clear by the context, we will neglect the distinction and identify one with the other.

For X and Y being the transformations of two positive Toeplitz matrices $X = [P_X, \mu_{X,1}, \dots, \mu_{X,n-1}]^T$ and $Y = [P_Y, \mu_{Y,1}, \dots, \mu_{Y,n-1}]^T$, the metric is

given by

$$d(X, Y) = \left(n\sigma(P_X, P_Y)^2 + \sum_{j=1}^{n-1} (n-j)\tau(\mu_{X,j}, \mu_{Y,j})^2 \right)^{1/2},$$

$$\sigma(P_X, P_Y) = \left| \log\left(\frac{P_Y}{P_X}\right) \right|, \quad \tau(\mu_{X,j}, \mu_{Y,j}) = \operatorname{atanh}\left(\left| \frac{\mu_{Y,j} - \mu_{X,j}}{1 - \mu_{X,j}\mu_{Y,j}} \right| \right),$$

where $\operatorname{atanh}(z) = \frac{1}{2} \log\left(\frac{1+z}{1-z}\right)$.

The barycenter of the positive Toeplitz matrices T_i , for $i = 1, \dots, m$, with respect to the Kähler metric will be denoted by $B(T_1, \dots, T_m) = [P_B, \mu_{B,1}, \dots, \mu_{B,n-1}]^T$. It is obtained in this transformed space by minimizing the function

$$f(X) = \sum_{i=1}^m d^2(X, T_i)$$

over $\mathbb{R}_+^* \times \mathbb{D}^{n-1}$. Notice that the problem of minimizing $f(X)$ can be decoupled into the problems of minimizing $\varphi_0(x) = \sum_{i=1}^m \sigma(x, P_{T_i})^2$ over \mathbb{R}_+^* , and the $n-1$ scalar functions

$$\varphi_j(z) = \sum_{i=1}^m \tau(z, \mu_{T_i,j})^2, \quad j = 1, \dots, n-1$$

over \mathbb{D} . The minimum of $\varphi_0(x)$ is easily obtained as $P_B = (P_{T_1} \cdots P_{T_m})^{1/m}$ by solving the equation $\nabla \varphi_0(x) = 0$. The minimum of $\varphi_j(z)$ is nothing else than the barycenter of $\mu_{T_1,j}, \dots, \mu_{T_m,j}$ with respect to the customary Poincaré metric on the unit disk and is the point where the gradient

$$\nabla \varphi_j(z) = 2(|z|^2 - 1) \sum_{i=1}^m \operatorname{sign}(c_{i,j}) \operatorname{atanh}(|c_{i,j}|), \quad c_{i,j} = \frac{\mu_{T_i,j} - z}{1 - \bar{z}\mu_{T_i,j}}, \quad (5.1)$$

equals zero.

In the real case we are able to find an explicit expression for this barycenter as well, since $\operatorname{sign}(c) \operatorname{atanh}(|c|) = \operatorname{atanh}(c)$ and after some manipulations we get

$$\mu_{X,j} = \mathcal{C}\left(\left(\mathcal{C}(\mu_{T_1,j}) \cdots \mathcal{C}(\mu_{T_m,j})\right)^{1/m}\right),$$

where $\mathcal{C}(z) = (1-z)/(1+z)$ is the Cayley transform.

In the complex case we were not able to find such an explicit formula but a quick numerical method can be devised using a gradient descent algorithm. We recall that the tangent space to the Poincaré disk can be identified with the complex plane and thus for a sufficiently small tangent vector $v \in \mathbb{C}$, one can consider the retraction $R_z(v) = z + v$, which captures the fact that the manifold is an open subset of the complex plane. The resulting algorithm to find the barycenter of $\mu_1, \dots, \mu_n \in \mathbb{C}$ is given by the iteration

$$z_{k+1} = z_k + t_k v_k, \quad v_k = (1 - |z_k|^2) \sum_{i=1}^n \operatorname{sign}(c_{i,k}) \operatorname{atanh}(|c_{i,k}|), \quad c_{i,k} = \frac{\mu_i - z_k}{1 - \bar{z}_k \mu_i}, \quad (5.2)$$

for a suitable initial value z_0 and a sufficiently small steplength t_k .

Another possibility is to consider the retraction

$$R_z(v) = \frac{z + e^{i\theta} + (z - e^{i\theta})e^{-s}}{1 + \bar{z}e^{i\theta} + (1 - \bar{z}e^{i\theta})e^{-s}}, \quad \theta = \arg v, \quad s = \frac{2|v|}{1 - |z|^2},$$

which corresponds to moving along the geodesics of the Poincaré disk. The corresponding gradient descent method is

$$z_{k+1} = R_{z_k}(t_k v_k),$$

with the same v_k as (5.2).

5.1 Properties of the Kähler barycenter

Regarding the properties of this barycenter, it is easily seen that it is permutation invariant, repetition invariant and idempotent (this holds for any barycenter). Moreover, for any $\alpha > 0$, it holds that the transformed values of αT are $[\alpha P_T, \mu_{T,1}, \dots, \mu_{T,m}]^T$ and from the explicit expression of P_B in the real case we get that $B(\alpha T_1, \alpha T_2, \dots, \alpha T_m) = \alpha^{1/m} B(T_1, \dots, T_m)$, that is, homogeneity holds.

Unfortunately, this new barycenter does not verify other properties as shown by the following example.

Example 5.1. From the explicit expression for the mean in the real case we get a simple formula for the Kähler barycenter of two 2×2 matrices

$$T_1 = \begin{bmatrix} x_1 & y_1 \\ y_1 & x_1 \end{bmatrix}, \quad T_2 = \begin{bmatrix} x_2 & y_2 \\ y_2 & x_2 \end{bmatrix},$$

namely

$$B(T_1, T_2) = \sqrt{x_1 x_2} \begin{bmatrix} 1 & \frac{a-b}{a+b} \\ \frac{a-b}{a+b} & 1 \end{bmatrix}, \quad \text{with } \begin{cases} a = \sqrt{(x_1 + y_1)(x_2 + y_2)} \\ b = \sqrt{(x_1 - y_1)(x_2 - y_2)} \end{cases}.$$

Now consider the following matrices

$$T_1 = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}, \quad \tilde{T}_1 = \begin{bmatrix} 4 & -1 \\ -1 & 4 \end{bmatrix}, \quad T_2 = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix},$$

with $\tilde{T}_1 \geq T_1$. By symbolic computation, one gets that

$$B(\tilde{T}_1, T_2) = \begin{bmatrix} 2\sqrt{2} & \sqrt{2}(\sqrt{5}-3) \\ \sqrt{2}(\sqrt{5}-3) & 2\sqrt{2} \end{bmatrix} \not\geq B(T_1, T_2) = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix},$$

in fact one eigenvalue of $B(\tilde{T}_1, T_2) - B(T_1, T_2)$ is $\lambda = \sqrt{10} - 2 - \sqrt{2} < 0$. Thus, we have proved that the Kähler barycenter is not monotonic. Moreover,

$$B(T_1, T_2) \neq (T_1 T_2)^{1/2} = \begin{bmatrix} \sqrt{3} & 0 \\ 0 & \sqrt{3} \end{bmatrix},$$

and hence the Kähler barycenter does not coincide with the Karcher mean for circulant matrices. In particular, it is not a structured geometric mean as defined in Section 1.

Observe that in the previous example $B(T_1, T_2)$ surprisingly coincides with the arithmetic mean of T_1 and T_2 . It is not difficult to construct examples where it is not true that $B(T_1, T_2) \leq (T_1 + T_2)/2$ as should be for a geometric mean.

6 Numerical experiments

In this section, the different algorithms proposed in Sections 4.2 and 4.3 will be compared w.r.t. speed and accuracy. The numerical experiments are confined to Toeplitz matrices, because of applicational interest in computing their structured matrix mean [24]. These matrices are constructed randomly, but with chosen condition number, using the technique described in [16]. Performance, accuracy and computational distance are subjects of the forthcoming investigations. For clarity we remind the reader that the Richardson-iteration corresponds to a projection technique on a manifold, with the classical Euclidean inner product. For all algorithms, the stopping criteria is based on checking the relative size of the gradient and on comparing two consecutive iteration points.

It is worth pointing out that, in spite of the lack of the proof of uniqueness for structured geometric mean in the Toeplitz case, for any fixed set of data matrices used in our experiments, any initial value and any algorithm yielded always the same structured geometric mean. This suggests the conjecture that in the Toeplitz case there is a unique structured geometric mean.

We have also compared the structured geometric mean obtained by our algorithms with the Kähler metric mean, getting in most experiments a relative difference of the order 10^{-1} , which indicates that these two means are relatively far from each other.

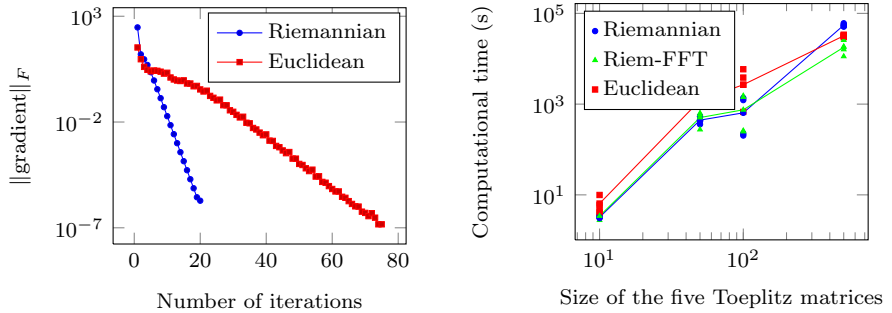
6.1 The projection methods

Performance The performance of the projection methods explained in Section 4 can be compared by looking at both the number of iterations the methods require and the total amount of computational time they need.

In Figure 3(a), the evolution of the gradient over the iterations is displayed for both techniques (and hence also the number of iterations). Using the projection method introduced in Section 4.3 gives a faster decrease of the gradient and results in fewer iteration steps. The number of iterations remains almost constant for this method as the size of the matrices increases. For the projection technique from Section 4.2 on the other hand, this number starts to increase when the matrix size grows.

However, comparing expression (4.6) and (4.8), it can be seen that the second one is computationally more expensive and hence the advantage of requiring fewer iterations could be nullified. Therefore, Figure 3(b) displays the total computational time of both methods for varying sizes of the matrices (both approaches from Section 4.4 are shown). The two methods based on Section 4.3 maintain an advantage despite their larger computational cost per iteration. Note that for the largest matrix size the computational time of the Euclidean based method appears less than one of the other methods. However, this is caused by the increasing number of iterations required by this Euclidean method. Consequently, the maximum number of iterations is reached before convergence and the algorithms is terminated prematurely. Concerning the operation count in Section 4.4, the advantage of the method based on FFT starts to appear when the matrices become sufficiently large.

Accuracy In order to analyze the accuracy of the projection methods, we implement a high precision version of the first algorithm in Section 4.4 using the



(a) Gradient evolution for five 50×50 matrices.

(b) Computational time for various matrix sizes. The mean of the samples is connected by a line.

Figure 3: Comparison of the projection methods for Toeplitz matrices. In the legends, Euclidean indicates the method of Section 4.2, Riemannian indicates the first approach described in Section 4.4, and Riem-FFT the second.

`vpa` functionality of `MATLAB`. The relative distance, based on the intrinsic distance (1.2), between this high precision computation and the result of the actual algorithms is shown in Figure 4. For small condition numbers, the accuracy of all methods is similar in average, but as the condition of the matrices becomes worse, the accuracy of the projection method based on Euclidean geometry deteriorates much faster than that of the method based on the Riemannian geometry. This first method even fails to converge when the condition number of the matrices becomes significantly large. The accuracy of the two approaches in Section 4.4 is similar and deteriorates steadily as the condition numbers of the matrices increase.

7 Conclusions

In this article a generalization of the Karcher mean for positive definite matrices to structured positive definite matrices was proposed. Besides a theoretical investigation and adaptation of the desired properties of such a mean, algorithms were proposed. In the design of the algorithms, two trajectories were put forward, one relying mostly on linear algebra, and one based on differential geometry. A convergence analysis has been performed showing the superiority of the algorithm based on differential geometry. Numerical experiments compared the accuracy and speed of the various techniques and confirmed the theoretical analysis.

In the case of Toeplitz matrices, we have considered also the Kähler metric mean [4], whose properties have been investigated, providing an explicit expression in the real case and a quick algorithm in the complex case. For Toeplitz matrices, both the new structured geometric mean and the Kähler metric mean have not completely satisfying properties, in fact they are not monotone, neither do they satisfy the arithmetic-geometric inequality. We wonder if it is possible to provide a definition of geometric mean for Toeplitz matrices which behaves well with respect to the ordering of positive matrices.

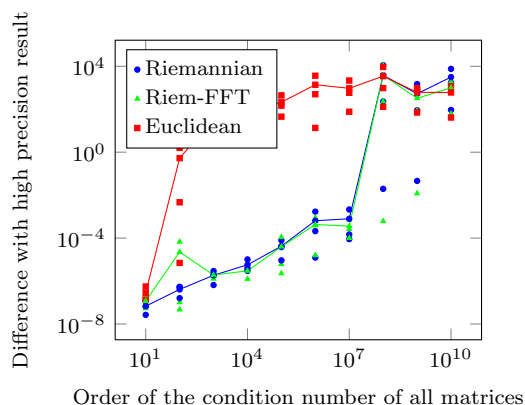


Figure 4: Accuracy of the projection methods when compared to a high precision version. The mean of the samples is connected by a line. In the legends, Euclidean indicates the method of Section 4.2, Riemannian indicates the first approach described in Section 4.4, and Riem-FFT the second.

References

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.
- [2] T. Ando, C. Li, and R. Mathias. Geometric means. *Linear Algebra and its Applications*, 385:305–334, 2004.
- [3] F. Barbaresco. Information Intrinsic Geometric Flows. In A. Mohammad-Djafari, editor, *Bayesian Inference and Maximum Entropy Methods In Science and Engineering*, volume 872 of *American Institute of Physics Conference Series*, pages 211–218, Nov. 2006.
- [4] F. Barbaresco. Interactions between symmetric cone and information geometries: Bruhat-tits and siegel spaces models for high resolution autoregressive doppler imagery. In F. Nielsen, editor, *Emerging Trends in Visual Computing*, volume 5416 of *Lecture Notes in Computer Science*, pages 124–163. Springer Berlin / Heidelberg, 2009.
- [5] P. G. Batchelor, M. Moakher, D. Atkinson, F. Calamante, and A. Connelly. A Rigorous Framework for Diffusion Tensor Calculus. *Magnetic Resonance in Medicine*, 53:221–225, 2005.
- [6] R. Bhatia. *Positive Definite Matrices*. Princeton Series in Applied Mathematics. Princeton University Press, 2007.
- [7] R. Bhatia and J. Holbrook. Riemannian geometry and matrix geometric means. *Linear Algebra and its Applications*, 413(2-3):594–618, 2006.
- [8] D. Bini, B. Meini, and F. Poloni. An effective matrix geometric mean satisfying the Ando–Li–Mathias properties. *Mathematics of Computation*, 79(269):437–452, 2010.

- [9] D. Bini and V. Pan. *Polynomial and Matrix Computations*. Birkhäuser, 1994.
- [10] D. A. Bini and B. Iannazzo. Computing the Karcher mean of symmetric positive definite matrices. *Linear Algebra and its Applications*, 2011. In Press, doi 10.1016/j.laa.2011.08.052.
- [11] D. A. Bini and B. Iannazzo. A note on computing matrix geometric means. *Advances in Computational Mathematics*, 35(2–4):175–192, 2011.
- [12] D. A. Bini and B. Iannazzo. Computing the Karcher mean of symmetric positive definite matrices. *Linear Algebra Appl.*, 2012. To appear.
- [13] S. Bonnabel and R. Sepulchre. Geometric distance and mean for positive semi-definite matrices of fixed rank. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1055–1070, 2009.
- [14] M. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer-Verlag, 1999.
- [15] E. Cartan. Leçons sur la géométrie des espaces de Riemann. *Uspekhi Matematicheskikh Nauk*, 3(3):218–222, 1948.
- [16] M. T. Chu and G. H. Golub. Structured inverse eigenvalue problems. *Acta Numerica*, 11:1–71, 2002.
- [17] Y. Eidelman and I. C. Gohberg. On a new class of structured matrices. *Integral Equations and Operator Theory*, 34:293–324, 1999.
- [18] S. Fiori. Learning the Fréchet mean over the Manifold of Symmetric Positive-Definite Matrices. *Springer Science, Business Media*, 2009.
- [19] N. J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2008.
- [20] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, Cambridge, 1991.
- [21] B. Iannazzo. The geometric mean of two matrices from a computational viewpoint. Technical report, Dipartimento di Matematica e Informatica, Università di Perugia, Perugia, Italy, 2011. arXiv:1201.0101v1 [math.NA].
- [22] B. Jeuris, R. Vandebril, and B. Vandereycken. A survey and comparison of contemporary algorithms for computing the matrix geometric mean. *Electronic Transactions on Numerical Analysis*, 39:379–402, 2012.
- [23] W. Kendall. Probability, convexity, and harmonic maps with small image I: uniqueness and fine existence. *Proceedings London Mathematical Society*, 61(2):371–406, 1990.
- [24] J. Lapuyade-Lahorgue and F. Barbaresco. Radar detection using Siegel distance between autoregressive processes, application to HF and X-band radar. In *Radar Conference, 2008. RADAR '08, IEEE*, Rome, May 2008.
- [25] J. Lawson and Y. Lim. Monotonic properties of the least squares mean. *Mathematische Annalen*, 351(2):267–279, 2011.

- [26] Y. Lim. Riemannian and Finsler structures of symmetric cones. *Trends in Mathematics*, 4(2):111–118, 2001.
- [27] M. Moakher. A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications*, 26(3):735–747, 2005.
- [28] M. Moakher. On the averaging of symmetric positive-definite tensors. *Journal of Elasticity*, 82(3):273–296, 2006.
- [29] K. Nakamura. Geometric means of positive operators. *Kyunpook Mathematical Journal*, 49:167–181, 2009.
- [30] M. Pálfi. The Riemann barycenter computation and means of several matrices. *International Journal of Computational and Mathematical Sciences*, 3(3):128–133, 2009.
- [31] C. Qi. *Numerical optimization methods on Riemannian manifolds*. PhD thesis, Florida State University, College of arts and sciences, 2011.
- [32] Q. Rentmeester and P.-A. Absil. Algorithm comparison for Karcher mean computation of rotation matrices and diffusion tensors. In *European Signal Processing, 2011, EUSIPCO. 19th conference on*, 2011.
- [33] R. Vandebril, M. Van Barel, and N. Mastronardi. *Matrix Computations and Semiseparable Matrices. Vol. 1, Linear Systems*. The Johns Hopkins University Press, 2008.
- [34] L. Yang, M. Arnaudon, and F. Barbaresco. Geometry of covariance matrices and computation of median. In *AIP Conference Proceedings, 30th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, volume 1305, pages 479–486, 2011.