Electronic Transactions on Numerical Analysis.
Volume 39, pp. 379-402, 2012.

ISSN 1068-9613.

ETNA
Kent State University
http://etna.math.kent.edu

A SURVEY AND COMPARISON OF CONTEMPORARY ALGORITHMS FOR COMPUTING THE MATRIX GEOMETRIC MEAN*

BEN JEURIS†, RAF VANDEBRIL†, AND BART VANDEREYCKEN‡

Abstract. In this paper we present a survey of various algorithms for computing matrix geometric means and derive new second-order optimization algorithms to compute the Karcher mean. These new algorithms are constructed using the standard definition of the Riemannian Hessian. The survey includes the ALM list of desired properties for a geometric mean, the analytical expression for the mean of two matrices, algorithms based on the centroid computation in Euclidean (flat) space, and Riemannian optimization techniques to compute the Karcher mean (preceded by a short introduction into differential geometry). A change of metric is considered in the optimization techniques to reduce the complexity of the structures used in these algorithms. Numerical experiments are presented to compare the existing and the newly developed algorithms. We conclude that currently first-order algorithms are best suited for this optimization problem as the size and/or number of the matrices increase.

Key words. matrix geometric mean, positive definite matrices, Karcher mean, Riemannian optimization

AMS subject classifications. 15A24, 53B21, 65K10

1. Introduction. A mean is, in general, simply a type of average subject to certain generic properties such as idempotency (the mean of $(A, \ldots, A)$ equals A), invariance under a permutation of the elements, and homogeneity (the mean of $(\lambda A_1, \ldots, \lambda A_K)$ equals λ times the mean of $(A_1, \ldots, A_K)$). However, these generic properties alone do not uniquely define a mean, so there can be many different types of means. In this paper we discuss the *geometric mean*, which for positive real numbers $(a_1, \ldots, a_K)$ is defined as

$$\mathbf{G}(a_1, \ldots, a_K) = (a_1 \cdots a_K)^{\frac{1}{K}}. \tag{1.1}$$

When conveying this definition to the set of symmetric positive definite $n \times n$ matrices $\mathbb{S}_+^n$, we see that the formula above can not be readily extended to matrices due to their non-commutativity. However, a list of desired properties for the general geometric mean can be derived from this scalar expression.

These properties (listed in Section 2) have proven to be useful in various applications, e.g., radar technology [6], medical imaging [19], mechanics [30], and image processing [37]. All these areas display situations in which the information about the current system is being represented in a collection of positive definite matrices. In order to perform calculations on these matrices, such as averaging and interpolation, we need algorithms that preserve the structure of the data, such as positive definiteness, which is one of the useful properties of the geometric mean. Another property of this mean provides advantages in the area of elasticity calculations of structures [30]. In these calculations, both a positive definite elasticity matrix and its inverse, the compliance matrix, are used. Hence, given a collection of these elasticity matrices and a collection consisting of the corresponding compliance matrices, the geometric means of both matrix collections will again be each others inverses (as stated in property 8 in Section 2).

Thanks to the wide range of practical and theoretical applications, matrix means have received a lot of attention from present-day scientists. Monotonicity, for example, is a property

* Received March 22, 2012. Accepted August 16, 2012. Published online on November 5, 2012. Recommended by L. Reichel.

†Department of Computer Science, KU Leuven, Celestijnenlaan 200A, 3001 Heverlee, Belgium (`{ben.jeuris, raf.vandebril}@cs.kuleuven.be`).

‡Chair of ANCHP, MATHICSE, Mathematics Section, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland (`bart.vandereycken@epfl.ch`).

only recently proven for the Karcher mean [8, 23], which is a specific instance of a geometric mean for positive definite matrices. A consequence of the diversity of application areas is the wide variety of approaches to define and compute the mean. Some constructions are based on intuitive interpretations of the geometric mean (Section 3), while others prefer to think of it as an optimization problem (Section 4). In this last approach, Riemannian optimization, which generalizes classical optimization techniques, is a popular concept.

The main contribution of this paper is to present a survey of algorithms for computing the matrix geometric mean. We recall the theoretical foundation for the analytically known mean of two matrices, the interpretations of the algorithms based on intuitive approaches, and a basic framework needed to understand the methods based on Riemannian optimization. We also introduce a new, explicit expression for the Riemannian Hessian and consider the use of a different inner product on the manifold of positive definite matrices $\mathbb{S}_+^n$, which leads to simpler optimization algorithms. Finally, a first-time application of the Riemannian BFGS method to the optimization problem is presented. Numerical experiments are performed to compare all these techniques.

The organization of this paper is as follows: we start by listing the desired properties of the geometric mean and the resulting unique definition in case of two matrices in Section 2. Next, in Section 3, we discuss some intuitively appealing algorithms based on planar approaches: the ALM, NBMP, and the CHEAP mean. However, these appealing interpretations will not lead to very efficient numerical algorithms. Finally, in Section 4, we examine Riemannian optimization algorithms for the Karcher mean, which is defined as the minimizer over all positive definite matrices of the sum of squared (intrinsic) distances to all matrices in the mean. The algorithms are adapted versions of the steepest descent, conjugate gradient, trust region, and BFGS methods generalized towards manifolds. This last section is more descriptive than the previous ones in order to properly introduce the Riemannian geometry and the construction of the generalized algorithms. Throughout the paper, we compare the performance of the algorithms discussed.

There have also been suggestions of other means, which do not satisfy all properties of the geometric mean, but which tend to give very close results. In [34], the mean of two matrices is used to construct the general mean of any number of matrices. However, the result of the algorithm depends on the ordering of the matrices, which causes the mean to be variable under permutation. A different class of means was presented in [28] satisfying a number of interesting properties. The means in this class depend on a parameter $t \in [-1, 1]\backslash\{0\}$, and it is shown that as t approaches zero, the corresponding mean approaches the Karcher mean. Another class of means with results close to the geometric mean is discussed in [27], where again the Karcher mean arises as a special case. These means are mentioned for completeness of the survey, but will not be discussed further.

2. The geometric mean of two matrices. The scalar geometric mean (1.1) can not be readily extended to positive definite matrices because the matrix product is not commutative. Indeed, $(A_1 \cdots A_K)^{\frac{1}{K}}$ is not invariant under permutation, which is one of the most basic properties of means. Hence, a list of desired features has been composed instead, often referred to as the ALM list [5, 11]. Because of the importance of these properties, we summarize them here using the partial ordering of symmetric matrices: a positive semidefinite matrix A is denoted by $A \geq 0$. Similarly, $B \geq C$ is a simplified notation for $B - C \geq 0$. The same approach is used for positive definiteness with the strict inequality. The ALM list, using positive definite matrices $A_1, \ldots, A_K$, where we denote the *geometric mean* by $\mathbf{G}(A_1, \ldots, A_K)$, is given by the following properties:

1. Consistency: if $A_1, \ldots, A_K$ commute, then $\mathbf{G}(A_1, \ldots, A_K) = (A_1 \cdots A_K)^{\frac{1}{K}}$.

2. Joint homogeneity:

$$\mathbf{G}(\alpha_1 A_1, \ldots, \alpha_K A_K) = (\alpha_1 \cdots \alpha_K)^{\frac{1}{K}} \mathbf{G}(A_1, \ldots, A_K), \qquad \alpha_1, \ldots, \alpha_K > 0.$$

3. Invariance under permutation: $\mathbf{G}(A_{\pi(1)}, \ldots, A_{\pi(K)}) = \mathbf{G}(A_1, \ldots, A_K)$ with π a permutation of $(1, \ldots, K)$.
4. Monotonicity: if $A_i \geq B_i$, for all i, then $\mathbf{G}(A_1, \ldots, A_K) \geq \mathbf{G}(B_1, \ldots, B_K)$.
5. Continuity from above: if for all fixed i, $A_i^{(j)}$ is a monotonically decreasing sequence of matrices converging to A_i^* for $j \to \infty$, then $\mathbf{G}(A_1^{(j)}, \ldots, A_K^{(j)})$ converges to $\mathbf{G}(A_1^*, \ldots, A_K^*)$.
6. Congruence invariance: for all invertible matrices $S \in \mathbb{R}^{n \times n}$,

$$\mathbf{G}(S^T A_1 S, \ldots, S^T A_K S) = S^T \mathbf{G}(A_1, \ldots, A_K) S.$$

7. Joint concavity:

$$\begin{aligned} &\mathbf{G}(\lambda A_1 + (1-\lambda)B_1, \ldots, \lambda A_K + (1-\lambda)B_K) \\ &\quad \geq \lambda \mathbf{G}(A_1, \ldots, A_K) + (1-\lambda)\mathbf{G}(B_1, \ldots, B_K), \qquad 0 < \lambda < 1. \end{aligned}$$

8. Invariance under inversion: $\mathbf{G}(A_1, \ldots, A_K) = \left(\mathbf{G}(A_1^{-1}, \ldots, A_K^{-1})\right)^{-1}$.
9. Determinant equality: $\det \mathbf{G}(A_1, \ldots, A_K) = (\det A_1 \cdots \det A_K)^{\frac{1}{K}}$.
10. Arithmetic-geometric-harmonic inequality:

$$\frac{1}{K}\sum_{i=1}^{K} A_i \geq \mathbf{G}(A_1, \ldots, A_K) \geq \left(\frac{1}{K}\sum_{i=1}^{K} A_i^{-1}\right)^{-1}.$$

Unfortunately, these properties do not result in a unique definition for the geometric mean of a general number of matrices. For the case of two matrices, however, the geometric mean is uniquely defined by properties 1 to 10 and given by the following expressions [7]

$$\mathbf{G}(A, B) = A(A^{-1}B)^{\frac{1}{2}} = A^{\frac{1}{2}}(A^{-\frac{1}{2}}BA^{-\frac{1}{2}})^{\frac{1}{2}}A^{\frac{1}{2}}. \tag{2.1}$$

Considering the manifold of symmetric positive definite $n \times n$ matrices $\mathbb{S}_+^n$, we can find another intuitively attractive interpretation of this result. We note that the intrinsic distance between $A, B \in \mathbb{S}_+^n$ (see Section 4.1) is given by

$$\delta(A, B) = ||\log(A^{-\frac{1}{2}}BA^{-\frac{1}{2}})||_F, \tag{2.2}$$

with $||\cdot||_F$ the Frobenius norm. Using this distance measure, we can determine the geodesic between A and B [7], i. e., the curve of shortest distance on the manifold between A and B, as

$$\begin{aligned} \gamma(t) &= A(A^{-1}B)^t = A^{\frac{1}{2}}(A^{-\frac{1}{2}}BA^{-\frac{1}{2}})^t A^{\frac{1}{2}} \\ &=: A\#_t B, \quad t \in [0, 1]. \end{aligned} \tag{2.3}$$

This shows that the geometric mean is exactly the midpoint on the geodesic (the notation in the last term will be used further in the text):

$$\mathbf{G}(A, B) = \gamma(\tfrac{1}{2}) = A\#_{\frac{1}{2}}B.$$

The subscript in the last term is often dropped when $t = \frac{1}{2}$.

3. Geometric means based on planar approaches. While the properties in the ALM list result in an explicit, unique definition for calculating the geometric mean of two matrices, this is not the case when dealing with more matrices. Considering the simplified case of a space with planar Euclidean geometry, the arithmetic mean of three matrices is the centroid of the triangle they form. There are various intuitively appealing techniques to determine this centroid that have been generalized to the non-planar (non-Euclidean) geometry of $\mathbb{S}_+^n$ [5, 9, 11, 31, 35], where the arithmetic mean is naturally generalized to the geometric mean. This generalization causes the need for the exact formulas for the centroid (in case of the NBMP and CHEAP mean) to be iterated. Throughout the different algorithms, we consistently notice a trade-off between the speed of convergence and the number of properties in the ALM list which the algorithms satisfy.

In this section, we discuss the geometric interpretation of the ALM [5], NBMP [11, 31], and CHEAP [9] mean and end with a comparison of these algorithms. Although one can obtain new geometric means by combining existing ones, in [35] it has been shown that there is no improvement for more than four matrices. Hence, we omit this approach in the survey.

All the numerical experiments were performed on an Intel® Core™i5-2540M CPU at 2.60 GHz with Matlab R2012a.

3.1. ALM mean. The ALM mean [5] is a geometric mean which, as the name implies, satisfies the desired properties enumerated in the ALM list. When taking the ALM mean of K matrices, recursion is used to define the iterations in which we replace $\left(A_1^{(j)},\ldots,A_K^{(j)}\right)$ by

$$\left(A_1^{(j+1)},\ldots,A_K^{(j+1)}\right)=\left(G_{ALM}((A_i^{(j)})_{i\neq 1}),\ldots,G_{ALM}((A_i^{(j)})_{i\neq K})\right),$$

where G_{ALM} denotes the recursively defined ALM mean of $K-1$ matrices with the known geometric mean of two matrices (2.1) as its base. In [5], all terms in these iterations are proven to converge towards the same limit and in Figure 3.1a, a planar simplification of this algorithm for three matrices is depicted.

3.2. NBMP mean. The NBMP mean [11, 31], just as the ALM mean, satisfies all properties in the ALM list. To compute the NBMP mean of K matrices, we use recursion to define the iterations in which we replace $\left(A_1^{(j)},\ldots,A_K^{(j)}\right)$ by $\left(A_1^{(j+1)},\ldots,A_K^{(j+1)}\right)$, which equals

$$\left(A_1^{(j)}\#_{\frac{K-1}{K}}G_{NBMP}((A_i^{(j)})_{i\neq 1}),\ldots,A_K^{(j)}\#_{\frac{K-1}{K}}G_{NBMP}((A_i^{(j)})_{i\neq K})\right),$$

where G_{NBMP} denotes the recursively defined NBMP mean of $K-1$ matrices with the geometric mean of two matrices (2.1) as its base. The notation from (2.3) was used to denote the point on the geodesic representing the weighted mean of the terms involved. In [11], all terms in these iterations are again proven to converge towards the same limit and in Figure 3.1b we show a simplified representation of how the algorithm operates for three matrices.

3.3. General class. We have encountered two means, both satisfying all properties in the ALM list but yielding different results as shown in the next example.

EXAMPLE 3.1. If we consider the matrices

$$\begin{bmatrix}25 & 4\\ 4 & 1\end{bmatrix},\begin{bmatrix}20 & 1\\ 1 & 1\end{bmatrix},\begin{bmatrix}1 & 1\\ 1 & 20\end{bmatrix},$$

the results for the ALM and NBMP algorithm are respectively

$$\begin{bmatrix}7.6943 & 0.9919\\ 0.9919 & 2.0528\end{bmatrix}\text{ and }\begin{bmatrix}7.7139 & 0.9719\\ 0.9719 & 2.0425\end{bmatrix},$$

ETNA
Kent State University
http://etna.math.kent.edu

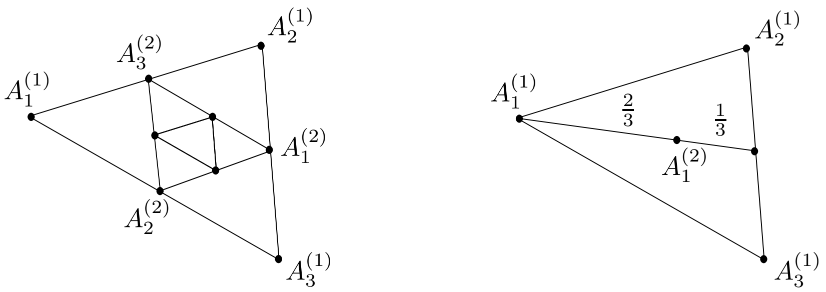

(a) The ALM mean. (b) The NBMP mean.

FIGURE 3.1. *Simplified representations of the algorithms for three matrices in a flat geometric space.*

which clearly shows that the results differ.

In fact, in [11, 25] it is shown that the ALM and NBMP mean are two instances of an entire class of means, all satisfying the required properties but with possibly different results. For K matrices, this *general mean* $G_{s_1,\dots,s_{K-1}}$ depends on $K-1$ parameters $(s_1,\dots,s_{K-1})$ and again recursion is used to define the iterations in which we replace $\left(A_1^{(j)},\dots,A_K^{(j)}\right)$ by $\left(A_1^{(j+1)},\dots,A_K^{(j+1)}\right)$ defined as

$$\left(A_1^{(j)}\#_{s_1}G_{s_2,\dots,s_{K-1}}((A_i^{(j)})_{i\neq 1}),\dots,A_K^{(j)}\#_{s_1}G_{s_2,\dots,s_{K-1}}((A_i^{(j)})_{i\neq K})\right).$$

For the ALM and NBMP mean, these parameters become respectively $(1,1,\dots,1,\frac{1}{2})$ and $(\frac{K-1}{K},\frac{K-2}{K-1},\dots,\frac{1}{2})$. This class illustrates that for a general number of matrices the geometric mean is not uniquely defined, not even starting from the ten desired properties. In Section 4, we investigate the Karcher mean, which also satisfies all properties but has a more appealing analogy with the arithmetic mean.

Currently, research is being conducted which indicates that all the means satisfying the properties in the ALM list have a close proximity to each other.

3.4. CHEAP mean. The CHEAP mean [9], unlike the previous algorithms, is no longer defined recursively. It also no longer satisfies all properties present in the ALM list, but as we will notice later, this will be compensated by its very cheap computational cost. The underlying idea is again to compute the centroid of a triangle (with vertices A, B, and C, see Figure 3.2) by the formula

$$A+\frac{1}{3}\big((B-A)+(C-A)\big).$$

The expression above can be interpreted as a step in an Euclidean space from vertex A into the direction of $\frac{1}{3}\left((B-A)+(C-A)\right)$, which is the arithmetic mean of the directions of vertex A to the three vertices A, B, and C (where the direction $A-A$ is trivially omitted). Generalizing the notions of a path (and consequently the direction) between two points and of taking a step into a certain direction to the manifold of positive definite matrices, we obtain the expression (see [9])

$$A\exp\left(\frac{1}{3}\left(\log(A^{-1}B)+\log(A^{-1}C)\right)\right).$$

ETNA
Kent State University
http://etna.math.kent.edu

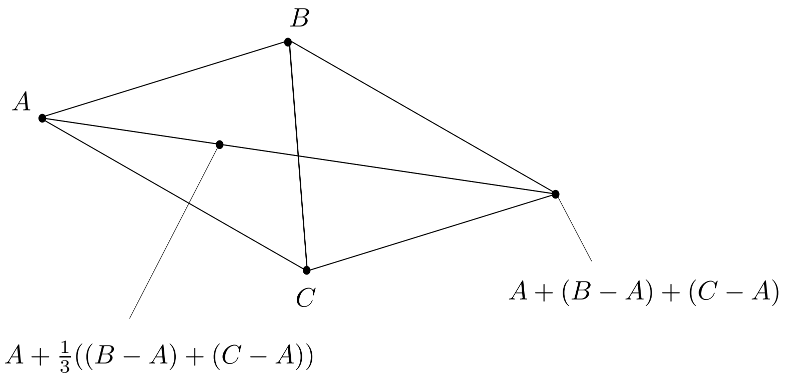

FIGURE 3.2. *Simplified representations of the CHEAP mean for three matrices.*

In the general case of K matrices, we replace in each iteration the matrices $(A_1^j, \dots, A_K^j)$ by $(A_1^{j+1}, \dots, A_K^{j+1})$, where

$$A_i^{j+1} = A_i^j \exp\left(\frac{1}{K} \sum_{\ell=1, \ell\neq i}^{K} \log\left(\left(A_i^j\right)^{-1} A_\ell^j\right)\right).$$

We iterate until convergence, although convergence is not always guaranteed for this algorithm, i. e., when the matrices are not sufficiently close to each other; see [9, Theorem 2.1] for an exact bound.

For the ALM and NBMP algorithms, the mean of two matrices is by definition known to be the analytical geometric mean, since they are recursively defined starting with this analytical expression. For the CHEAP mean, this consistency is less obvious but it is nonetheless still present. If we examine the CHEAP mean of two matrices by applying one iteration of the algorithm, we get

$$A_1^{(0)} \to A_1^{(1)} = A_1^{(0)} \exp\left(\frac{1}{2}\log\left((A_1^{(0)})^{-1}A_2^{(0)}\right)\right) = A_1^{(0)}\left((A_1^{(0)})^{-1}A_2^{(0)}\right)^{\frac{1}{2}},$$
$$A_2^{(0)} \to A_2^{(1)} = A_2^{(0)} \exp\left(\frac{1}{2}\log\left((A_2^{(0)})^{-1}A_1^{(0)}\right)\right) = A_2^{(0)}\left((A_2^{(0)})^{-1}A_1^{(0)}\right)^{\frac{1}{2}},$$

which are two equivalent expressions for the geometric mean of $A_1^{(0)}$ and $A_2^{(0)}$.

3.5. Comparison. In Figure 3.3a we show the required computational time of all the above algorithms as the number of 30×30 well-conditioned matrices in the mean increases. The random matrices throughout the paper are constructed in MATLAB as follows, with n being the size of the matrix, K the number of matrices, and f the order of magnitude of the condition number.

```
for i=1:K
  [Q,~]=qr(rand(n)); D=diag([[rand(1,n-1)+1],10^(-f)]);
  A{i}=Q*D*Q';
end
```

The stopping criterion for all three algorithms is satisfied when the difference between two consecutive iteration points becomes less than a specific tolerance.

While the ALM mean is proven to converge linearly [5] and the NBMP mean superlinearly of order 3 [11], both have rapidly increasing computational time as the number of matrices increases. The number of operations for both algorithms equals $\mathcal{O}(n^3 K! \prod_{i=3}^{K} p_i)$,

ETNA
Kent State University
http://etna.math.kent.edu

in which n denotes the size of the matrices, K the number of matrices, and p_i the average number of iterations required to compute the ALM and NBMP mean of i matrices. The advantage of the superlinear convergence of the NBMP algorithm over the linear convergence of the ALM algorithm is found in the p_i factors since they will be much smaller for the former. The problem for both, however, lies in the significant $K!$ factor, which grows tremendously fast as K increases. Despite the worse performance, it is still interesting to examine these means since they were the first algorithms devised to compute the matrix geometric mean of a general number of matrices.

For the CHEAP mean, however, the number of operations equals $\mathcal{O}(n^3K^2p_K)$, in which the factor K^2 is a vast improvement over $K!$. Of course, this increased speed of the CHEAP mean comes at a price. It no longer satisfies all properties in the ALM list and can therefore no longer be considered an actual geometric mean. We therefore compare the results of the different algorithms by taking the means of three 30×30 matrices, for which the condition number of each matrix is varied. In Figure 3.3b, the intrinsic distances (2.2) between the results are shown, and it is clear that the ALM and NBMP mean are more similar to each other than to the CHEAP mean, especially as the condition number of the matrices increases. However, the CHEAP mean can still be found to be in the vicinity of the other means when the condition number of the matrices is not too large.

A similar figure could be obtained by displaying the classical Euclidean distance between the results, but the distances are consistently smaller than for the intrinsic distance (2.2). This difference is explained by the fact that the intrinsic distance is measured on a curved manifold, while the Euclidean one measures the distance on a straight line in the enveloping vector space $\mathbb{S}^n$, the set of symmetric $n \times n$ matrices.

The accuracy of the methods is harder to verify since we need a reference solution to compare the results of the algorithms with. By constructing the matrices in the mean as described above with the same matrix Q, we obtain a set of simultaneously diagonalizable, commuting matrices. Of these we know the exact geometric mean using the first property in the ALM list, so we can use this as our reference solution. However, the CHEAP mean is shown [9] to converge in one iteration to the exact solution when the matrices commute. Hence, this test is only meaningful for the ALM and NBMP mean, of which we show the results in Figure 3.3c. The relative intrinsic distance

$$\frac{||\log(A^{-\frac{1}{2}}GA^{-\frac{1}{2}})||_F}{||G||_F}, \tag{3.1}$$

with A being the result of one of the algorithms and G the exact solution, is used to display their deviation for different condition numbers of the matrices. Recall that the numerator is the intrinsic distance (2.2) between A and G. The accuracy of both algorithms is very similar and deteriorates steadily as the condition number of the matrices increases. We note that when using the classical Euclidean distance, the deviations are almost at machine precision for all condition numbers, and this difference can again be explained by the curvature of the manifold.

For matrices which are not simultaneously diagonalizable, the exact geometric mean is not known, hence we need a different model solution to compare our results with. A high precision version of each of the algorithms is designed using the `vpa` functionality of MATLAB with 32 digits of accuracy. The relative intrinsic distance (3.1) between results from the original algorithms and the high precision versions is displayed in Figure 3.3d (using three random 10×10 symmetric positive definite matrices). Again the steady deterioration of the accuracy can be observed as the condition number of each of the matrices increases. However, it is clear that the CHEAP algorithm is more sensitive to this condition number than the ALM

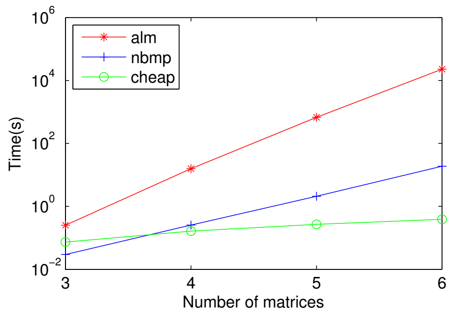

(a) Required time for different number of matrices.

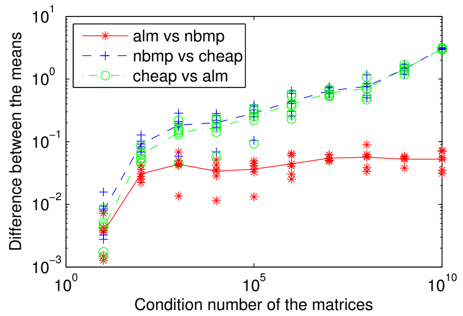

(b) Distance between the different means for the intrinsic measure (2.2). The mean of the samples is indicated by the connecting lines.

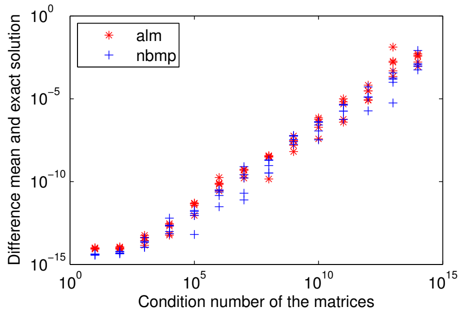

(c) Accuracy for different condition numbers of the three simultaneously diagonalizable matrices.

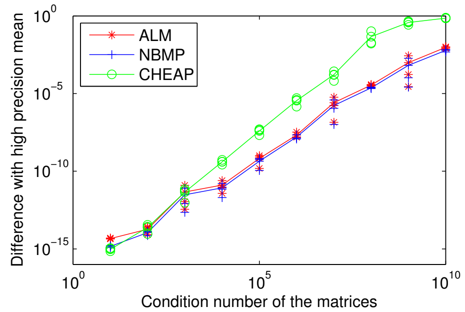

(d) Accuracy for different condition numbers of the three random matrices. The mean of the samples is indicated by the connecting lines.

FIGURE 3.3. *Comparison of the different planar algorithms.*

and NBMP algorithms.

4. The Karcher mean. As mentioned in the previous section, the properties in the ALM list fail to specify a unique definition for the geometric mean. The ALM and NBMP are only two examples of a general class of means satisfying all of them. Another mean that satisfies all the necessary properties is the Karcher mean. It is defined as the minimizer

$$\mathbf{K}(A_1, \ldots, A_K) = \operatorname*{argmin}_{X \in \mathbb{S}_+^n} \sum_{i=1}^{K} \delta^2(A_i, X),$$

where $\mathbb{S}_+^n$ represents the set of symmetric positive definite $n \times n$ matrices, $\delta(A, X)$ is the intrinsic distance on this manifold as given in (2.2), and A_i are the matrices of which we want to find the Karcher mean. In terms of an optimization problem, this translates to a cost function f

$$f(X) = \sum_{i=1}^{K} || \log(A_i^{-\frac{1}{2}} X A_i^{-\frac{1}{2}}) ||_F^2. \tag{4.1}$$

For this mean to be well-defined, the minimizer of the cost function should be unique. When the manifold $\mathbb{S}_+^n$ is endowed with its natural inner product (4.3) (see Section 4.1.2), the cost function f is strictly geodesically convex, which is a generalization of the classical convexity as follows: let $X, Y \in \mathbb{S}_+^n$, and $t \in [0, 1]$. Then

$$f(X \#_t Y) \leq (1 - t) f(X) + t f(Y).$$

Combining this with the convexity of the set $\mathbb{S}_+^n$ itself, the minimizer can be proven to be unique [7, 12].

The uniqueness of this minimizer can also be verified in a different manner. It is known from [15, 26] that $\mathbb{S}_+^n$ with the natural inner product (4.3) forms a Cartan–Hadamard manifold, which is a complete and simply connected Riemannian manifold with non-positive sectional curvature everywhere. On this type of manifold, the so-called Riemannian center-of-mass, which in this case is exactly the Karcher mean, is known to be unique [16, 22].

The Karcher mean satisfies all properties in the ALM list, of which the monotonicity has only very recently been proven; see [8, 23, 30]. Moreover, this mean is found to be appealing because of its analogy with the arithmetic mean, which can be seen as a similar minimizer by using the standard Euclidean distance.

Since the Karcher mean will be computed iteratively as the solution of an optimization problem, we need a good starting guess. We will use the CHEAP mean in all experiments, since it shows good computational speed and reasonable accuracy as discussed before.

4.1. Differential geometry. Calculating the Karcher mean involves solving an optimization problem on a manifold, which requires a more general approach than in the traditional case of vector spaces. We need to introduce some new concepts to perform this generalization, but we only briefly discuss these matters here; for a more thorough discussion of the subject we refer to any introductory book on differential geometry [14, 24] and to [1] for the optimization perspective. After the introduction, we will use these generalized concepts to implement a number of optimization techniques, more specifically, the steepest descent, conjugate gradient, trust region, and BFGS algorithms. The type of generalized optimization used here is often referred to as retraction-based optimization [1, 2], indicating that the concept of retractions (Section 4.1.3) lies at the foundation of these techniques.

The general concepts discussed here are found in [1], and many of these structures have already been derived for $\mathbb{S}_+^n$ endowed with its natural metric [18, 33, 38]. We add an explicit expression for the Levi–Civita connection (Section 4.1.4) on this manifold (and consequently for the Riemannian Hessian) and a derivation of all these structures for the manifold endowed with the inner product inherited from $\mathbb{S}^n$ (see Section 4.1.2).

4.1.1. Manifold and tangent space. So far we have been calling the space of positive definite $n \times n$ matrices $\mathbb{S}_+^n$ a manifold, without specifying what this means exactly. In order not to get too caught up in details, we give a more intuitive definition: a *manifold* is a set which can locally be mapped one-to-one to $\mathbb{R}^d$ (where d is the dimension of the manifold). In order to get a smooth (C^∞) manifold, we also require these mappings to transition smoothly onto each other in case their domains overlap. The space $\mathbb{S}_+^n$ is well-known to be a smooth manifold [29, 40].

Another important concept is the *tangent space* to a manifold in a certain point, which is basically a first-order (vector space) approximation of the manifold at this point. For $\mathbb{S}_+^n$, the tangent space at each point X, denoted by $T_X\,\mathbb{S}_+^n$, can be identified with the vector space of symmetric matrices $\mathbb{S}^n$

$$T_X\,\mathbb{S}_+^n \simeq \mathbb{S}^n\,.$$

Applying a tangent vector $\xi_X \in T_X\,\mathbb{S}_+^n$ at a point $X \in \mathbb{S}_+^n$ to a differentiable function $f : \mathbb{S}_+^n \to \mathbb{R}$ is defined to be

$$\xi_X f = \mathrm{D}f(X)[\xi_X],$$

where ξ_X on the right-hand side is simply a symmetric matrix and $\mathrm{D}f$ denotes the classical

Fréchet derivative of f. For the cost function f in (4.1), this differential is given by [7]

$$\mathrm{D}f(X)[\xi_X] = 2\sum_{i=1}^{K} \mathrm{tr}\left(X^{-1}\log(XA_i^{-1})\xi_X\right), \tag{4.2}$$

with A_i being the matrices in the mean and $\mathrm{tr}(\cdot)$ the matrix trace.

A *vector field* is a construction that associates with each point on the manifold a tangent vector in its tangent space. Suppose ξ is a vector field on a manifold and f is a real-valued function on this manifold, then ξf is again a real-valued function on the manifold defined by

$$\xi f : \mathbb{S}_+^n \to \mathbb{R} : X \mapsto \xi_X f.$$

We also apply a more general version of differentiation in this paper, namely that of functions between manifolds. This *differential* gives the change of the tangent vectors throughout the function.

4.1.2. Inner product and gradient. Gradient-based optimization requires the notions of a gradient and an inner product, which will be introduced here for $\mathbb{S}_+^n$. In fact, we consider two inner products. The first one is the inner product most frequently associated with $\mathbb{S}_+^n$: for $\xi_X, \eta_X \in T_X\,\mathbb{S}_+^n$, we have

$$\langle \xi_X, \eta_X \rangle_X^{\text{pd}} = \mathrm{tr}(\xi_X X^{-1}\eta_X X^{-1}), \tag{4.3}$$

which leads to the intrinsic distance measure (2.2) and geodesics of the form (2.3). Another benefit of this inner product is that the corresponding geodesics are complete, meaning that any geodesic segment can be extended indefinitely. For the second inner product we take the same as the enveloping space $\mathbb{S}^n$: suppose again $\xi_X, \eta_X \in T_X\,\mathbb{S}_+^n$, then

$$\langle \xi_X, \eta_X \rangle_X^{\text{sym}} = \mathrm{tr}(\xi_X \eta_X). \tag{4.4}$$

As a consequence, the intrinsic distance and expression of the geodesics become the same as in $\mathbb{S}^n$:

$$\begin{aligned} \delta^{\text{sym}}(A,B) &= ||B-A||_F, \\ \gamma^{\text{sym}}(t) &= A + t(B-A), \end{aligned} \tag{4.5}$$

with $A, B \in \mathbb{S}_+^n$. These geodesics are no longer infinitely extendable since it is possible for some A, B, t that the matrix $\gamma^{\text{sym}}(t)$ in (4.5) is no longer positive definite and thus not an element of $\mathbb{S}_+^n$. However, for sufficiently small t, $\gamma^{\text{sym}}(t)$ is in $\mathbb{S}_+^n$ and it appears to be computationally more efficient than the more involved expression γ^{pd} in (2.3).

Furthermore, the gradient of a cost function gives the direction of steepest ascent. It can be defined at each point X as the tangent vector $\mathrm{grad}\, f(X) \in T_X M$ such that

$$\langle \mathrm{grad}\, f(X), \xi_X \rangle_X = \mathrm{D}f(X)[\xi_X], \quad \forall \xi_X \in T_X M.$$

Using (4.2) we find for our current setting when using the inner product in (4.3)

$$\mathrm{grad}^{\text{pd}} f(X) = 2\sum_{i=1}^{K} X^{\frac{1}{2}} \log(X^{\frac{1}{2}} A_i^{-1} X^{\frac{1}{2}}) X^{\frac{1}{2}} = 2\sum_{i=1}^{K} X \log(A_i^{-1}X), \tag{4.6}$$

and

$$\mathrm{grad}^{\text{sym}} f(X) = 2\sum_{i=1}^{K} X^{-\frac{1}{2}} \log(X^{\frac{1}{2}} A_i^{-1} X^{\frac{1}{2}}) X^{-\frac{1}{2}} = 2\sum_{i=1}^{K} \log(A_i^{-1}X) X^{-1}, \tag{4.7}$$

when using (4.4). Note the slight difference between both expressions in the sign of the power of the outer X-factors.

From the geodesic convexity of f with the inner product (4.3), we know that the optimization problem, or equivalently $\mathrm{grad}^{\mathrm{pd}} f(X) = 0$, always admits a unique solution. Since (4.7) is obtained from (4.6) after a full-rank transformation, setting this gradient to zero results in the same unique solution.

4.1.3. Retraction and vector transport. Our optimization algorithms also require a map $R_X : T_X\,\mathbb{S}_+^n \to \mathbb{S}_+^n$ called *retraction* that locally maps $T_X\,\mathbb{S}_+^n$ onto the manifold $\mathbb{S}_+^n$ itself while preserving the first-order information of the tangent space in this point (see Figure 4.1a). This means that a step of size zero stays at the same point X and the differential of the retraction at this origin is the identity mapping. An interpretation of these retractions is that they take a unit step along a geodesic or an approximation thereof on the manifold into the direction specified by the argument. We consider three retractions:

$$R_X^{\mathrm{sym}}(\xi) = X + \xi, \tag{4.8}$$

$$R_X^{\mathrm{pd}}(\xi) = X^{\frac{1}{2}} \exp(X^{-\frac{1}{2}} \xi X^{-\frac{1}{2}}) X^{\frac{1}{2}}, \tag{4.9}$$

$$R_X^{\mathrm{pd'}}(\xi) = X + \xi + \frac{1}{2}\xi X^{-1}\xi. \tag{4.10}$$

Note that we omitted the subscript of the tangent vector $\xi \in T_X\,\mathbb{S}_+^n$ for clarity. The first of these is a unit step along the geodesic (4.5) and can thus be considered to be a natural retraction with respect to the inner product (4.4). When the manifold is endowed with the inner product (4.3), R_X^{sym} is a first-order retraction. As mentioned there, precaution has to be taken to assure that the result of the retraction is still positive definite. We do this by reducing our step size when necessary. The second one is the retraction that naturally arises when the manifold is endowed with the inner product (4.3). Recall that the geodesic between $A, B \in \mathbb{S}_+^n$ is given by

$$\begin{aligned} \gamma(t) &= A^{\frac{1}{2}} (A^{-\frac{1}{2}} B A^{-\frac{1}{2}})^t A^{\frac{1}{2}} \\ &= A^{\frac{1}{2}} \exp\left(t \log(A^{-\frac{1}{2}} B A^{-\frac{1}{2}})\right) A^{\frac{1}{2}}, \quad t \in [0, 1]. \end{aligned}$$

We obtain (4.9) as $\gamma(t)$ evaluated at $t = 1$ with $\xi_X = A^{\frac{1}{2}} \log(A^{-\frac{1}{2}} B A^{-\frac{1}{2}}) A^{\frac{1}{2}}$ and $A = X$. The last retraction is the second-order approximation to this second retraction, which can easily be seen by using the relation

$$\exp(X) = I + X + \frac{1}{2} X^2 + \mathcal{O}(X^3), \quad X \to 0.$$

Next, in order to perform, among others, the conjugate gradient algorithm, we need to somehow relate a tangent vector at some point $X \in \mathbb{S}_+^n$ to another point $Y \in \mathbb{S}_+^n$. This leads to the concept of a *vector transport* [1] (Figure 4.1b). We consider two vector transports: for $X \in \mathbb{S}_+^n, \xi_X, \eta_X \in T_X\,\mathbb{S}_+^n$,

$$\mathcal{T}_{\eta_X}^{\mathrm{sym}}(\xi_X) = \xi_X, \tag{4.11}$$

$$\mathcal{T}_{\eta_X}^{\mathrm{pd}}(\xi_X) = X^{\frac{1}{2}} \exp\left(\frac{X^{-\frac{1}{2}} \eta_X X^{-\frac{1}{2}}}{2}\right) X^{-\frac{1}{2}} \xi_X X^{-\frac{1}{2}} \exp\left(\frac{X^{-\frac{1}{2}} \eta_X X^{-\frac{1}{2}}}{2}\right) X^{\frac{1}{2}}, \tag{4.12}$$

where $\mathcal{T}_{\eta_X}^{(\cdot)}(\xi_X)$ denotes the vector transport of ξ_X over η_X. The definition of a vector transport [1, Definition 8.1.1] states that it has to be linear in ξ_X, and if η_X is the zero element,

ETNA
Kent State University
http://etna.math.kent.edu

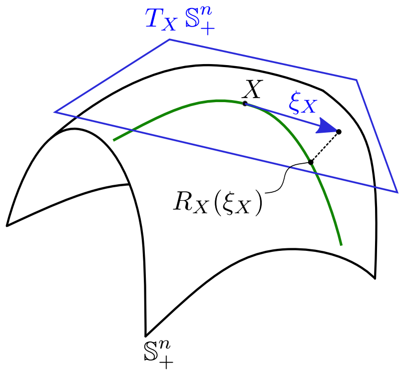

(a) Retraction.

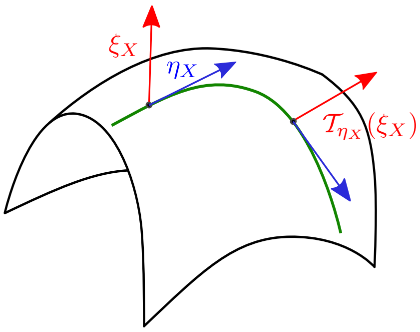

(b) Vector transport.

FIGURE 4.1. *Simplified representations of a retraction and a vector transport.*

the vector transport must be the identical mapping. Both conditions are easily checked for the expressions above. The definition also states that a vector transport has an associated retraction, meaning that the tangent vector $\mathcal{T}_{\eta_X}^{(\cdot)}(\xi_X)$ should be an element of the tangent space at $R_X(\eta_X)$, for some retraction R. The vector transport (4.11) is associated with the retraction R_X^{sym} in (4.8), since these structures naturally arise when $\mathbb{S}_+^n$ is endowed with the inner product (4.4) of the enveloping vector space $\mathbb{S}^n$. Such a natural vector transport is often referred to as parallel transport. The structure of the vector transport (4.12) suggests that it is associated with R_X^{pd} in (4.9), which is stated in [18]. Note that it is also possible to find a vector transport associated with $R_X^{\text{pd}'}$ in (4.10), but we decided to restrict our attention to the two more interesting vector transports mentioned above.

4.1.4. The Levi–Civita connection and the Riemannian Hessian. Some of our optimization methods require second-order information about the system, which is provided by the Hessian operator. The *Riemannian Hessian* of a real-valued function f at a point X on the manifold is a linear, symmetric mapping from the tangent space into itself given by

$$\operatorname{Hess} f(X) : T_X\,\mathbb{S}_+^n \to T_X\,\mathbb{S}_+^n : \xi_X \mapsto \operatorname{Hess} f(X)[\xi_X] = \nabla_{\xi_X} \operatorname{grad} f, \tag{4.13}$$

where ∇ is the so-called *Levi-Civita connection*, which depends on the inner product, hence the Hessian will also depend on the inner product.

When endowed with the inner product (4.4), the manifold is a dense Riemannian submanifold of $\mathbb{S}^n$, which is a vector space. Hence, the Levi–Civita connection is given by

$$\nabla_{\zeta_X}^{\text{sym}} \xi = \mathrm{D}(\xi)(X)[\zeta_X],$$

which is simply the derivative as in the vector space.

For the inner product (4.3), however, this connection is more complicated. It can be shown that

$$\nabla_{\zeta_X}^{\text{pd}} \xi = \mathrm{D}(\xi)(X)[\zeta_X] - \frac{1}{2}\left(\zeta_X X^{-1} \xi_X + \xi_X X^{-1} \zeta_X\right)$$

satisfies all properties of the Levi–Civita connection. A straightforward way to do this is by checking that it satisfies the Koszul formula [1], which at a point $X \in \mathbb{S}_+^n$ is given by

$$\begin{aligned} 2\langle \nabla_{\zeta_X}\eta, \xi_X\rangle_X = & \zeta_X\langle \eta, \xi\rangle + \eta_X\langle \xi, \zeta\rangle - \xi_X\langle \zeta, \eta\rangle \\ & - \langle \zeta_X, [\eta, \xi]_X\rangle_X + \langle \eta_X, [\xi, \zeta]_X\rangle_X + \langle \xi_X, [\zeta, \eta]_X\rangle_X. \end{aligned}$$

The actual computation of the Hessian will be discussed later for each second order method separately.

ETNA
Kent State University
http://etna.math.kent.edu

4.2. First-order implementations. We are ready to use all the building blocks of the previous section to assemble a number of optimization methods. They will range from simple first-order techniques such as steepest descent to more advanced second-order methods such as trust region algorithms. In this section, we start with discussing the first-order algorithms, while the second-order techniques are examined subsequently.

The methods in this section have already been derived in various papers [10, 18, 33, 38], both in their standard form as well as in some approximated manner. As an example of these approximated approaches, we found a Richardson-like iteration in [10], which is based on the standard steepest descent method with the inner product (4.3), except that a linearization of the natural retraction (4.9) is used. But this approximation is exactly the steepest descent algorithm using the retraction (4.8), hence the technique discussed in the paper can still be interpreted as a steepest descent technique on the manifold.

4.2.1. The steepest descent method. In a first attempt, we combine the elements of the previous section into the steepest descent algorithm, which takes in each iteration a step into the direction of $-\operatorname{grad} f(x)$, the direction of steepest descent. The step size is determined using Armijo line search [1], which is a standard backtracking technique, starting from step size 1 and iteratively multiplying it with a factor $\frac{1}{2}$ until an acceptable decrease of the cost function, relative to the step size, is obtained. Algorithm 4.1 contains the steepest descent method when setting the parameter $\beta = 0$. The convergence conditions checked in this algorithm are satisfied when the Armijo step size or the absolute or relative difference between two consecutive iterations are smaller than their respective tolerances. We consider all three retractions (4.8), (4.9), and (4.10).

4.2.2. The conjugate gradient method. In Figure 4.2a, we show the typical zigzag-pattern that arises for the steepest descent method. The pattern was activated by using five 3×3 random matrices for which the smallest eigenvalue is roughly 10^3 times smaller than the others. Algorithm 4.1 shows the conjugate gradient algorithm which helps us to deal with this problem, as can be seen in Figure 4.2b. The amount of influence of the previous search direction in the conjugate gradient algorithm is determined by the β factor, for which we consider three different formulas given in [32] and denoted by $\beta^{(\mathrm{fr})}$ (Fletcher–Reeves), $\beta^{(\mathrm{pr})}$ (Polak–Ribière), and $\beta^{(\mathrm{hs})}$ (Hestenes–Stiefel). As mentioned, this previous search direction is transported between different tangent spaces by vector transports. To get the best affinity between the vector transports and the retractions, we work with the natural retractions R^{sym} (with corresponding vector transport $\mathcal{T}^{\mathrm{sym}}$) and R^{pd} (with $\mathcal{T}^{\mathrm{pd}}$) here.

4.2.3. Comparisons. When comparing the overall performance of the steepest descent and conjugate gradient algorithms, we notice that the influence of choosing the inner products (4.3) or (4.4) is far greater than the impact of the chosen retraction or, in case of the conjugate gradient algorithm, the β type. In fact, when only varying the retraction and the β type, the results are all very similar. The speed-up of the conjugate gradient technique over steepest descent is also hardly noticeable in general, which is explained by the presence of a sufficiently good initial guess, the CHEAP mean. When this initial point is sufficiently close to the solution of the problem, the cost function will behave nicely in this neighbourhood and the zigzag-pattern mentioned before is less likely to occur.

To investigate the accuracy of these first order methods, the results of the algorithms are compared with a high precision computation of the Karcher mean (again using the `vpa` functionality of MATLAB with 32 digits of accuracy). In Figure 4.3a, the relative intrinsic distance (3.1) between the results of the algorithms (for three 10×10 matrices) and this high precision solution is given as function of the condition number of the matrices. The algorithms using the natural inner product (4.3) show in general better results than those

ETNA
Kent State University
http://etna.math.kent.edu

Algorithm 4.1 The Karcher mean using the conjugate gradient method.

Input: matrices $A_1, \ldots, A_K$, $K > 2$, initial guess X_0, retraction and vector transport type R^i and $\mathcal{T}^i$ (R^{sym} and $\mathcal{T}^{\mathrm{sym}}$ or R^{pd} and $\mathcal{T}^{\mathrm{pd}}$), β type β^j ($\beta^{(\mathrm{fr})}$, $\beta^{(\mathrm{pr})}$ or $\beta^{(\mathrm{hs})}$ or 0)
Output: Karcher mean $\mathbf{K}(A_1, \ldots, A_K)$
$k \leftarrow 0$
$grad_k \leftarrow \mathrm{grad}\, f(X_k)$ $\{$ f is our cost function $\}$
$\xi_k \leftarrow -grad$ $\{\xi_k$ is the search direction$\}$
while not converged **do**
  $X_{k+1} \leftarrow R^i_{X_k}(t^A \xi_k)$ $\{$ with t^A the Armijo step size [1]$\}$
  $grad_{k+1} \leftarrow \mathrm{grad}\, f(X_{k+1})$
  $\xi_{old} \leftarrow \mathcal{T}^i_{t^A \xi_k}(\xi_k)$ $\{$vector transport of the old search direction$\}$
  Determine β according to the given type
  $\xi_{k+1} \leftarrow -grad_{k+1} + \beta \xi_{old}$
  if ξ_{k+1} not a descent direction **then**
    $\xi_{k+1} \leftarrow -grad_{k+1}$
  end if
  $k \leftarrow k + 1$
end while

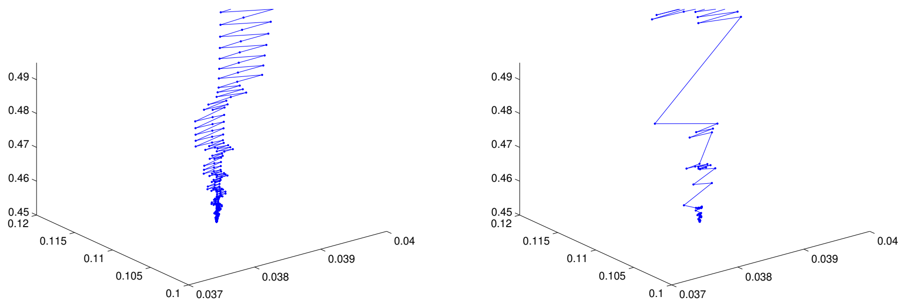

(a) Using steepest descent. (b) Using conjugate gradient.

FIGURE 4.2. *Evolution of the eigenvalues of the consecutive iteration points in calculating the Karcher mean of five* 3×3 *matrices for which the zigzag-pattern appears for steepest descent.*

using (4.4).

The speed of the algorithms is tested both for an increasing number of matrices (with the size fixed to 10×10 matrices) and for varying sizes of the matrices in the mean (with the number of matrices fixed to five). Again we notice the advantage of the structures associated with the inner product (4.3), as these result in the algorithms requiring less iterations. The computational cost of these structures, however, is usually higher than those related to the inner product (4.4), which causes the overall computational time to be very similar. In Figure 4.3b, the number of iterations is displayed for the algorithms as a function of the size of the matrices, which clearly shows the distinction between the two inner products.

We can conclude that structures associated with the inner product (4.3) are best suited for our problem since the resulting algorithms require less iterations. It therefore seems interesting to consider the performance of a steepest descent algorithm when the inner product

$$\langle \xi_X, \eta_X \rangle_X^{(\alpha)} = \mathrm{tr}\left(\xi_X X^{-\alpha} \eta_X X^{-\alpha}\right) \tag{4.14}$$

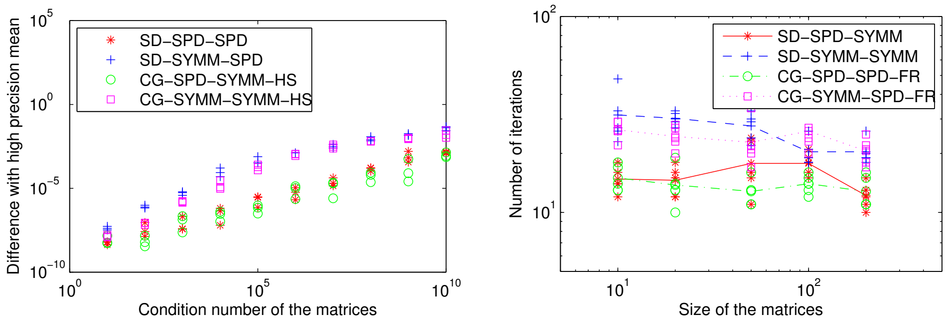

(a) Accuracy for different condition numbers.

(b) Required iterations for different sizes. The mean of the samples is indicated by the connecting lines.

FIGURE 4.3. *Comparison of the accuracy and computational speed of Algorithm 4.1 using the steepest descent and conjugate gradient algorithm. In the legends, we first indicate whether the Steepest Descent (SD) or Conjugate Gradient (CG) technique is used, next, which of the inner products* (4.3) *(SPD) or* (4.4) *(SYMM) is used, and finally whether retraction R_X^{sym} (SYMM) or R_X^{pd} (SPD) was taken. In case of CG, we also indicate which β type is used.*

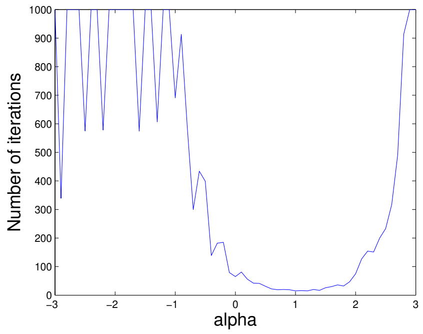

FIGURE 4.4. *Number of iterations required for a steepest descent algorithm when using structures related to the inner product* (4.14) *for different values of α.*

and its related structures are used, which reduces to the previous cases when $\alpha = 0$ or $\alpha = 1$. In Figure 4.4, we display the number of iterations that such an algorithm requires for different values of α. The figure displayed has been constructed using the retraction R_X^{pd}, but a nearly identical figure was obtained using R_X^{sym}, again indicating that the influence of the inner product is far greater than that of the retraction. It is obvious from the figure that the inner product (4.3) is still most natural to the manifold, since this is the one corresponding to $\alpha = 1$.

4.3. Second-order implementations. The goal of second-order optimization techniques is to use (an approximation of) the Hessian of the cost function to obtain a quadratically (or at least superlinearly) convergent algorithm. In the following, we discuss a number of attempts to accomplish this and compare their performance. This Hessian has, however, typically a higher computational complexity, which means it is yet to be determined whether these algorithms are more efficient than the first-order techniques. In the discussions of the trust region method, we focus on the computation of the Hessian while the actual implementation of the method is performed using Algorithms 10 and 11 from [1].

In the existing literature, the Riemannian Hessian is computed as it is defined in Section 4.3.2 [17, 38]. We derive the Hessian according to the classical definition (4.13) (using the Levi–Civita connection) in Section 4.3.1 and an approximation in Section 4.3.3. The

Riemannian BFGS method in Section 4.3.4 is the application of the existing generalized algorithm [36, 39] to $\mathbb{S}^n_+$.

4.3.1. A trust region method: exact Hessian. In Section 4.1.4, we derived all the components needed to determine the Hessian of our cost function f and noticed that the result again depends on the inner product. Hence using the definition of the Hessian for the inner product (4.4) with ∇^{sym} gives

$$\mathrm{Hess}^{\mathrm{sym}} f(X)[\xi_X] = \mathrm{D}(\mathrm{grad}^{\mathrm{sym}} f)(X)[\xi_X].$$

Using (4.7) leads to

$$\begin{aligned}\mathrm{Hess}^{\mathrm{sym}} f(X)[\xi_X] =& 2\sum_{i=1}^{K} \mathrm{D}(\log)(A_i^{-1}X)[A_i^{-1}\xi_X]X^{-1} \\ & - 2\sum_{i=1}^{K} \log(A_i^{-1}X)X^{-1}\xi_X X^{-1},\end{aligned}$$

where we used the product and chain rules of differentiation as well as the differential of the matrix inverse function. We also recognize the differential of the matrix logarithm function at the point $A_i^{-1}X$ into the direction of the tangent vector $A_i^{-1}\xi_X$, which is computed using the algorithm presented in [3, 4, 21].

When using the inner product (4.3) with ∇^{pd}, the Hessian becomes

$$\begin{aligned}\mathrm{Hess}^{\mathrm{pd}} f(X)[\xi_X] =& \mathrm{D}(\mathrm{grad}^{\mathrm{pd}} f)(X)[\xi_X] \\ & - \frac{1}{2}\left(\xi_X X^{-1}\,\mathrm{grad}^{\mathrm{pd}} f(X) + \mathrm{grad}^{\mathrm{pd}} f(X)X^{-1}\xi_X\right).\end{aligned} \tag{4.15}$$

In this case using (4.6) leads to

$$\begin{aligned}\mathrm{Hess}^{\mathrm{pd}} f(X)[\xi_X] =& 2\sum_{i=1}^{K} \xi_X \log(A_i^{-1}X) + 2\sum_{i=1}^{K} X\,\mathrm{D}(\log)(A_i^{-1}X)[A_i^{-1}\xi_X] \\ & - \left(\sum_{i=1}^{K} \xi_X \log(A_i^{-1}X) + \sum_{i=1}^{K} \log(XA_i^{-1})\xi_X\right) \\ =& \sum_{i=1}^{K} \xi_X \log(A_i^{-1}X) - \sum_{i=1}^{K} \log(XA_i^{-1})\xi_X \\ & + 2\sum_{i=1}^{K} X\,\mathrm{D}(\log)(A_i^{-1}X)[A_i^{-1}\xi_X],\end{aligned} \tag{4.16}$$

where again we need the differential of the matrix logarithm. Note that the first two terms in (4.16) are each others transpose (except for the minus sign), which can be exploited in the computation. At first sight this seems somewhat peculiar since this subtraction produces a skew-symmetric matrix while the result of the Hessian is supposed to be symmetric. However, looking at (4.15), we can see that the differential of the gradient is symmetric and the second part, as the sum of a matrix and its transpose, is symmetric as well, proving that $\mathrm{Hess}^{\mathrm{pd}} f(X)[\xi_X]$ is an element of $\mathbb{S}^n$.

ETNA
Kent State University
http://etna.math.kent.edu

4.3.2. A trust region method: Hessian by decomposition. Another way to compute the Hessian of our cost function is described in [17] and more explicitly for the current cost function in [38]. The proposed procedure consists of determining the components of a decomposition and then combining them to form the actual Hessian. A downside, however, is that when determining the Karcher mean of K $n \times n$ matrices, this sum consists of $Kn(n+1)/2$ terms, which number grows rapidly as n increases (in [38] only 3×3 matrices were considered). Two other important remarks are that this technique is derived for the manifold endowed with the inner product (4.3), and that the computation of the terms as in [38] is only valid when we take the Hessian at the identity matrix $X = I$. This causes the need to translate the problem in each iteration step to ensure this position for the current iteration point, which is done by applying the mapping $Y \mapsto X_k^{-\frac{1}{2}} Y X_k^{-\frac{1}{2}}$, in which X_k is the newly found iteration point, to all matrices in the mean and to X_k itself. After convergence, we apply the inverse mapping to translate the identity matrix to the actual Karcher mean of the original matrices. Theoretically, this need for a translation is not a downside, in fact, it can even simplify notations and required structures. Computationally, however, it could cause problems when working, e. g., with ill-conditioned matrices.

The expression for the Hessian is given by

$$\begin{aligned}\operatorname{Hess}^{\mathrm{pd}} f(X)[\xi_X] =& 2\sum_{l=1}^{m} \langle \xi_X, E_l\rangle_X w_l(1) E_l \\ =& 2\sum_{l=1}^{m} \operatorname{tr}(\xi_X X^{-1} E_l X^{-1}) w_l(1) E_l,\end{aligned}$$

where $m = Kn(n+1)/2$, E_l, and w_l are defined as in [38]. Remember that the decomposition is only valid at $X = I$, hence the expression can be further simplified to

$$\operatorname{Hess}^{\mathrm{pd}} f(X)[\xi_X] = 2\sum_{l=1}^{m} \operatorname{tr}(\xi_X E_l) w_l(1) E_l. \tag{4.17}$$

When computing the Hessian, the rank-1 structure of E_l can be exploited to limit the required amount of operations. However, if we compare the cost to compute the Hessians (4.16) and (4.17), which are theoretically equal, we still find the number of operations for the first one to be $\mathcal{O}(n^3)$ and that of the second $\mathcal{O}(n^4)$.

4.3.3. A trust region method: Hessian by approximation. In the calculations of the exact Hessian, determining the differential of the matrix logarithm function appears to be a serious computational cost. A perhaps less elegant but sometimes advantageous solution is to replace the matrix logarithm by (a truncation of) its Taylor series at $X = I$ given by

$$\sum_{m=1}^{\infty} \frac{(-1)^{m+1}}{m}(X - I)^m,$$

which converges to $\log(X)$ for all $\rho(X - I) < 1, X \neq 0$, where ρ denotes the spectral radius. Truncating this series after the second term and entering the result into the expressions for the gradient (4.6) and (4.7) corresponding to inner products (4.3) and (4.4), we obtain

$$\begin{aligned}\operatorname{grad}_2^{\mathrm{pd}} f(X) &= 2\sum_{i=1}^{K}(2XA_i^{-1}X - \frac{3}{2}X - \frac{1}{2}XA_i^{-1}XA_i^{-1}X) \text{ and} \\ \operatorname{grad}_2^{\mathrm{sym}} f(X) &= 2\sum_{i=1}^{K}(2A_i^{-1} - \frac{3}{2}X^{-1} - \frac{1}{2}A_i^{-1}XA_i^{-1}),\end{aligned}$$

ETNA
Kent State University
http://etna.math.kent.edu

respectively, as an approximation to the gradient. Note that we only use this approximation to derive the approximated Hessian, while the actual gradient will be used in the final algorithm. Applying the definition of the Hessian to these expressions results in

$$\begin{aligned}
\mathrm{Hess}_2^{\mathrm{pd}} f(X)[\xi_X] &= \mathrm{D}(\mathrm{grad}_{spd,2} f)(X)[\xi_X] \\
&\quad - \frac{1}{2}\left(\xi_X X^{-1} \mathrm{grad}_2^{\mathrm{pd}} f(X) + \mathrm{grad}_2^{\mathrm{pd}} f(X) X^{-1} \xi_X\right) \\
&= \sum_{i=1}^{K} \Big(2XA_i^{-1}\xi_X + 2\xi_X A_i^{-1} X - \frac{1}{2}\xi_X A_i^{-1} X A_i^{-1} X \\
&\quad - XA_i^{-1}\xi_X A_i^{-1} X - \frac{1}{2} X A_i^{-1} X A_i^{-1} \xi_X\Big), \\
\mathrm{Hess}_2^{\mathrm{sym}} f(X)[\xi_X] &= \mathrm{D}(\mathrm{grad}_2^{\mathrm{sym}} f)(X)[\xi_X] \\
&= \sum_{i=1}^{K} \left(3X^{-1}\xi_X X^{-1} - A_i^{-1}\xi_X A_i^{-1}\right).
\end{aligned}$$

The finite spectral radius of the Taylor series indicates that when the matrices in the mean lie close to each other (and to the current estimate), a trust region method using one of these Hessians is expected to work well. However, since these are only approximations to the Hessian, quadratic convergence of the algorithm is no longer guaranteed as will be shown in Section 4.3.5.

4.3.4. A Riemannian BFGS method. Finally, we test a Riemannian generalization of the classical BFGS method [36, 39], which updates an estimation of the Hessian throughout the algorithm instead of solving systems with the Hessians exactly. The main point of interest is how this update is done. Suppose we know the estimate $\mathcal{B}_k$ in iteration step k, which is assumed to be a linear operator from $T_{X_k}\mathbb{S}_+^n$ onto itself and can thus be represented by an $(n(n+1)/2)\times(n(n+1)/2)$ matrix B_k. The linear operator $\mathcal{B}_{k+1} : T_{X_{k+1}}\mathbb{S}_+^n \to T_{X_{k+1}}\mathbb{S}_+^n$ is then defined by

$$\mathcal{B}_{k+1}p = \widetilde{\mathcal{B}}_k p - \frac{\langle s_k, \widetilde{\mathcal{B}}_k p\rangle_{X_{k+1}}}{\langle s_k, \widetilde{\mathcal{B}}_k s_k\rangle_{X_{k+1}}}\widetilde{\mathcal{B}}_k s_k + \frac{\langle y_k, p\rangle_{X_{k+1}}}{\langle y_k, s_k\rangle_{X_{k+1}}} y_k, \qquad \forall p \in T_{X_{k+1}}\mathbb{S}_+^n \tag{4.18}$$

$$\widetilde{\mathcal{B}}_k = \mathcal{T}_{t^A\eta_k} \circ \mathcal{B}_k \circ \left(\mathcal{T}_{t^A\eta_k}\right)^{-1}, \tag{4.19}$$

in which η_k is the current search direction, t^A is the Armijo step size, s_k is the vector transport of the search direction to the new iteration point, y_k is a measure for the change of the gradient over the iteration step (formal expressions can be found in Algorithm 4.2), and $\mathcal{T}$ is the vector transport. Since an inner product and a vector transport are present in these expressions, there are again different situations to investigate. We will test the algorithm for each of the two inner products combined with their natural retraction and vector transport. Note that in order to evaluate $\mathcal{B}_{k+1}p$ using the matrix representation B_{k+1}, we need an $n(n+1)/2$ vector representation of p. This is done by the half-vectorization operator vech, which stacks the elements of the upper triangular part of p columnwise. Using this representation, the matrix-vector product $B_{k+1}\,\mathrm{vech}(p)$ returns an $(n(n+1)/2)$ vector which is the half-vectorization of the matrix $\mathcal{B}_{k+1}p$.

ETNA
Kent State University
http://etna.math.kent.edu

In the simpler case of the inner product (4.4), expressions (4.18) and (4.19) become

$$\begin{aligned} B_{k+1}\operatorname{vech}(p) =& \widetilde{B}_k \operatorname{vech}(p) - \widetilde{B}_k \operatorname{vech}(s_k)\frac{\operatorname{tr}(s_k\widetilde{\mathcal{B}}_k p)}{\operatorname{tr}(s_k\widetilde{\mathcal{B}}_k s_k)} \\ &+ \operatorname{vech}(y_k)\frac{\operatorname{tr}(y_k p)}{\operatorname{tr}(y_k s_k)}, \qquad \forall p \in T_{X_{k+1}}\mathbb{S}^n_+, \\ \widetilde{\mathcal{B}}_k =& \mathcal{B}_k. \end{aligned}$$

To remove p from this expression, the matrix traces have to be split up by using the property $\operatorname{tr}(AB) = \operatorname{vec}(A)^T \operatorname{vec}(B)$ with A and B being symmetric matrices. For the expression above, however, the second matrix should be half-vectorized, which needs to be compensated for in the first vectorization. To this end, we also change the vectorization of the first matrix to half-vectorization, but with the adaptation that each off-diagonal element is doubled and this operation is denoted by vech_2. This yields $\operatorname{tr}(AB) = \operatorname{vech}_2(A)^T \operatorname{vech}(B)$ and our update formula becomes

$$\begin{aligned} B_{k+1} =& B_k - \frac{1}{\operatorname{tr}(s_k\mathcal{B}_k s_k)} B_k \operatorname{vech}(s_k)\operatorname{vech}_2(s_k)^T B_k \\ &+ \frac{1}{\operatorname{tr}(y_k s_k)}\operatorname{vech}(y_k)\operatorname{vech}_2(y_k)^T. \end{aligned} \tag{4.20}$$

Note that this update differs in only two rank-1 terms from the previous estimation of the Hessian, which can be exploited in the implementation.

For the inner product (4.3), the calculation of $\widetilde{\mathcal{B}}_k$ is no longer so straightforward. Entering the vector transport $\mathcal{T}^{\mathrm{pd}}$ (4.12) into equation (4.19), we obtain

$$\widetilde{\mathcal{B}}_k p = Q\mathcal{B}_k(Q^{-1}pQ^{-T})Q^T, \qquad Q = X_k^{\frac{1}{2}} \exp\left(\frac{X_k^{-\frac{1}{2}} t^A \eta_k X_k^{-\frac{1}{2}}}{2}\right) X_k^{-\frac{1}{2}}.$$

To extract p from this expression, we want to use the property $\operatorname{vec}(ABC) = (C^T \otimes A)\operatorname{vec}(B)$ for general matrices A, B, and C, with $\otimes$ the Kronecker product, but this property cannot be used when half-vectorization is applied. Therefore, let us pretend for a moment that $\mathcal{B}_k$ and $\widetilde{\mathcal{B}}_k$ are represented by $n^2 \times n^2$ matrices and apply this rule to the above expression:

$$\begin{aligned} \widetilde{B}_k \operatorname{vec}(p) &= \operatorname{vec}(Q\mathcal{B}_k(Q^{-1}pQ^{-T})Q^T) \\ &= (Q \otimes Q)B_k \operatorname{vec}(Q^{-1}pQ^{-T}) \\ &= (Q \otimes Q)B_k (Q^{-1} \otimes Q^{-1})\operatorname{vec}(p). \end{aligned}$$

Changing back to half-vectorization can be accomplished by using the so-called duplication and elimination matrices D_n and E_n, which are simple matrices for transforming respectively a half-vectorization into a normal vectorization and vice versa. Equation (4.18) can be tackled in the same fashion as before, where we only need to pay attention to the extra factors in the current inner product. The total update procedure now becomes

$$\begin{aligned} \widetilde{B}_k =& E_n(Q\otimes Q)D_n B_k E_n (Q^{-1}\otimes Q^{-1})D_n, \\ B_{k+1} =& \widetilde{B}_k - \frac{1}{\operatorname{tr}(X_{k+1}^{-1}s_k X_{k+1}^{-1}\widetilde{\mathcal{B}}_k s_k)}\widetilde{B}_k \operatorname{vech}(s_k)\operatorname{vech}_2(X_{k+1}^{-1}s_k X_{k+1}^{-1})^T \widetilde{B}_k \\ &+ \frac{1}{\operatorname{tr}(X_{k+1}^{-1}y_k X_{k+1}^{-1}s_k)}\operatorname{vech}(y_k)\operatorname{vech}_2(X_{k+1}^{-1}y_k X_{k+1}^{-1})^T. \end{aligned} \tag{4.21}$$

Note that there are again two rank-1 terms present in the update. However, the update from B_k to $\widetilde{B}_k$ is no longer the identity and can in this case be seen as an update of the rank-1 terms which originated in the previous iterations.

Now that the techniques to update the estimate of the Hessian are specified, we display the entire Riemannian BFGS method in Algorithm 4.2. The convergence criteria are specified as before by testing whether the Armijo step size or the absolute or relative difference between two consecutive iterates are smaller than their respective tolerances.

Algorithm 4.2 The Karcher mean using the Riemannian BFGS method.

Input: matrices $A_1, \ldots, A_K$, $K > 2$, initial guess X_0, initial Hessian approximation B_0, retraction and vector transport type R^i and $\mathcal{T}^i$ (R^{sym} and $\mathcal{T}^{\text{sym}}$ or R^{pd} and $\mathcal{T}^{\text{pd}}$)
Output: Karcher mean $\mathbf{K}(A_1, \ldots, A_K)$

  $k \leftarrow 0$
  $grad_k \leftarrow \operatorname{grad} f(X_k)$ $\{f$ is our cost function$\}$
  while not converged **do**
    Obtain η_k: Solve the system: $B_k \operatorname{vech}(\eta_k) = -\operatorname{vech}(grad)$ $\{$The search direction$\}$
    if η_k not a descent direction **then**
      $\eta_k \leftarrow -grad_k$
    end if
    $X_{k+1} \leftarrow R^i_{X_k}(t^A \eta_k)$ $\{$with t^A the Armijo step size [1]$\}$
    $grad_{k+1} \leftarrow \operatorname{grad} f(X_{k+1})$
    $s_k \leftarrow \mathcal{T}^i_{t^A \eta_k}(t^A \eta_k)$
    $y_k \leftarrow grad_{k+1} - \mathcal{T}^i_{t^A \eta_k}(grad_k)$
    Update B_k to B_{k+1} using (4.20) or (4.21), depending on the inner product
    $k \leftarrow k + 1$
  end while

4.3.5. Comparison. We start by applying the second-order techniques to three 10×10 matrices and compare the results to a high precision computation of the Karcher mean, which is displayed in Figure 4.5a. The trust region methods from Sections 4.3.1 and 4.3.2 show in general better accuracy results than the previous first-order techniques, even though there are occasional outliers that shift the mean values in Figure 4.5a. The Riemannian BFGS method on the other hand displays an accuracy similar to the first-order methods. The trust region method using the approximated Hessian from Section 4.3.3 shows very bad accuracy, which is why we will not discuss this method any further. Overall, we again noticed slightly better results for the techniques corresponding to the inner product (4.3).

To test the speed of the algorithms, the size of the matrices in the mean is again varied (when taking 5 matrices) as well as the number of matrices (where we fix the size to 10×10 matrices). We notice once more the smaller number of iterations required by techniques based on the inner product (4.3), although the corresponding structures will in general be more expensive to compute. In Figure 4.5b, the number of iterations are shown for some of the methods as the number of matrices in the mean varies, where the difference in the iterations for the Riemannian BFGS methods confirms our statement. The advantage of the smaller number of iterations when using the inner product (4.3) can be seen when comparing the computational time of the two resulting trust region methods (see Figure 4.6a: TR-SPD and TR-SYMM). In fact, the trust region method using the exact Hessian corresponding to the inner product (4.3) (TR-SPD) gives very good results, approaching even the computational time of the first-order methods.

Another remarkable result in Figure 4.6a is the performance of the Riemannian BFGS

ETNA
Kent State University
http://etna.math.kent.edu

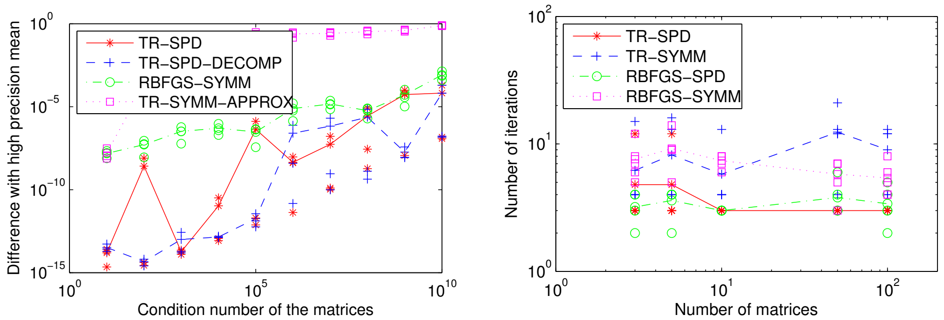

(a) Accuracy for different condition numbers of the matrices.

(b) Required iterations for a different number of matrices.

FIGURE 4.5. *Comparison of the accuracy and computational speed of the Karcher mean using the various trust region methods and the Riemannian BFGS algorithm. In the legends, we first indicate whether the trust region (TR) or Riemannian BFGS (RBFGS) technique is used, next, which of the inner products (4.3) (SPD) or (4.4) (SYMM) is used and if a non-standard method was used (DECOMP for the technique in Section 4.3.2 and APPROX for those in Section 4.3.3). The mean of the samples is in both figures indicated by the connecting lines.*

method, which shows a lower computational time than the steepest descent algorithm. We do note that this test is performed for a varying number of 10×10 matrices, and as the size of the matrices starts to increase, the Riemannian BFGS method as well as the other second-order techniques are outperformed by the steepest descent and conjugate gradient algorithm (see Figure 4.6b). Figure 4.6a also shows that the Riemannian BFGS method based on the inner product (4.3) is faster than the one using the inner product (4.4), even though the corresponding structures are more expensive. As Figure 4.5b indicates, the number of iterations is significantly lower for the first, which makes up for this extra computational cost. However, we note again that as the size of the matrices increases, the method using the inner product (4.4) becomes faster due to its less expensive update formula.

Furthermore, we mention that while the technique in Section 4.3.2 theoretically works with the same Hessian as the classical method using the standard definition of the Hessian (4.15) (for the inner product (4.3)), the number of iterations required for this first technique is generally larger than for the second, indicating a lower stability and causing a longer computational time.

Figure 4.6c shows the evolution of the gradient for all algorithms based on the inner product (4.3). The quadratic convergence of the trust region algorithm is clearly visible as well as a superlinear convergence for the Riemannian BFGS method. The steepest descent and conjugate gradient algorithm show a very similar (linear) convergence since the problem is well-behaved, eliminating the need for the conjugate gradient technique to be activated. Finally, the trust region algorithm using the approximated Hessian (Section 4.3.3) has lost all quadratic convergence and displays an even slower convergence than the steepest descent method. We note that the techniques that use the Armijo line search technique to determine the next iteration point stop when the norm of the gradient is about the square root of the machine precision. This is caused by the use of the squared norm of the gradient in the Armijo condition.

5. Conclusions. This paper has demonstrated various techniques to compute a matrix geometric mean. The lack of uniqueness of the definition was mostly overcome by the appealing analogy of the Karcher mean with the arithmetic mean. The convergence of the first-order optimization techniques for computing this Karcher mean can easily be verified using Corollary 4.3.2 and Theorems 4.4.1 and 4.4.2 in [1] and exploiting the convexity of the

ETNA
Kent State University
http://etna.math.kent.edu

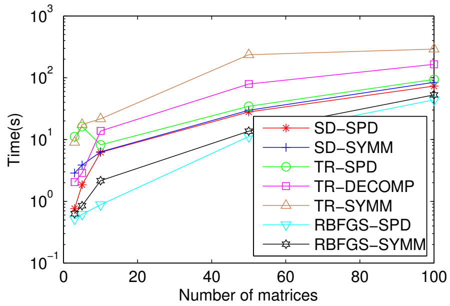

(a) Required time for different methods using the two inner products for a varying number of matrices.

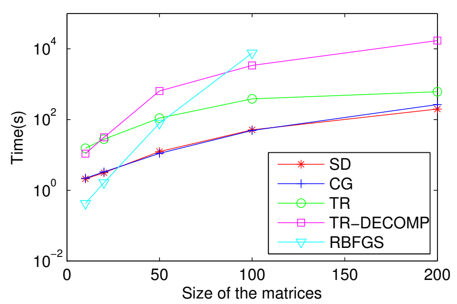

(b) Required time for the different methods using the inner product (4.3) for varying sizes.

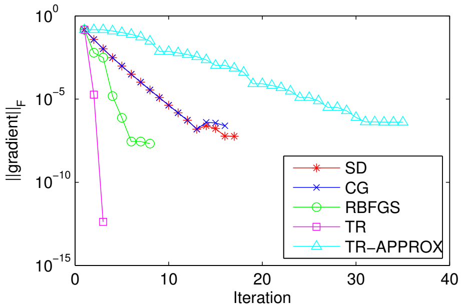

(c) Norm of the gradient throughout the iterations.

FIGURE 4.6. *Comparison of all the discussed algorithms. In the legends, the abbreviations SD (Steepest Descent), CG (Conjugate Gradient), TR (Trust Region), and RBFGS (Riemannian BFGS) are used to denote the applied technique and when necessary, the used inner product is indicated with SPD* (4.3) *or SYMM* (4.4)*. For TR the suffix DECOMP or APPROX is added to indicate the techniques in Section 4.3.2 and 4.3.3, respectively. In the first figure, the influence of the inner products is compared. In the second, we compare the algorithms which use the inner product* (4.3)*. In the last figure, the evolution of the gradient is depicted for all techniques using the inner product* (4.3)*.*

problem. The convergence of the second-order optimization methods, although predicted by the experiments, is theoretically not so easily guaranteed and will be treated as future work (indications are present in [1, 20]).

We noticed that while the second-order techniques required less iterations, the computational cost associated with each of these iterations was higher than that of the first-order algorithms, nullifying the advantage of a quadratic convergence. This effect was most visible when the size of the matrices increased. Hence, we conclude that for the current algorithms on the manifold $\mathbb{S}_+^n$, it is more advantageous to work with first-order optimization techniques when the size of the matrices increases (already at $n = 10$). It is possible to produce more efficient second-order optimization algorithms if we were able to reduce our search space, the manifold of interest, to a certain subset which structure can be further exploited. For example, the geometry of the manifold of larger matrices of fixed, low rank has already been extensively researched [13, 41] and can be used to apply the optimization techniques in this paper. This will also be a topic of future research.

The MATLAB code used to produce the experiments in this paper is available at `http://people.cs.kuleuven.be/~raf.vandebril/`.

Acknowledgments. The authors thank Samuel Relton for providing his MATLAB implementation of [4]. We also acknowledge Bruno Iannazzo and Daniel Kressner for the fruitful discussions and updates about new developments, and we thank Bruno Iannazzo for his hospitality during the research visit.

This research was partially supported 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).

REFERENCES

[1] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, Princeton, 2008.
[2] R. L. ADLER, J.-P. DEDIEU, J. Y. MARGULIES, M. MARTENS, AND M. SHUB, *Newton's method on Riemannian manifolds and a geometric model for the human spine*, IMA J. Numer. Anal., 22 (2002), pp. 359–390.
[3] A. H. AL-MOHY AND N. J. HIGHAM, *Improved inverse scaling and squaring algorithms for the matrix logarithm*, SIAM J. Sci. Comput., 34 (2012), pp. C153–C169.
[4] A. H. AL-MOHY, N. J. HIGHAM, AND S. D. RELTON, *Computing the Fréchet derivative of the matrix logarithm and estimating the condition number*, MIMS EPrint 2012.72, Manchester Institute for Mathematical Sciences, University of Manchester, Manchester, 2012.
[5] T. ANDO, C.-K. LI, AND R. MATHIAS, *Geometric means*, Linear Algebra Appl., 385 (2004), pp. 305–334.
[6] F. BARBARESCO, *Innovative Tools for Radar Signal Processing Based on Cartan's Geometry of SPD Matrices and Information Geometry*, in RADAR '08 IEEE Radar Conference 2008, Rome, IEEE Conference Proceedings, 2008, pp. 1–6.
[7] R. BHATIA, *Positive Definite Matrices*, Princeton University Press, Princeton, 2007.
[8] R. BHATIA AND R. L. KARANDIKAR, *Monotonicity of the matrix geometric mean*, Math. Ann., 353 (2012), pp. 1453–1467.
[9] D. BINI AND B. IANNAZZO, *A note on computing matrix geometric means*, Adv. Comput. Math., 35 (2011), pp. 175–192.
[10] ———, *Computing the Karcher mean of symmetric positive definite matrices*, Linear Algebra Appl., in press, 2012.
[11] D. BINI, B. MEINI, AND F. POLONI, *An effective matrix geometric mean satisfying the Ando–Li–Mathias properties*, Math. Comp., 79 (2010), pp. 437–452.
[12] R. L. BISHOP AND B. O'NEILL, *Manifolds of negative curvature*, Trans. Amer. Math. Soc., 145 (1969), pp. 1–49.
[13] S. BONNABEL AND R. SEPULCHRE, *Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank*, SIAM Matrix Anal. Appl., 31 (2009), pp. 1055–1070.
[14] W. M. BOOTHBY, *An Introduction to Differentiable Manifolds and Riemannian Geometry*, Academic Press, New York, 1975.
[15] M. BRIDSON AND A. HAEFLIGER, *Metric Spaces of Non-Positive Curvature*, Springer, Berlin, 1999.
[16] E. CARTAN, *Leçons sur la Géométrie des Espaces de Riemann*, Gauthier-Villars, Paris, 1946.
[17] R. FERREIRA, J. XAVIER, J. COSTEIRA, AND V. BAROSSO, *Newton algorithms for Riemannian distance related problems on connected locally symmetric manifolds*, Tech. Report, Institute for Systems and Robotics, IST Lisbon, Lisbon, November 2008.
[18] ———, *Newton method for Riemannian centroid computation in naturally reductive homogeneous spaces*, in Proc. IEEE International Conference on Acoustics, Speech and Signal 2006 Vol. 3, IEEE Conference Proceedings, 2006, pp.704–706.
[19] P. FLETCHER AND S. JOSHI, *Riemannian geometry for the statistical analysis of diffusion tensor data*, Signal Process., 87 (2007), pp. 250–262.
[20] K. A. GALLIVAN, C. QI, AND P.-A. ABSIL, *A Riemannian Dennis–Moré condition*, in High-Performance Scientific Computing – Algorithms and Applications, M. Berry, K. Gallivan, E. Gallopoulos, A. Grama, B. Philippe, Y. Saad, and F. Saied, eds., Springer, London, 2012, pp. 281–293.
[21] N. J. HIGHAM, *Functions of Matrices*, SIAM, Philadelphia, 2008.
[22] W. KENDALL, *Probability, convexity, and harmonic maps with small image I: uniqueness and fine existence*, Proc. London Math. Soc., 61 (1990), pp. 371–406.
[23] J. LAWSON AND Y. LIM, *Monotonic properties of the least squares mean*, Math. Ann., 351 (2011), pp. 267–279.
[24] J. M. LEE, *Manifolds and Differential Geometry*, AMS, Providence, 2009.

[25] H. LEE, Y. LIM, AND T. YAMAZAKI, *Multi-variable weighted geometric means of positive definite matrices*, Linear Algebra Appl., 435 (2011), pp. 307–322.
[26] Y. LIM, *Riemannian and Finsler structures of symmetric cones*, Trends in Mathematics, 4 (2001), pp. 111–118.
[27] ———, *Factorizations and geometric means of positive definite matrices*, Linear Algebra Appl., 437 (2012), pp. 2159–2172.
[28] Y. LIM AND M. PALFIA, *Matrix power means and the Karcher mean*, J. Funct. Anal., 262 (2012), pp. 1498–1514.
[29] M. MOAKHER, *A differential geometric approach to the geometric mean of symmetric positive-definite matrices*, SIAM J. Matrix Anal. Appl., 26 (2005), pp. 735–747.
[30] ———, *On the averaging of symmetric positive-definite tensors*, J. Elasticity, 82 (2006), pp. 273–296.
[31] K. NAKAMURA, *Geometric means of positive operators*, Kyunpook Math. J., 49 (2009), pp. 167–181.
[32] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer, New York, 1999.
[33] M. PALFIA, *The Riemannian barycenter computation and means of several matrices*, Int. J. Comput. Math. Sci., 3 (2009), pp. 128–133.
[34] ———, *A multivariable extension of two-variable matrix means*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 385–393.
[35] F. POLONI, *Constructing matrix geometric means*, Electron. J. Linear Algebra, 20 (2010), pp. 419–435.
[36] C. QI, *Numerical optimization methods on Riemannian manifolds*, PhD Thesis, Dept. of Mathematics, Florida State University, 2011.
[37] Y. RATHI, A. TANNENBAUM, AND O. MICHAILOVICH, *Segmenting images on the tensor manifold*, in IEEE Conference on Computer Vision and Pattern Recognition 2007, IEEE Conference Proceedings, pp. 1–8.
[38] Q. RENTMEESTER AND P.-A. ABSIL, *Algorithm comparison for Karcher mean computation of rotation matrices and diffusion tensors*, in Proc. 19th European Signal Processing Conference, EUSIPCO 2011, EURASIP, 2011, pp. 2229–2233.
[39] B. SAVAS AND L.-H. LIM, *Quasi-Newton methods on Grassmannians and multilinear approximations of tensors*, SIAM J. Sci. Comput., 32 (2010), pp. 3352–3393.
[40] B. VANDEREYCKEN, P.-A. ABSIL, AND S. VANDEWALLE, *Embedded geometry of the set of symmetric positive semidefinite matrices of fixed rank*, in Proc. of the IEEE 15th Workshop on Statistical Signal Processing 2009, IEEE Conference Proceedings, 2009, pp. 389–392.
[41] B. VANDEREYCKEN, P.-A. ABSIL, AND S. VANDEWALLE, *A Riemannian geometry with complete geodesics for the set of positive semidefinite matrices of fixed rank*, IMA J. Numer. Anal., in press, 2012.