The Abdus Salam International Centre for Theoretical Physics

2470-7

Advanced School and Workshop on Matrix Geometries and Applications

1 - 12 July 2013

Computational issues of matrix geometric means (4)

Dario Bini
Università di Pisa
Italy

Strada Costiera, 11 - 34151 - Trieste - Italy • Tel. +39 0402240111 • Fax. +39 040224163 • sci_info@ictp.it • www.ictp.it
ICTP is governed by UNESCO, IAEA, and Italy, and it is a UNESCO Category 1 Institute

Computational issues of matrix geometric means
Part 4

D.A. Bini
Università di Pisa

Trieste – July 2013

Contents

The Karcher mean

Definition: The Karcher mean $G = G(A_1, \ldots, A_k)$ is the matrix G where the following function takes its minimum.

$$f(X) = \sum_{i=1}^{k} d(X, A_i)^2 = \sum_{i=1}^{k} \| \log(A_i^{-\frac{1}{2}} X A_i^{-\frac{1}{2}}) \|^2$$

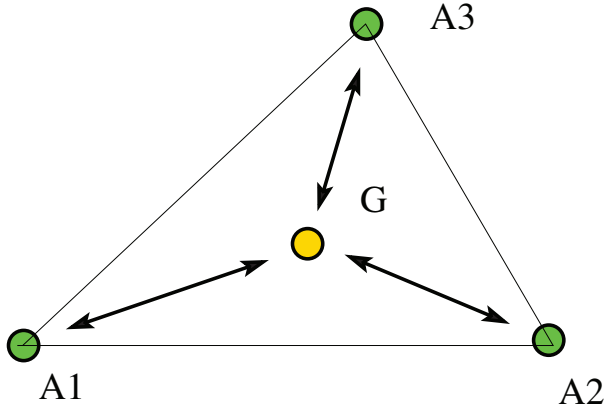

Property: The Karcher mean is unique and satisfies the 10 ALM axioms
It is also called least squares mean, or Riemannian mean [Moakher], [Bhatia], [Holbrook], [Jeuris, Vandebril, Vandereycken], [Lawson, Lim].

Existence and uniqueness follow from the fact that $f(X)$ is *geodesically convex*, i.e., it is convex along any geodesic

The Karcher mean

The point where the function takes its minimum is the unique positive definite matrix which solves the equation $\nabla f(X) = 0$, where $\nabla f(X)$ is the gradient of $f(x)$.

The gradient of $f(X)$ is

$$\nabla f(X) = 2X^{-1} \sum_{i=1}^{k} \log(XA_i^{-1}) = 2 \sum_{i=1}^{k} \log(A_i^{-1}X)X^{-1}$$

$$= 2X^{-\frac{1}{2}} \sum_{i=1}^{k} \log(X^{\frac{1}{2}} A_i^{-1} X^{\frac{1}{2}}) X^{-\frac{1}{2}}$$

This way, the Karcher mean can be viewed as the unique positive solution of the matrix equation $\nabla f(X) = 0$, that is

$$\sum_{i=1}^{k} \log(A_i^{-1}X) = 0$$

Solving matrix equations: Fixed point algorithms

A class of algorithms to solve a matrix equation of the kind $H(X) = 0$ is based on generating a matrix sequence as follows

$$X_{\nu+1} = \varphi(X_\nu), \qquad \nu = 0, 1, 2, \ldots$$
$$X_0 \quad \text{given}$$

where the matrix function $\varphi(X)$ has a fixed point G, i.e., such that

$$\varphi(G) = G$$

if G is such that $H(G) = 0$. Possible examples

$$\varphi_1(X) = X - \theta H(X), \qquad \varphi_2(X) = X \exp(-\theta H(X))$$

where $\theta \neq 0$, and for φ_2, X has real eigenvalues.

This way, if $\varphi(X)$ is continuous and the sequence X_ν converges to a limit G then $G = \varphi(G)$ and G is solution to the matrix equation

Their convergence analysis

Let $\varphi(X)$ be a matrix function and $L_\varphi(Y, E)$ its Fréchet derivative at Y. Then

$$\varphi(Y + E) - \varphi(Y) = L_\varphi(Y, E) + o(\|E\|) \doteq L_\varphi(Y, E)$$

Let G be such that $G = \varphi(G)$ and set $Y = G$, $E = X - G$. Then

$$\varphi(X) - G = \varphi(X) - \varphi(G) \doteq L_\varphi(G, X - G)$$

Therefore, denoting $E_\nu = X_\nu - G$ the error at step ν then

$$E_{\nu+1} \doteq L_\varphi(G, E_\nu)$$

Local convergence depends on the linear map $E \to L_\varphi(G, E)$.

There exists a neighborhood $\mathcal{U}$ of G such that for any $X_0 \in \mathcal{U}$ the sequence X_ν converges to G if the spectral radius of the linear map $E \to L_\varphi(G, E)$ is less than 1.

Let $e_\nu = \mathrm{vec}(E_\nu)$ and rewrite equation

$$E_{\nu+1} \doteq L_\varphi(G, E_\nu)$$

in matrix-vector form as

$$e_{\nu+1} \doteq K_\varphi(G) e_\nu$$

where $K_\varphi(G)$ is the $n^2 \times n^2$ matrix which represents the linear map $E \to L_\varphi(G, E)$ in terms of $\mathrm{vec}(E)$. Then

- $e_\nu \doteq K_\varphi(G)^\nu e_0$
- $\|K_\varphi(G)\| < 1$, for some operator norm $\|\cdot\|$, provides local convergence;
- the reduction of the norm error at any step is determined by $\|K_\varphi(G)\|$
- $\rho(K_\varphi(G))$ gives the asymptotic average of the error reduction per step in any vector norm
- if $\rho(K_\varphi(G)) = 0$ and the function is sufficiently regular, then convergence is at least quadratic

Algorithm 1

For $\theta > 0$, define

$$\varphi(X) = X - \theta X \sum_{i=1}^{k} \log(A_i^{-1} X) = X - \theta X^{\frac{1}{2}} \left(\sum_{i=1}^{k} \log(X^{\frac{1}{2}} A_i^{-1} X^{\frac{1}{2}}) \right) X^{\frac{1}{2}}$$

Remarks

- The correction coincides with $\frac{1}{2}\theta X \nabla_f(X) X$
- If θ is small enough, the sequence X_ν is formed by positive definite matrices
- $X = \varphi(X)$ if and only if $X \sum_{i=1}^{k} \log(A_i^{-1} X) = 0$
- Therefore, $\varphi(X)$ has only one positive definite fixed point X. It coincides with the solution of the equation $\sum_{i=1}^{k} \log(A_i^{-1} X) = 0$
- If $\lim_\nu X_\nu = \bar{X}$ then $\bar{X} \sum_{i=1}^{k} \log(A_i^{-1} \bar{X}) = 0$
- If $\bar{X}$ is positive definite then $\sum_{i=1}^{k} \log(A_i^{-1} \bar{X}) = 0$

Algorithm 1

In the case of scalars, choosing $\theta = \frac{1}{k}$, the function $\varphi(X)$ turns into

$$\varphi(x) = x(1 - \log \frac{x}{g}), \quad g = \left(\prod_{i=1}^{k} a_i \right)^{\frac{1}{k}}$$

The geometric mean g is the positive fixed point of $\varphi(x)$.

Moreover, since $\varphi'(x) = -\log \frac{x}{g}$, it holds that

$$\varphi'(g) = 0$$

Therefore if convergent, **the sequence x_ν converges quadratically to g**

Is this property true also in the matrix case?

Another algorithm: Algorithm 2

For $\theta > 0$ define

$$\varphi(X) = X \exp\left(-\theta \sum_{i=1}^{k} \log(A_i^{-1} X)\right)$$

$$= X^{\frac{1}{2}} \exp\left(-\theta \sum_{i=1}^{k} \log(X^{\frac{1}{2}} A_i^{-1} X^{\frac{1}{2}})\right) X^{\frac{1}{2}}$$

Remarks

- the sequence is formed by positive definite matrices
- $X = \varphi(X)$ if an only if $X \sum_{i=1}^{k} \log(A_i^{-1} X) = 0$
- Therefore, $\varphi(X)$ has only one positive definite fixed point X. It coincides with the solution of the equation $\sum_{i=1}^{k} \log(A_i^{-1} X) = 0$
- If $\lim_\nu X_\nu = \bar{X}$ then $\sum_{i=1}^{k} \log(A_i^{-1} \bar{X}) = 0$
- Approximating $\exp(Y)$ with $I + Y$ in $\varphi(X)$ provides the function φ of Algorithm 1.

Another algorithm: Algorithm 2

In the scalar case, with $\theta = 1/k$, the function $\varphi(x)$ turns into

$$\varphi(x) = xe^{-\log \frac{x}{g}} = x\frac{g}{x} = g$$

Therefore, in one step the sequence arrives at convergence

What happens in the matrix case?

Yet another algorithm

The matrix equation

$$\sum_{i=1}^{k} \log(A_i^{-1} X) = 0$$

can be viewed as a system of n^2 equations in n^2 unknowns

$$f(x) = 0, \qquad x \in \mathbb{R}^{n^2}, \quad f : \mathbb{R}^{n^2} \to \mathbb{R}^{n^2}$$

We can apply Newton's iteration to this system and get

$$x_{\nu+1} = x_\nu - J(x_\nu)^{-1} f(x_\nu)$$

If the Jacobian $J(x)$ of f is nonsingular at the solution g then the sequence x_ν has a local convergence which is at least quadratic.

The drawback of this approach is that at each step we have to solve a linear system of size n^2. If we do not exploit specific structures of the Jacobian, then the computational cost grows as $O(n^6)$.

Yet another algorithm

For the function $\sum_{i=1}^{k} \log(A_i^{-1}X)$, the Jacobian has a structure which is hardly exploitable.

Using the expression of the Fréchet derivative of the matrix logarithm [Higham, Function of Matrices 2008], one can prove that

$$J = \sum_{i=1}^{k} (I \otimes (X^{-1}A_i)) h(\log(A_i^{-1}X)^T \otimes I - I \otimes \log A_i^{-1}X)(I \otimes A_i^{-1}),$$

$$h(t) = \frac{t}{e^t - 1}$$

Question: Is it possible to solve an $n^2 \times n^2$ system with the coefficient matrix J in $O(n^3)$ operations? It is very unlikely

Analysis of Algorithm 1

Recall that $E_{\nu+1} \doteq L_\varphi(G, E_\nu)$, $E_\nu = X_\nu - G$

Problem: evaluate the Fréchet derivative $L_\varphi(G, E)$ of $\varphi(X)$ at $X = G$ in the direction of E.

Properties of the Fréchet derivative:

$$L_{g*h}(X, E) = L_g(X, E)h(X) + g(X)L_h(X, E) \qquad \text{product rule}$$

$$L_{g\circ h}(X, E) = L_g(h(X), L_h(X, E)) \qquad \text{chain rule}$$

$$L_{f^{-1}}(f(X), E) = L_f^{-1}(X, E) \qquad \text{derivative of inverse functior}$$

$$\varphi(X) = X - \theta X \sum_{i=1}^{k} \log(A_i^{-1}X) =: X - \theta \sum_{i=1}^{k} \psi_i(X),$$

$$\psi_i(X) = X \log(A_i^{-1}X)$$

For simplicity consider $\psi(X) = X \log(A^{-1}X)$ and compute $L_\psi(X, E)$.

Properties:

$$L_{g*h}(X,E) = L_g(X,E)h(X) + g(X)L_h(X,E) \quad \text{product rule}$$
$$L_{g\circ h}(X,E) = L_g(h(X), L_h(X,E)) \quad \text{chain rule}$$

For simplicity consider $\psi(X) = X\log(A^{-1}X)$ and compute $L_\psi(x,E)$.

Applying the product rule with $g(X) = X$ and $h(X) = \log(A^{-1}X)$ and then the chain rule with $h(X) = A^{-1}X$, $g(X) = \log(X)$ yields

$$\begin{aligned} L_\psi(X,E) &= L_X(X,E)\log(A^{-1}X) + X\,L_{\log(A^{-1}X)}(X,E) && \text{product rule} \\ &= E\,\log(A^{-1}X) + X\,L_{\log(A^{-1}X)}(X,E) \\ &= E\,\log(A^{-1}X) + X\,L_{\log}(A^{-1}X, A^{-1}E) && \text{chain rule} \end{aligned}$$

$$L_{\psi_i}(X,E) = E\ \log(A_i^{-1}X) + X\, L_{\log}(A_i^{-1}X, A_i^{-1}E)$$

Recall that

$$\varphi(X) = X - \theta \sum_{i=1}^{k} \psi_i(X)$$

so that

$$L_{\varphi}(G,E) = E - \theta E \sum_{i=1}^{k} \log(A_i^{-1}G) - \theta G \sum_{i=1}^{k} L_{\log}(A_i^{-1}G, A_i^{-1}E)$$

But G solves the matrix equation so that the blue part is zero. Thus,

$$L_{\varphi}(G,E) = E - \theta G \sum_{i=1}^{k} L_{\log}(A_i^{-1}X, A_i^{-1}E)$$

In matrix-vector form $\mathrm{vec}(L_{\varphi}(G,E)) = K_{\varphi}(G)\mathrm{vec}(E)$

$$K_{\varphi}(G) = I - \theta(I \otimes G) \sum_{i=1}^{k} K_{\log}(A_i^{-1}G)(I \otimes A_i^{-1})$$

It remains to evaluate $K_{\log}(A_i^{-1}G)$. From the rule of inverse function derivative one has $K_{\log}(X) = K_{\exp}^{-1}(Y),\ Y = \log X$

From [Higham, Function of Matrices 2008] one has

$$K_{\exp}(Y) = \gamma(Y^T \otimes I - I \otimes Y)(I \otimes \exp(Y)), \quad \gamma(t) = (e^t - 1)/t$$

so that

$$K_{\log}(X) = (I \otimes X^{-1})\beta(\log X^T \otimes I - I \otimes \log X), \quad \beta(t) = \frac{1}{\gamma(t)} = t/(e^t - 1)$$

replacing this expression in

$$K_\varphi(G) = I - \theta(I \otimes G)\sum_{i=1}^{k} K_{\log}(A_i^{-1}G)(I \otimes A_i^{-1})$$

yields

$$K_\varphi(G) = I - \theta\sum_{i=1}^{k} \beta(\log(GA_i^{-1}) \otimes I - I \otimes \log(GA_i^{-1})), \quad \beta(t) = t/(e^t - 1)$$

$$K_\varphi(G) = I - \theta \sum_{i=1}^{k} \beta(\log(GA_i^{-1}) \otimes I - I \otimes \log(GA_i^{-1})), \quad \beta(t) = t/(e^t - 1)$$

$$e_{\nu+1} = K_\varphi(G) e_\nu \quad e_\nu = \mathrm{vec}(E_\nu), \quad E_\nu = X_\nu - G$$

Introduce $\widehat{E}_\nu = G^{-\frac{1}{2}} E_\nu G^{-\frac{1}{2}} = G^{-\frac{1}{2}} X_\nu G^{-\frac{1}{2}} - I$ so that for $\widehat{e}_\nu = \mathrm{vec}(\widehat{E}_\nu)$ it holds that

$$\widehat{e}_{\nu+1} = (G^{-\frac{1}{2}} \otimes G^{-\frac{1}{2}}) K_\varphi(G) (G^{-\frac{1}{2}} \otimes G^{-\frac{1}{2}}) \widehat{e}_\nu =: \widehat{K}_\varphi(G) \widehat{e}_\nu$$

$$\widehat{K}_\varphi(G) = I - \theta \sum_{i=1}^{k} \beta(M_i \otimes I - I \otimes M_i), \quad M_i = G^{-\frac{1}{2}} A_i^{-1} G^{-\frac{1}{2}}$$

wich shows that $\widehat{K}_\varphi$ is **symmetric**. Moreover, M_i are positive definite and since the function $\beta(t)$ takes positive values also the matrices in the summation are positive definite as well their sum.

Theorem

Let G be the Karcher mean of $A_1, \ldots, A_k$, let X_ν be the sequence defined by $X_{\nu+1} = \varphi(X_\nu)$ with

$$\varphi(X) = X - \theta X \sum_{i=1}^{k} \log(A_i^{-1} X)$$

Define $\widehat{E}_\nu = G^{-1/2}(X_\nu - G)G^{-1/2}$, $\widehat{e}_\nu = vec(\widehat{E}_\nu)$. Then

$$\widehat{e}_{\nu+1} \doteq (I - \theta H)\widehat{e}_\nu \quad H = \textstyle\sum_{i=1}^{k} H_i$$

$$H_i = \beta(\log(M_i) \otimes I - I \otimes \log(M_i)), \quad \beta(t) = t/(e^t - 1)$$

$$M_i = G^{1/2} A_i^{-1} G^{1/2}$$

Remark

M_i are symmetric positive definite as well as H_i and H. Therefore, for $\theta > 0$ small enough, $I - \theta H$ has spectral radius less than 1 and the method is locally convergent

Question:

What is the best choice of θ which minimizes the spectral radius of $I - \theta H$?

- The eigenvalues of $I - \theta H$ are $1 - \theta\lambda_j(H)$, $j = 1, \ldots, n^2$. If $m \leq M$ are the minimum and maximum eigenvalues of H, respectively, then for $\theta = \frac{2}{M+m}$ we get the minimal spectral radius that is $\frac{M-m}{M+m}$

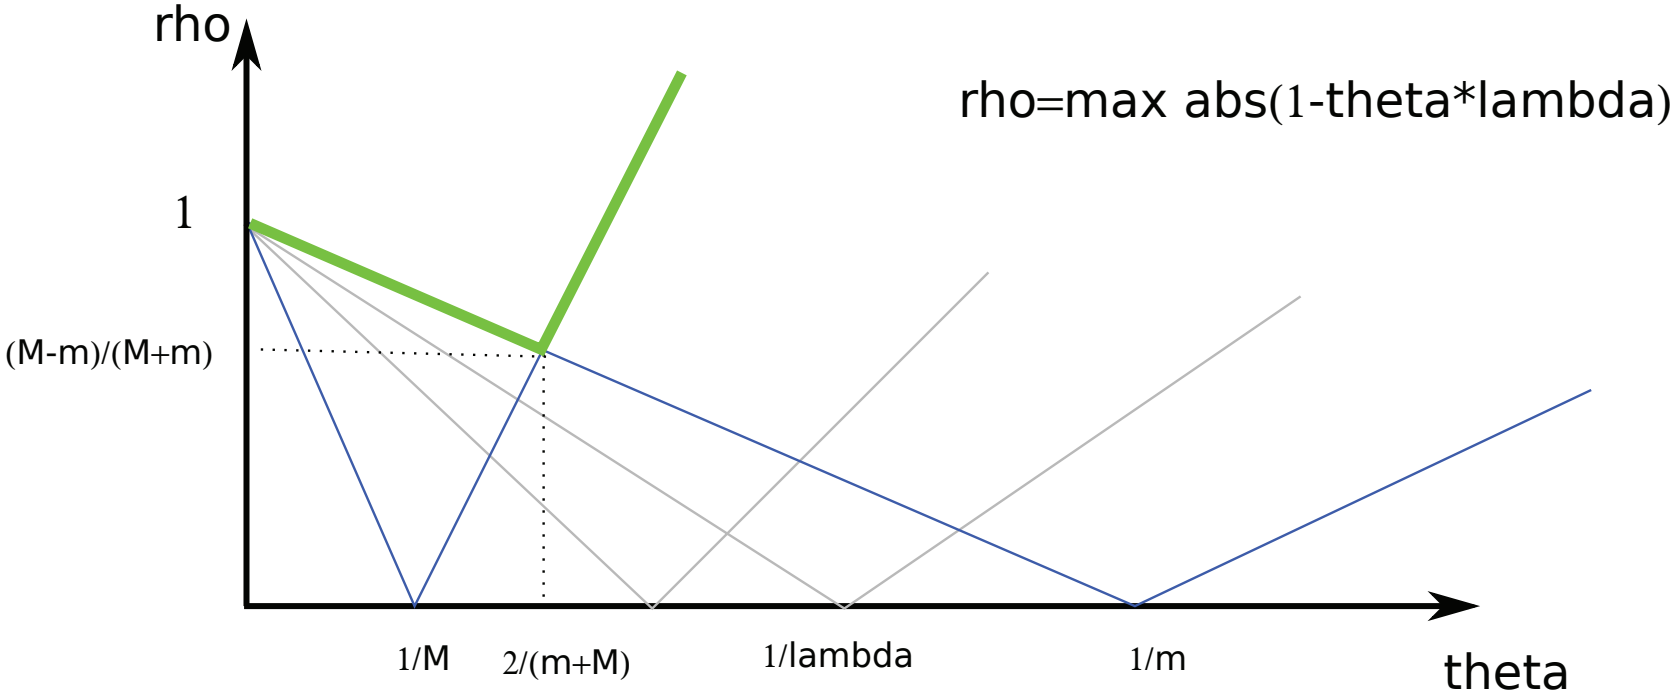

- From the Courant-Fischer theorem it holds that for any pair A, B of real symmetric matrices

$$\lambda_{\min}(A)+\lambda_{\min}(B) \leq \lambda_{\min}(A+B) \leq \lambda_{\max}(A+B) \leq \lambda_{\max}(A)+\lambda_{\max}(B)$$

Therefore

$$\sum_{i=1}^{k} \lambda_{\min}(H_i) \leq m \leq M \leq \sum_{i=1}^{k} \lambda_{\max}(H_i)$$

- Since the function $\beta(t)$ is decreasing, then

$$\lambda_{\min}(H_i) = \frac{\log c_i}{c_i - 1}, \quad \lambda_{\max}(H_i) = c_i \frac{\log c_i}{c_i - 1}, \quad c_i = \frac{\lambda_{\max}(M_i)}{\lambda_{\min}(M_i)}$$

Combining the above properties one gets the optimal value

$$\rho(I - \theta H) \leq \frac{\sum_{i=1}^{k} \log c_i}{\sum_{i=1}^{k} \frac{c_i+1}{c_i-1} \log c_i}, \quad \text{for} \quad \theta = \frac{2}{\sum_{i=1}^{k} \frac{c_i+1}{c_i-1} \log c_i}$$

where

$$c_i = \frac{\lambda_{\max}(M_i)}{\lambda_{\min}(M_i)}, \quad M_i = G^{\frac{1}{2}} A_i^{-1} G^{\frac{1}{2}}$$

Remarks:

- c_i close to 1 for any i, i.e., M_i well conditioned, implies ρ close to 0, that is fast convergence
- Recall that $d(A_i, G) = \| \log M_i \|$. If for any i, A_i is close to G in the Riemannian metric then M_i is close to I, c_i is close to 1 that is, convergence is fast.
- If A_i commute, then $H_i = I$ so that for $\theta = 1/k$ one has $I - \theta H = 0$. That is convergence is at least quadratic
- If A_i almost commute but are "far" from each other, then $\theta = 1/k$ leads to failure

Some numerical experiments

Number of iterations for k random matrices with condition number 10^2 and 10^4, respectively, where we set X_0 equal to the Cheap mean.

k	3	4	5	6	7	8	9	10
cond=10^2	17	17	16	16	15	15	14	14
cond=10^4	41	37	35	31	29	29	29	28

Number of iterations for 10 random matrices with condition number 20 and 10^5, respectively, lying in a neighborhood of radius ϵ.

ϵ	0.5	10^{-1}	10^{-2}	10^{-3}	10^{-4}
cond= 20	5	4	2	1	1
cond= 10^5	22	19	14	12	6

Implementation

The choice of the optimal value of θ can be performed by replacing G with the current approximation X_ν in the expression

$$c_i = \mathrm{cond}(M_i), \quad M_i = G^{\frac{1}{2}} A_i^{-1} G^{\frac{1}{2}}$$

Since $\mathrm{cond}(M_i) = \mathrm{cond}(M_i^{-1})$ we use the expression
$c_i = \mathrm{cond}(G^{-\frac{1}{2}} A_i G^{-\frac{1}{2}})$

The computation of the correction (scaled gradient) and of the optimal value of θ can be performed simultaneously by the following function

```
function [w,theta]=correction(varargin)
% [w,theta]=correction(A1,...,Ak,X)
% compute the correction
% -S*sum(log(SI*A_i*SI))S, S=X^{1/2}, SI=inv(S)
% at X and provides the optimal theta
% theta=2/(sum log(c_i)(c_i+1)/(c_i-1)), c_i=cond(SI*A_i*SI)
```

Implementation: Computing the correction and the optimal θ

```
A=varargin; p=length(A); X=A{p};  k=p-1;
n=length(A{1});
sqx=sqrtm(X);  isqx=inv(sqx); v=zeros(n);  theta=0;
for i=1:k
   M=isqx*A{i}*isqx;
   c=cond(M);
   if c != 1
     theta=theta+(c+1)*log(c)/(c-1);
   else
     theta=theta+2;
   end;
   v=v+logm(M);
end
w=-sqx*v*sqx;
theta=2/theta;
```

Implementation: the Karcher mean

```
function [G,iter]=karcher(varargin)
% [G,iter]=karcher(A_1,...,A_k) computes the Karcher mean by the
% Richardson-like iteration X_new= X+theta*S*sum(log(SI*A_i*SI))S
% S=X^{1/2}, SI=inv(S), with optimal theta
%    A1,...,Ak: positive definite matrices
%    G: the Karcher mean of A1,...,Ak
%    iter: the number of iterations needed

A=varargin; k=length(A);  n=length(A{1});
tol=1d-11; maxiter=1000;  X=cheap(A{});
for iter=1:maxiter
  A{k+1}=X;  [C,theta]=correction(A{});  err=norm(C);  X=X-theta*C;
  if err<tol  break  end
end
G=X;
if iter==maxiter
  disp('Reached the maximum number of iterations');
end
```

A cheaper formula

In order to avoid inversion of A_i we used the formula

$$\log(X^{\frac{1}{2}} A_i^{-1} X^{\frac{1}{2}}) = -\log(X^{-\frac{1}{2}} A_i X^{-\frac{1}{2}})$$

A less expensive formula can be derived by replacing X with the Cholesky factorization $X = LL^T$, where L is lower triangular:

$$\begin{aligned}
\varphi(X) &= X - \theta X \sum_{i=1}^{k} \log(A_i^{-1} X) \\
&= X - \theta L \sum_{i=1}^{k} \log(L^T A_i^{-1} L) L^T \\
&= X + \theta L \sum_{i=1}^{k} \log(L^{-1} A_i L^{-T}) L^T
\end{aligned}$$

Computing L and L^{-1} is much cheaper than computing $X^{1/2}$ and its inverse. Moreover, $L^{-1}A_iL^{-T}$ has the same eigenvalues of $X^{-\frac{1}{2}}A_iX^{-\frac{1}{2}}$ so that

$$c_i = \mathrm{cond}(M_i) = \mathrm{cond}(L^{-1}A_iL^{-T})$$

Proof:

$$\begin{aligned}\det(L^{-1}A_iL^{-T} - \lambda_i) &= \det(A_i - \lambda LL^T) = \det(A_i - \lambda X) \\ &= \det(A_i - \lambda X^{\frac{1}{2}}X^{\frac{1}{2}}) = \det(X^{-\frac{1}{2}}A_iX^{-\frac{1}{2}} - \lambda I)\det X\end{aligned}$$

Therefore in the `correction` function it is enough to replace

`sqx=sqrtm(X);` with `sqx=chol(X)';`

and to add some trasposition, i.e., to replace

`M=isqx*Ai*isqx;` with `M=isqx*Ai*isqx';`

and `w=-sqx*v*sqx;` with `w=-sqx*v*sqx';`

Analysis of Algorithm 2

$$\varphi(X) = X \exp\left(-\theta \sum_{i=1}^{k} \log(A_i^{-1} X)\right)$$
$$= X^{\frac{1}{2}} \exp\left(-\theta \sum_{i=1}^{k} \log(X^{\frac{1}{2}} A_i^{-1} X^{\frac{1}{2}})\right) X^{\frac{1}{2}}$$

Since $\exp(X) = I + X + O(X^2)$, the function $\varphi(X)$ can be written as

$$\varphi(X) = X - \theta X \sum_{i=1}^{k} \log(A_i^{-1} X) + O(X^2)$$

so that $\varphi(X)$ has the same Fréchet derivative as the function $\varphi(X)$ of Algorithm 1.
Therefore, the same analysis applies.

Other algorithms

Inductive Mean (Spiral Descent) [Holbrook]

$$S_{\nu+1} = S_\nu \#_{\frac{1}{\nu}} A_{1+(\nu \mod n)}, \quad S_1 \in \mathcal{P}_n$$

$$G = \lim_\nu S_\nu$$

The t-power mean [Lim, Palfia] is the unique solution $G_t \in \mathcal{P}_n$ to the the matrix equation

$$X = \frac{1}{k} \sum_{i=1}^{k} X \#_t A_i$$

It holds that

$$G = \lim_{t \to 0} X_t$$

The t-power mean can be computed by using the iteration

$$X_{\nu+1} = \frac{1}{k} \sum_{i=1}^{k} X_\nu \#_t A_i$$

Convergence is linear with reduction of the error per step $1 - t$.

```
function [G,iter]=inductive_mean(varargin)
A=varargin; k=length(A); n=size(A{1}); tol=1.e-7; maxiter=100
S=eye(n);
for iter=1:maxiter
   h=1+mod(iter,k);
   S1=sharp(S,A{h},1/iter);
   if norm(S-S1)<tol  break  end
   S=S1;
end
G=S1;
if iter==maxiter disp("reached the max number of iterations")
end
```

Structured mean: definition and algorithms

Fact: Unfortunately, the Karcher mean does not preserve structure.
$A_1, \ldots, A_k \in \mathcal{A} \subset \mathcal{P}_n \not\Rightarrow G(A_1, \ldots, A_k) \in \mathcal{A}$

Example: $A = \begin{bmatrix} 2 & 1 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 0 & 1 & 2 & 1 \\ 0 & 0 & 1 & 2 \end{bmatrix}$, $B = I$, $G = \begin{bmatrix} 1.3590 & 0.3860 & -0.0611 & 0.0173 \\ 0.3860 & 1.2978 & 0.4034 & -0.0611 \\ -0.0611 & 0.4034 & 1.2978 & 0.3860 \\ 0.0173 & -0.0611 & 0.3860 & 1.3590 \end{bmatrix}$

A different definition is needed [B., Iannazzo, Jeuris, Vandebril 2013]:

Let $\sigma(t) : \mathbb{R}^q \to \mathbb{R}^{n \times n}$ be a differentiable map and define $\mathcal{A} = \sigma(\mathbb{R}^q) \cap \mathcal{P}_n$, where $\mathcal{P}_n$ is the set of positive definite matrices.

Definition (Structured mean)

The **structured mean with respect to** $\mathcal{A}$ of the matrices

$$A_i = \sigma(a_i) \in \mathcal{A}, \quad a_i \in \mathbb{R}^q, \quad i = 1, \ldots, k$$

is the set

$$G_{\mathcal{A}} = \{X = \sigma(t) \in \mathcal{A} : \quad f(X) = \inf_{Y \in \mathcal{A}} f(Y)\}$$

Properties that we can prove

- $G_{\mathcal{A}}$ is not empty if $\mathcal{A}$ is a linear space
- If $\mathcal{A}$ is geodesically convex, then the structured mean coincides with the Karcher mean.
- If $\sigma(\mathbb{R}^q)$ is a matrix algebra then it is geodesically convex
- the structured mean is not generally unique

Example of non uniqueness

Let $A = I$, $B = \text{diag}(\alpha, \alpha^{-1})$, $\sigma(t) = A + t(B - A)$, $0 \leq t \leq 1$.

the function $f(t) = \delta^2(\sigma(t), A) + \delta^2(\sigma(t), B)$

is symmetric w.r.t. $t = 1/2$ and for $\alpha = 100$ is such that

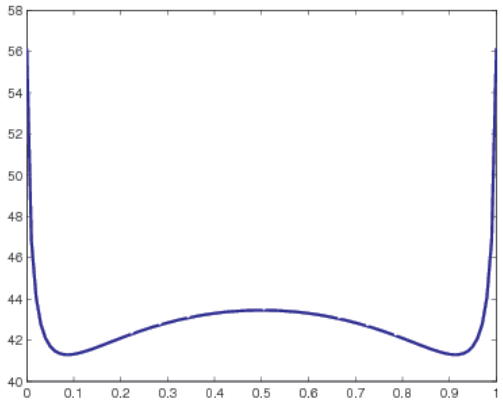

Properties that we can prove

Self duality $G(A_1, \dots, A_k)^{-1} = G(A_1^{-1}, \dots, A_k^{-1})$ is satisfied in the following form

$$G_{\mathcal{A}}(A_1, \dots, A_k)^{-1} = G_{\mathcal{A}^{-1}}(A_1^{-1}, \dots, A_k^{-1}), \quad \mathcal{A}^{-1} = \{A^{-1} : \quad A \in \mathcal{A}\}$$

The inverse of the structured mean w.r.t. $\mathcal{A}$ coincides with the structured mean of the inverses w.r.t. $\mathcal{A}^{-1}$

Congruence invariance $G(S^T A_1 S, \dots, S^T A_k S) = S^T G(A_1, \dots, A_k) S$ is satisfied in the following form

$$G_{S^T \mathcal{A} S}(S^T A_1 S, \dots, S^T A_k S) = G_{\mathcal{A}}(A_1, \dots, A_k)$$

Properties that we can prove

Joint homogeneity $G_{\mathcal{A}}(\alpha_1 A_1, \ldots, \alpha_k A_k) = (\prod_i \alpha_i)^{1/k} G(A_1, \ldots, A_k)$ holds if $\mathcal{A}$ is a linear space

Permutation invariance

Repetition invariance: $G(A_1, \ldots, A_k, A_1, \ldots, A_k) = G(A_1, \ldots, A_k)$

Monotonicity is false: there are counterexamples

Computing the structured mean: A vector equation

The set of structured means $G = \sigma(g)$ is a subset of the set of stationary points for $f(\sigma(t))$, i.e., $\nabla_t f(\sigma(t)) = 0$

The vectors g are the solutions of the vector equation $\nabla_t f(t; a_1, \ldots, a_k) = 0$ such that $\sigma(g)$ is positive definite

From the chain rule of derivatives, this leads to a vector equation which, for $\mathcal{A}$ linear space, takes the form

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

where $\sigma(t) : \mathbb{R}^q \to \mathbb{R}^{n \times n}$ is linear and $\mathrm{vec}(\sigma(t)) = Ut$ for $U \in \mathbb{R}^{q \times n^2}$

An example: Toeplitz matrices

A matrix $A = (a_{i,j})$ is Toeplitz if $a_{i,j} = t_{i-j}$

$$\sigma((t_0, t_1, t_2)) = \begin{bmatrix} t_0 & t_1 & t_2 \\ t_1 & t_0 & t_1 \\ t_2 & t_1 & t_0 \end{bmatrix}$$

$$\mathrm{vec}(\sigma(t)) = (t_0, t_1, t_2, t_1, t_0, t_1, t_2, t_1, t_0)^T$$

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

Algorithms for solving the matrix/vector equation

In the **structured case** we consider iterations of the kind

$$t^{(\nu+1)} = \varphi(t^{(\nu)}), \quad \varphi(t) = t - \theta V(t)^{-1} \nabla_t f(t)$$

where $V(t)$ is a suitable invertible matrix.

Remarks:

- If $V(t)$ is the Jacobian matrix of $\nabla_t f(t)$ then the algorithm is Newton's iteration
- The Jacobian of $\nabla_t f(t)$ can be obtained as $U^T K_\Gamma U$ where K_Γ is the matrix which represents the Fréchet derivative of $\Gamma(X)$. Using the same technique as before yields

$$K_\Gamma(X) = -X^{-1} \sum_{i=1}^{k} \log(XA_i^{-1}) \otimes X^{-1}$$

$$+ \sum_{i=1}^{k} (A_i^{-1} \otimes X^{-1}) \beta(W_i)(I \otimes AX^{-1})$$

$$W_i = \log(XA^{-1}) \otimes I - I \otimes \log(XA^{-1}), \qquad \beta(z) = z/(e^z - 1)$$

Newton's iteration is very expensive: the Jacobian depends on all the matrices $A_1, \ldots, A_n$

Cheaper iterations can be obtained by looking for a matrix V which depends only on the current approximation. Here are some possibilities:

1. $V(t) = D, \quad D = U^T U$
2. $V(t) = \left[U^T(\sigma(t)^{-1} \otimes \sigma(t)^{-1})U\right]$
3. $V^{-1}(t) = D^{-1}U^T(\sigma(t) \otimes \sigma(t))UD^{-1}$

Motivation:

1. projected gradient descent method w.r.t. the Euclidean inner product $\langle A, B\rangle = \mathrm{trace}(AB)$
2. projected gradient descent method w.r.t. the "natural" scalar product $\langle A, B\rangle_X = \mathrm{trace}(AX^{-1}BX^{-1})$
3. "projected version" of the transformation performed in the unstructured case: $\nabla_X f \quad \to \quad X(\nabla_X f)X$

Convergence analysis

Convergence analysis can be performed by computing the Jacobian of $\varphi(\sigma(t))$ through the Fréchet derivative of $\Gamma(X)$

By performing a convergence analysis similar to that of the unstructured case one finds that the Jacobian of $\varphi(t)$ in G is

$$J = I - \theta V^{-1} U (I \otimes G) H (I \otimes G^{-1}) U$$

where H is the same as in the unstructured case, that is

$$H = \sum_{i=1}^{k} H_i \qquad H_i = \beta(\log(M_i) \otimes I - I \otimes \log(M_i))$$

$$\beta(t) = t/(e^t - 1) \qquad M_i = G^{1/2} A_i^{-1} G^{1/2}$$

In the case of the second algorithm we find exactly the same bounds of the unstructured case:

$$\rho(J) \leq \frac{\sum_{i=1}^{k} \log c_i}{\sum_{i=1}^{k} \frac{c_i+1}{c_i-1} \log c_i}, \quad \text{for} \quad \theta = \frac{2}{\sum_{i=1}^{k} \frac{c_i+1}{c_i-1} \log c_i}$$

where

$$c_i = \frac{\lambda_{\max}(M_i)}{\lambda_{\min}(M_i)}, \quad M_i = G^{\frac{1}{2}} A_i^{-1} G^{\frac{1}{2}}$$

Some experiments

Case 1: $n = 10,\ k = 5,\ \epsilon = 10^{-3}$

$$A_i = H + \epsilon \texttt{Toep}(\texttt{rand}(\texttt{n}, 1)), \quad i = 1, \ldots, k, \qquad H = \texttt{Toep}([5\ 1 \ldots 1])$$

$\text{cond}(H) = 2.5$

Case 2: $n = 10,\ k = 5,\ \epsilon = 10^{-3}$

$$A_i = H + \epsilon \texttt{Toep}(\texttt{rand}(\texttt{n}, 1)), \quad i = 1, \ldots, k, \qquad H = \texttt{Toep}([\texttt{n} : -1 : 1])$$

$\text{cond}(H) = 132.36$

Case 3: $n = 10,\ k = 5,\ \epsilon = 10^{-3}$ $A_i = \texttt{Toep}(t)$,

```
t= rand(n,1); t(1)=t(1)-min(eig(Toep(t)))+1.e-3;
```

Number of iterations

Case	Iter. 1	Iter 2.	Iter. 3
1	33	3	11
2	> 1000	3	47
3	> 1000	32	183

Implementation: case of Toeplitz matrices

We need some functions:

```
function X=toep(t)
% X:  symmetric toeplitz matrix whose first column is t
n=max(size(t));   X=zeros(n);
for i=1:n
  for j=1:n
    X(i,j)=t(abs(i-j)+1);
  end
end
```

The matrix U

```
function U=Utoep(n)
% U = Utoep computes the n^2 x n matrix U such that Ut=vec(toep(t))
ey=eye(n);
U=vec(toep(ey(1,:)));
for i=2:n
  U=[U,vec(toep(ey(i,:)))];
end;
```

Implementation: case of Toeplitz matrices

Computing the projected gradient

```
function [w,theta]=projgrad(varargin)
% w=projgrad(t1,...,tk,x,U)  computes the projected gradient
% w = [U^T(X \otimes X)^(-1)U]^(-1) U^T vec(grad(X))
% at x of the Toeplitz matrices toep(t1),...,toep(tk)
% U is such that vec(Ut)=toep(t)

t=varagin; n=length(t{1}); p=length(t); U=t{p}; k=p-2;
for i=1:k; ai{i}=toep(ti{i});end
x=toep(ti{k+1});  ai{k+1}=x;
[grd,theta]=grad(ai{});
utgrad=U'*vec(grd);
xi=inv(x); vv=U'*kron(xi,xi)*U;
w=(z\utgrad)';
```

Implementation: Computing the gradient

```
function [w,theta]=grad(varargin)
% [w,theta]=grad(A1,...,Ak,X) computes the gradient
% w=-SI*sum(log(SI*A_i*SI))*SI, S=X^{1/2}, SI=inv(S) at X and
% provides the optimal theta theta=2/(sum log(c_i)(c_i+1)/(c_i-1)),
% c_i=cond(SI*A_i*SI)

A=varargin; p=length(A); X=A{p};  k=p-1;
n=length(A{1}); sqx=sqrtm(X); isqx=inv(sqx); v=zeros(n); theta=0;
for i=1:k
   M=isqx*A{i}*isqx;   v=v+logm(M);  c=cond(M);
   if c != 1
     theta=theta+(c+1)*log(c)/(c-1);
   else
     theta=theta+2;
   end;
end
w=-isqx*v*isqx;  theta=2/theta;
```

Implementation: case of Toeplitz matrices

Computing the structured mean

```
function [t, it]=toeplitzmean(varargin)
  % t=toeplitzmean(t1,t2,...,tk,t0)
  % G=toep(t) is the structured geometric mean of the Toeplitz
  % matrices A1=toep(t1), ..., Ak=toep(tk) starting from toep(t0)

A=varargin; n=length(A{1});  p=length(A); itmax=100;
t=A{p}; k=p-1;
U=Utoep(n); A{k+2}=U;
for it=1:itmax
   A{k+1}=t;  [pgrd,theta]=projgrad(A{});
   t=t-theta*pgrd;   err=max(abs(pgrd));
   if err<1.0e-12
      break
   end
end
if it==itmax
  disp("Warning: reached the max number of iterations")
end
```

Conclusions

- Matrix geometric means (MGM) are needed in the applications
- For two positive definite matrices there is a unique definition of MGM
- For $k > 2$ matrices there are many MGMs
- The mean based on medians can be computed faster than the ALM mean, however, the cost grows exponentially with k
- The cheap mean has a polynomial cost but does not satisfy the monotonicity property
- The Karcher mean has the good properties: it requires the solution of a matrix equation.
- Effective algorithms exist for solving this matrix equation
- A structured mean has been introduced with the property of preserving structures
- An effective algorithm for its computation has been provided

Some open issues and things to do

- Global convergence of the Cheap mean
- Monotonicity of the Cheap mean for close input matrices
- Global convergence of the Richardson iteration for the Karcher mean
- Analysis of the distances of the different geometric means
- Conditions for the uniqueness of the structured mean
- Structured means through positive parametrizations

Bibliography

- T. Ando, C. Li, and R. Mathias. Geometric means. *Linear Algebra and its Applications*, 385:305–334, 2004.
- R. Bhatia. The Riemannian mean of positive definite matrices. Matrix Information Geometry, (Eds.) F. Nielsen, R. Bhatia, Springer 2013.
- R. Bhatia. *Positive Definite Matrices*. Princeton Series in Applied Mathematics. Princeton University Press, 2007.
- R. Bhatia and J. Holbrook. Riemannian geometry and matrix geometric means. *Linear Algebra and its Applications*, 413(2-3):594–618, 2006.
- 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.
- D. A. Bini and B. Iannazzo. A note on computing matrix geometric means. *Advances in Computational Mathematics*, 35(2-4):175–192, 2011.
- D. A. Bini, B. Iannazzo. Computing the Karcher mean of symmetric positive definite matrices. Linear Algebra and Appl., 438(4):1700–1710, 2013.
- D. A. Bini, B. Iannazzo, B. Jeuris, R. Vandebril, Geometric means of structured matrices, BIT, in press.

Bibliography

- N. J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2008.
- J. Holbrook, No dice: a deterministic approach to the Cartan centroid, J. Ramanujan Math. Soc., Vol. 27 No. 4 – Dec 2012.
- B. Iannazzo. The geometric mean of two matrices from a computational viewpoint. Preprint, available on arXiv, 2011.
- B. Iannazzo, B. Meini. Palindromic matrix polyomials, matrix functions and integral representations, Linear Algebra Appl., 434-1 (2011), pp. 174-184.
- 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.
- 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.
- J. Lawson and Y. Lim. Monotonic properties of the least squares mean. *Mathematische Annalen*, 351(2):267–279, 2011.

Bibliography

- Y. Lim, Y., M. Pálfia, The matrix power means and the Karcher mean. J. Funct. Anal. 262, 1498–1514 (2012)
- 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.
- M. Moakher. On the averaging of symmetric positive-definite tensors. *Journal of Elasticity*, 82(3):273–296, 2006.
- K. Nakamura. Geometric means of positive operators. *Kyunpook Mathematical Journal*, 49:167–181, 2009.
- M. Pálfia. The Riemann barycenter computation and means of several matrices. *Int. J. Comput. Math. Sci*, 3(3):128–133, 2009.