# Riemannian Optimization with its Application to Averaging Positive Definite Matrices

Wen Huang

Xiamen University

October 23, 2019

# Riemannian Optimization

**Problem:** Given $f(x) : \mathcal{M} \to \mathbb{R}$, solve

$$\min_{x \in \mathcal{M}} f(x)$$

where $\mathcal{M}$ is a Riemannian manifold.

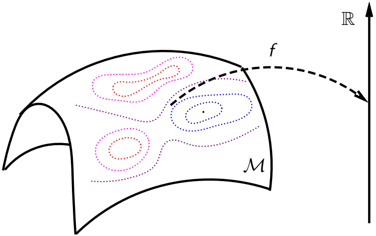

# Examples of Manifolds

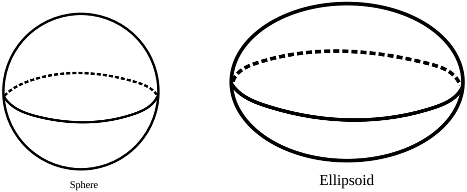


- Stiefel manifold: $\mathrm{St}(p, n) = \{X \in \mathbb{R}^{n \times p} | X^T X = I_p\}$
- Grassmann manifold: Set of all $p$-dimensional subspaces of $\mathbb{R}^n$
- Set of fixed rank $m$-by-$n$ matrices
- And many more

# Riemannian Manifolds

Roughly, a Riemannian manifold $\mathcal{M}$ is a smooth set with a smoothly-varying inner product on the tangent spaces.

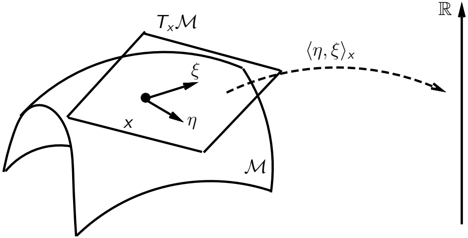

# Applications

Four applications are used to demonstrate the importances of the Riemannian optimization:

- Independent component analysis [CS93]
- Matrix completion problem [Van13]
- Geometric mean of symmetric positive definite matrices [ALM04, JVV12a, CS15]
- Elastic shape analysis of curves [SKJJ11, HGSA15]

# Application: Independent Component Analysis

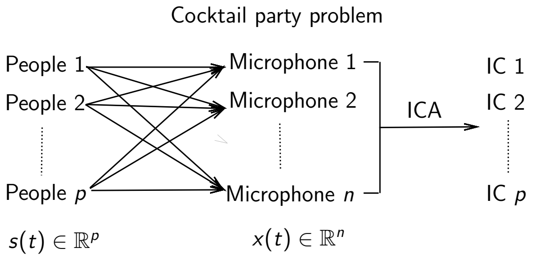


- Observed signal is $x(t) = As(t)$
- One approach:
  - Assumption: $E\{s(t)s(t+\tau)\}$ is diagonal for all $\tau$
  - $C_\tau(x) := E\{x(t)x(x+\tau)^T\} = AE\{s(t)s(t+\tau)^T\}A^T$

# Application: Independent Component Analysis

- Minimize joint diagonalization cost function on the Stiefel manifold [TI06]:

$$f : \mathrm{St}(p, n) \to \mathbb{R} : V \mapsto \sum_{i=1}^{N} \|V^T C_i V - \mathrm{diag}(V^T C_i V)\|_F^2.$$

- $C_1, \dots, C_N$ are covariance matrices and $\mathrm{St}(p, n) = \{X \in \mathbb{R}^{n \times p} | X^T X = I_p\}$.

# Application: Matrix Completion Problem

Matrix completion problem

| | Movie 1 | Movie 2 | | | | Movie $n$ |
|---|---|---|---|---|---|---|
| User 1 | | 1 | | | 4 | |
| User 2 | 3 | 5 | | | 4 | |
| | | | 5 | 1 | | |
| User $m$ | | 2 | | | 5 | 3 |

Rate matrix $M$

- The matrix $M$ is sparse
- The goal: complete the matrix $M$

# Application: Matrix Completion Problem

movies meta-user meta-movie

$$\begin{pmatrix} a_{11} & & & a_{14} \\ & & & a_{24} \\ & & a_{33} & \\ a_{41} & & & \\ & a_{52} & a_{53} & \end{pmatrix} = \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \\ b_{41} & b_{42} \\ b_{51} & b_{52} \end{pmatrix} \begin{pmatrix} c_{11} & c_{12} & c_{13} & c_{14} \\ c_{21} & c_{22} & c_{23} & c_{24} \end{pmatrix}$$

- Minimize the cost function

$$f : \mathbb{R}_r^{m\times n} \to \mathbb{R} : X \mapsto f(X) = \|P_\Omega M - P_\Omega X\|_F^2.$$

- $\mathbb{R}_r^{m\times n}$ is the set of $m$-by-$n$ matrices with rank $r$. It is known to be a Riemannian manifold.

# Application: Geometric Mean of Symmetric Positive Definite (SPD) Matrices

Computing the mean of a population of SPD matrices is important in medical imaging, image processing, radar signal processing, and elasticity. The desired properties are given in the ALM[1] list, some of which are

- if $A_1, \ldots, A_k$ commute, then $G(A_1, \ldots, A_k) = (A_1 \ldots A_k)^{\frac{1}{k}}$;
- $G(A_{\pi(1)}, \ldots, A_{\pi(k)}) = G(A_1, \ldots, A_k)$, with $\pi$ a permutation of $(1, \ldots, k)$;
- $G(A_1, \ldots, A_k) = G\left(A_1^{-1}, \ldots A_k^{-1}\right)^{-1}$;
- $\det G(A_1, \ldots, A_k) = (\det A_1 \ldots \det A_k)^{\frac{1}{k}}$;

where $A_1, \ldots, A_k$ are SPD matrices, and $G(\cdot, \ldots, \cdot)$ denotes the geometric mean of arguments.

[1] T. Ando, C.-K. Li, and R. Mathias, Geometric means, *Linear Algebra and Its Applications*, 385:305-334, 2004

# Application: Geometric Mean of Symmetric Positive Definite Matrices

One geometric mean is the Karcher mean of the manifold of SPD matrices with the affine invariant metric, i.e.,

$$G(A_1, \dots, A_k) = \arg \min_{X \in \mathbb{S}_+^n} \frac{1}{2k} \sum_{i=1}^{k} \mathrm{dist}^2(X, A_i),$$

where $\mathrm{dist}(X, Y) = \| \log(X^{-1/2} Y X^{-1/2}) \|_F$ is the distance under the Riemannian metric

$$g(\eta_X, \xi_X) = \mathrm{trace}(\eta_X X^{-1} \xi_X X^{-1}).$$

# Application: Elastic Shape Analysis of Curves

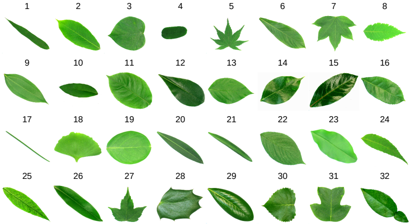


- Classification [LKS+12, HGSA15]
- Face recognition [DBS+13]

# Application: Elastic Shape Analysis of Curves

- Elastic shape analysis invariants:
    - Rescaling
    - Translation
    - Rotation
    - Reparametrization
- The shape space is a quotient space

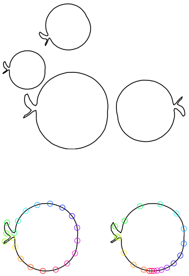

Figure: All are the same shape.

# Application: Elastic Shape Analysis of Curves

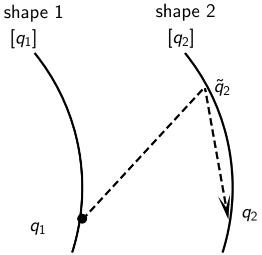


- Optimization problem $\min_{q_2 \in [q_2]} \mathrm{dist}(q_1, q_2)$ is defined on a Riemannian manifold
- Computation of a geodesic between two shapes
- Computation of Karcher mean of a population of shapes

# More Applications

- Role model extraction [MHB+16]
- Computations on SPD matrices [YHAG17]
- Phase retrieval problem [HGZ17]
- Blind deconvolution [HH17]
- Synchronization of rotations [Hua13]
- Computations on low-rank tensor
- Low-rank approximate solution for Lyapunov equation

# Iterations on the Manifold

Consider the following generic update for an iterative Euclidean optimization algorithm:

$$x_{k+1} = x_k + \Delta x_k = x_k + \alpha_k s_k \ .$$

This iteration is implemented in numerous ways, e.g.:

- Steepest descent: $x_{k+1} = x_k - \alpha_k \nabla f(x_k)$
- Newton's method: $x_{k+1} = x_k - \left[\nabla^2 f(x_k)\right]^{-1} \nabla f(x_k)$
- Trust region method: $\Delta x_k$ is set by optimizing a local model.

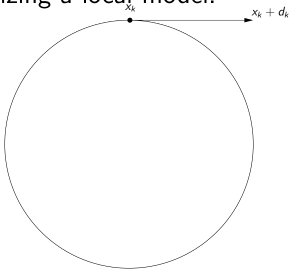


## Riemannian Manifolds Provide

- Riemannian concepts describing directions and movement on the manifold
- Riemannian analogues for gradient and Hessian

# Riemannian gradient and Riemannian Hessian

## Definition

The Riemannian gradient of $f$ at $x$ is the unique tangent vector in $T_x M$ satisfying $\forall \eta \in T_x M$, the directional derivative

$$\mathrm{D}\, f(x)[\eta] = \langle \operatorname{grad} f(x), \eta \rangle$$

and $\operatorname{grad} f(x)$ is the direction of steepest ascent.

## Definition

The Riemannian Hessian of $f$ at $x$ is a symmetric linear operator from $T_x M$ to $T_x M$ defined as

$$\operatorname{Hess} f(x) : T_x M \to T_x M : \eta \to \nabla_\eta \operatorname{grad} f,$$

where $\nabla$ is the affine connection.

# Retractions

| Euclidean | Riemannian |
|---|---|
| $x_{k+1} = x_k + \alpha_k d_k$ | $x_{k+1} = R_{x_k}(\alpha_k \eta_k)$ |

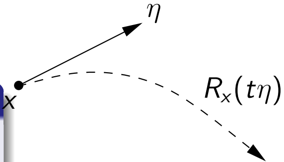


**Definition**

A retraction is a mapping $R$ from $TM$ to $M$ satisfying the following:

- $R$ is continuously differentiable
- $R_x(0) = x$
- $\mathrm{D}\, R_x(0)[\eta] = \eta$

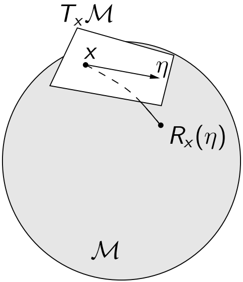


- maps tangent vectors back to the manifold
- defines curves in a direction

# Categories of Riemannian optimization methods

**Retraction-based: local information only**

Line search-based: use local tangent vector and $R_x(t\eta)$ to define line

- Steepest decent
- Newton

Local model-based: series of flat space problems

- Riemannian trust region Newton (RTR)
- Riemannian adaptive cubic overestimation (RACO)

# Categories of Riemannian optimization methods

**Retraction and transport-based: information from multiple tangent spaces**

- Nonlinear conjugate gradient: multiple tangent vectors
- Quasi-Newton e.g. Riemannian BFGS: transport operators between tangent spaces

Additional element required for optimizing a cost function $(M, g)$:

- formulas for combining information from multiple tangent spaces.

# Vector Transports

## Vector Transport

- Vector transport: Transport a tangent vector from one tangent space to another
- $\mathcal{T}_{\eta_x}\xi_x$, denotes transport of $\xi_x$ to tangent space of $R_x(\eta_x)$. $R$ is a retraction associated with $\mathcal{T}$

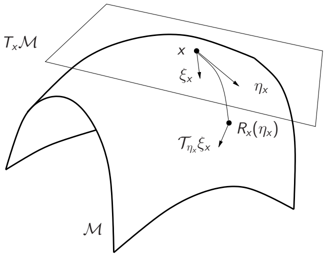


Figure: Vector transport.

# Retraction/Transport-based Riemannian Optimization

Given a retraction and a vector transport, we can generalize many Euclidean methods to the Riemannian setting. Do the Riemannian versions of the methods work well?

# Retraction/Transport-based Riemannian Optimization

Given a retraction and a vector transport, we can generalize many Euclidean methods to the Riemannian setting. Do the Riemannian versions of the methods work well?

No

- Lose many theoretical results and important properties;
- Impose restrictions on retraction/vector transport;

# Riemannian optimization methods

Elements required for optimizing a cost function $(M, g)$:

- an representation for points $x$ on $M$, for tangent spaces $T_x M$, and for the inner products $g_x(\cdot, \cdot)$ on $T_x M$;
- choice of a retraction $R_x : T_x M \to M$;
- formulas for $f(x), \mathrm{grad}\ f(x)$ and $\mathrm{Hess}\ f(x)$ (or its action);
- Computational and storage efficiency;

# Comparison with Constrained Optimization

- All iterates on the manifold
- Convergence properties of unconstrained optimization algorithms
- No need to consider Lagrange multipliers or penalty functions
- Exploit the structure of the constrained set

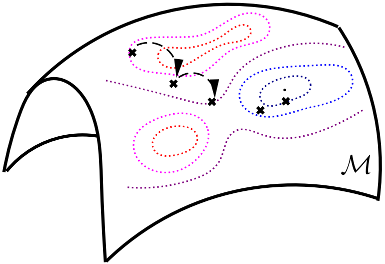

# Retraction/Transport-based Riemannian Optimization

## Benefits

- Increased generality does not compromise the important theory
- Less expensive than or similar to previous approaches
- May provide theory to explain behavior of algorithms specifically developed for a particular application – or closely related ones

## Possible Problems

- May be inefficient compared to algorithms that exploit application details

# Some History of Optimization On Manifolds (I)

Luenberger (1973), *Introduction to linear and nonlinear programming*. Luenberger mentions the idea of performing line search along geodesics, "which we would use if it were computationally feasible (which it definitely is not)". Rosen (1961) essentially anticipated this but was not explicit in his Gradient Projection Algorithm.

Gabay (1982), *Minimizing a differentiable function over a differential manifold*. Steepest descent along geodesics; Newton's method along geodesics; Quasi-Newton methods along geodesics. On Riemannian submanifolds of $\mathbb{R}^n$.

Smith (1993-94), *Optimization techniques on Riemannian manifolds*. Levi-Civita connection $\nabla$; Riemannian exponential mapping; parallel translation.

# Some History of Optimization On Manifolds (II)

The "pragmatic era" begins:

Manton (2002), *Optimization algorithms exploiting unitary constraints*
"The present paper breaks with tradition by not moving along geodesics". The geodesic update $\mathrm{Exp}_x \eta$ is replaced by a projective update $\pi(x + \eta)$, the *projection* of the point $x + \eta$ onto the manifold.

Adler, Dedieu, Shub, et al. (2002), *Newton's method on Riemannian manifolds and a geometric model for the human spine*. The exponential update is relaxed to the general notion of *retraction*. The geodesic can be replaced by any (smoothly prescribed) curve tangent to the search direction.

# Some History of Optimization On Manifolds (III)

Theory, efficiency, and library design improve dramatically:

Absil, Baker, Gallivan (2004-07), Theory and implementations of Riemannian Trust Region method. Retraction-based approach. Matrix manifold problems, software repository:
`http://www.math.fsu.edu/~cbaker/GenRTR`
Anasazi Eigenproblem package in Trilinos Library at Sandia National Laboratory

Ring and With (2012), combination of differentiated retraction and isometric vector transport for convergence analysis of RBFGS

Absil, Gallivan, Huang (2009-2017), Complete theory of Riemannian Quasi-Newton and related transport/retraction conditions, Riemannian SR1 with trust-region, RBFGS on partly smooth problems, A C++ library: `http://www.math.fsu.edu/~whuang2/ROPTLIB`

# Some History of Optimization On Manifolds (IV)

Absil, Mahony, Sepulchre (2007) Nonlinear conjugate gradient using retractions.

Ring and With (2012), Global convergence analysis for Fletcher-Reeves Riemannian nonlinear CG method with the strong wolfe conditions under a strong assumption.

Sato, Iwai (2013-2015), Global convergence analysis for Fletcher-Reeves type Riemannian nonlinear CG method with the strong wolfe conditions under a mild assumption; and global convergence for Dai-Yuan type Riemannian nonlinear CG method with the weak wolfe conditions under mild assumptions.

Zhu (2017), Global convergence for Riemannian version of Dai's nonmonotone nonlinear CG method.

# Some History of Optimization On Manifolds (V)

Bonnabel (2011), Riemannian stochastic gradient descent method.

Sato, Kasai, Mishra(2017), Riemannian stochastic gradient descent method using variance reduction or quasi-Newton.

Becigneul, Ganea(2018), Riemannian versions of ADAM, ADAGRAD, and AMSGRAD for geodesically convex functions.

Zhang, Sra(2016-2018), Riemannian first-order methods for geodesically convex optimization.

Liu, Boumal(2019), Riemannian optimization with constraints.

# Some History of Optimization On Manifolds (VI)

Hosseini, Grohs, Huang, Uschmajew, Boumal, (2015-2016), Lipschitz-continuous functions on Riemannian manifolds

Zhang, Sra(2016-2018), Riemannian first-order methods for geodesically convex optimization.

Bento, Ferreira, Melo(2017), Riemannian proximal point method for geodesically convex optimization.

Chen, Ma, So, Zhang(2018), Riemannian proximal gradient method.

Many people  Application interests start to increase noticeably

# Riemannian Optimization Libraries

Riemannian optimization libraries for general problems:

- Boumal, Mishra, Absil, Sepulchre(2014)
  Manopt (Matlab library)
- Townsend, Koep, Weichwald (2016)
  Pymanopt (Python version of manopt)
- Huang, Absil, Gallivan, Hand (2018)
  ROPTLIB (C++ library, interfaces to Matlab and Julia)
- Martin, Raim, Huang, Adragni(2018)
  ManifoldOptim (R wrapper of ROPTLIB)
- Meghwanshi, Jawanpuria, Kunchukuttan, Kasai, Mishra (2018)
  McTorch (Riemannian optimization for deep learning)

# Application: Averaging Symmetric Positive Definite matrices

Joint work with:

- Pierre-Antoine Absil, Professor of Mathematical Engineering, *Université catholique de Louvain*
- Kyle A. Gallivan, Professor of Mathematics, *Florida State University*
- Xinru Yuan, Data scientist, *Esurance*

# Application: Averaging Symmetric Positive Definite matrices

## Definition

*A symmetric matrix $A$ is called* positive definite *$A \succ 0$ iff all its eigenvalues are positive.*

$$\mathcal{S}^{\mathrm{n}}_{++} = \{A \in \mathbb{R}^{n \times n} : A = A^T, A \succ 0\}$$

$2 \times 2$ SPD matrix

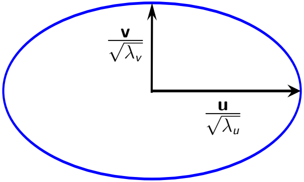


$3 \times 3$ SPD matrix

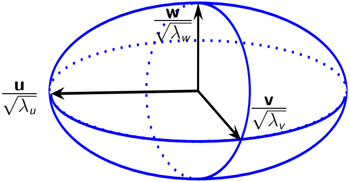

# Motivation of Averaging SPD Matrices

- Possible applications of SPD matrices
  - Diffusion tensors in medical imaging [CSV12, FJ07, RTM07]
  - Describing images and video [LWM13, SFD02, ASF+05, TPM06, HWSC15]

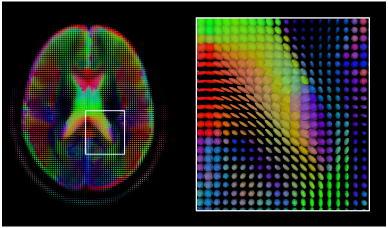

- Motivation of averaging SPD matrices
  - Aggregate several noisy measurements of the same object
  - Subtask in interpolation methods, segmentation, and clustering

# Averaging Schemes: from Scalars to Matrices

Let $A_1, \ldots, A_K$ be SPD matrices.

- Generalized arithmetic mean: $\frac{1}{K}\sum_{i=1}^{K} A_i$

  $\rightarrow$ Not appropriate in many practical applications

# Averaging Schemes: from Scalars to Matrices

Let $A_1, \ldots, A_K$ be SPD matrices.

- Generalized arithmetic mean: $\frac{1}{K} \sum_{i=1}^{K} A_i$

  $\rightarrow$ Not appropriate in many practical applications

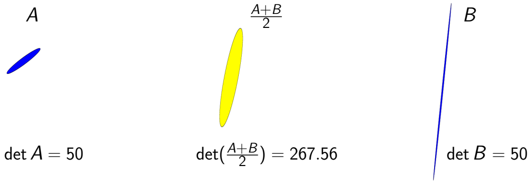

# Averaging Schemes: from Scalars to Matrices

Let $A_1, \ldots, A_K$ be SPD matrices.

- Generalized arithmetic mean: $\frac{1}{K} \sum_{i=1}^{K} A_i$

  $\rightarrow$ Not appropriate in many practical applications

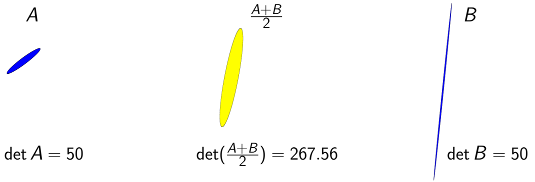


- Generalized geometric mean: $(A_1 \cdots A_K)^{1/K}$

  $\rightarrow$ Not appropriate due to non-commutativity

  $\rightarrow$ How to define a matrix geometric mean?

# Desired Properties of a Matrix Geometric Mean

The desired properties are given in the ALM list[2], some of which are:

- $G(A_{\pi(1)}, \ldots, A_{\pi(K)}) = G(A_1, \ldots, A_K)$ with $\pi$ a permutation of $(1, \ldots, K)$
- if $A_1, \ldots, A_K$ commute, then $G(A_1, \ldots, A_K) = (A_1, \ldots, A_K)^{1/K}$
- $G(A_1, \ldots, A_K)^{-1} = G(A_1^{-1}, \ldots, A_K^{-1})$
- $\det(G(A_1, \ldots, A_K)) = (\det(A_1) \cdots \det(A_K))^{1/K}$

[2] T. Ando, C.-K. Li, and R. Mathias, *Geometric means*, Linear Algebra and Its Applications, 385:305-334, 2004

# Geometric Mean of SPD Matrices

- A well-known mean on the manifold of SPD matrices is the **Karcher mean** [Kar77]:

$$G(A_1, \ldots, A_K) = \arg \min_{X \in \mathcal{S}_{++}^n} \frac{1}{2K} \sum_{i=1}^{K} \delta^2(X, A_i), \qquad (1)$$

where $\delta(X, Y) = \| \log(X^{-1/2} Y X^{-1/2}) \|_F$ is the geodesic distance under the affine-invariant metric

$$g(\eta_X, \xi_X) = \mathrm{trace}(\eta_X X^{-1} \xi_X X^{-1})$$

- The Karcher mean defined in (1) satisfies all the geometric properties in the ALM list [LL11]

# Algorithms

$$G(A_1, \ldots, A_k) = \operatorname*{argmin}_{X \in \mathcal{S}^n_{++}} \frac{1}{2K} \sum_{i=1}^{K} \delta^2(X, A_i),$$

- Riemannian steepest descent [RA11, Ren13]
- Riemannian Barzilai-Borwein method [IP15]
- Riemannian Newton method [RA11]
- Richardson-like iteration [BI13]
- Riemannian steepest descent, conjugate gradient, BFGS, and trust region Newton methods [JVV12b]
- Limited-memory Riemannian BFGS method [YHAG16]

# Remarks

Previous work:

- Riemannian steepest descent and Riemannian CG methods are preferred in terms of computational time
- High rate of convergence of Riemannian Newton method does not make up for extra complexity

New results:

- Explain the preference for the first order methods
- More options of retractions and vector transports
- More efficient implementation
- Limited-memory Riemannian BFGS method

# Conditioning of the Objective Function

Hemstitching phenomenon
for steepest descent

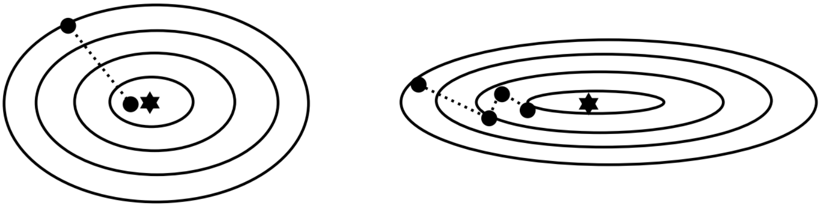

well-conditioned Hessian

ill-conditioned Hessian

- Small condition number $\Rightarrow$ fast convergence
- Large condition number $\Rightarrow$ slow convergence

# Conditioning of the Karcher Mean Objective Function

- **Riemannian metric**:

$$g_X(\xi,\eta) = \mathrm{trace}(\xi X^{-1}\eta X^{-1})$$

- **Euclidean metric**:

$$g_X(\xi,\eta) = \mathrm{trace}(\xi\eta)$$

**Condition number $\kappa$ of Hessian at the minimizer $\mu$:**

- Hessian of Riemannian metric:
  - $\kappa(H^R) \leq 1 + \frac{\ln(\max \kappa_i)}{2}$, where $\kappa_i = \kappa(\mu^{-1/2} A_i \mu^{-1/2})$
  - $\kappa(H^R) \leq 20$ if $\max(\kappa_i) = 10^{16}$

- Hessian of Euclidean metric:
  - $\frac{\kappa^2(\mu)}{\kappa(H^{\mathrm{R}})} \leq \kappa(H^{\mathrm{E}}) \leq \kappa(H^{\mathrm{R}})\kappa^2(\mu)$
  - $\kappa(H^E) \geq \kappa^2(\mu)/20$

# Implementations

- Retraction
  - Exponential mapping: $\mathrm{Exp}_X(\xi) = X^{1/2} \exp(X^{-1/2} \xi X^{-1/2}) X^{1/2}$
  - Second order approximation retraction [JVV12b]:
    $R_X(\xi) = X + \xi + \frac{1}{2} \xi X^{-1} \xi$
- Vector transport
  - Parallel translation: $\mathcal{T}_{p_\eta}(\xi) = Q \xi Q^T$, with $Q = X^{\frac{1}{2}} \exp(\frac{X^{-\frac{1}{2}} \eta X^{-\frac{1}{2}}}{2}) X^{-\frac{1}{2}}$
  - Vector transport by parallelization [HAG16]: essentially an identity
  - Requires orthogonal basis for tangent spaces

# Implementations

- Cholesky $X_k = L_k L_k^T$ assumed to be computed on each step
- $B_X$ of $\mathrm{T}_X \mathcal{S}_{++}^n$, the orthonormal basis of $\mathrm{T}_X \mathcal{S}_{++}^n$

$$B_X = \{L e_i e_i^T L^T : i = 1, \dots, n\} \cup \{\frac{1}{\sqrt{2}} L(e_i e_j^T + e_j e_i^T) L^T,$$
$$i < j, \ i = 1, \dots, n, \ j = 1, \dots, n\},$$

!where $\{e_i, \dots, e_n\}$ is the standard basis of $n$-dimensional Euclidean space.

- orthonormal under $g_X(\xi_X, \eta_X)$.
- $\xi_X = B_X \hat{\xi}_X \leftrightarrow \xi_X = LSL^T$, where $S$ is symmetric and constains scale coefficients.
- intrinsic representation of tangent vectors is easily maintained.

# Numerical Results: $K = 100$, size $= 3 \times 3$, $d = 6$

- $1 \leq \kappa(A_i) \leq 200$

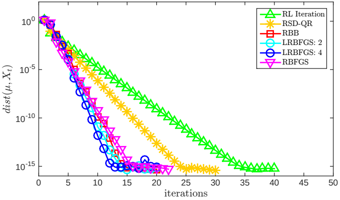


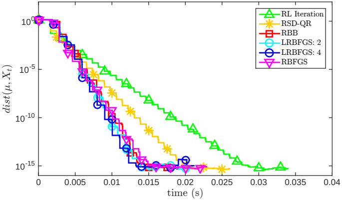


- $10^3 \leq \kappa(A_i) \leq 10^7$

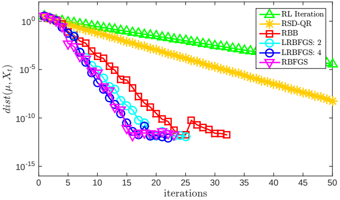


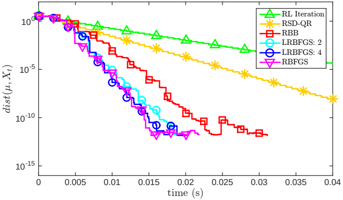

# Numerical Results: $K = 30$, size $= 100 \times 100$, $d = 5050$

- $1 \leq \kappa(A_i) \leq 20$

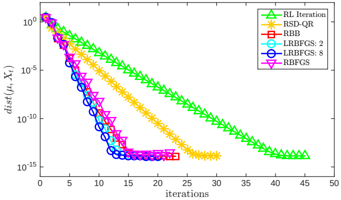


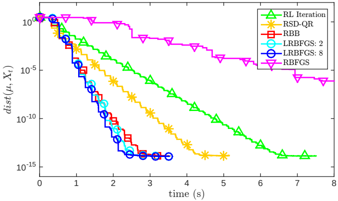


- $10^4 \leq \kappa(A_i) \leq 10^7$

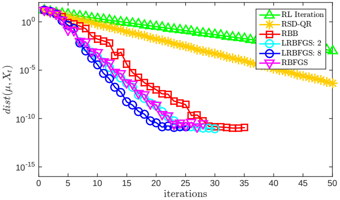


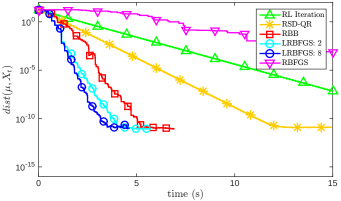

# Numerical Results: Riemannian vs. Euclidean Metrics

- $K = 100$, $n = 3$, and $1 \leq \kappa(A_i) \leq 10^6$.

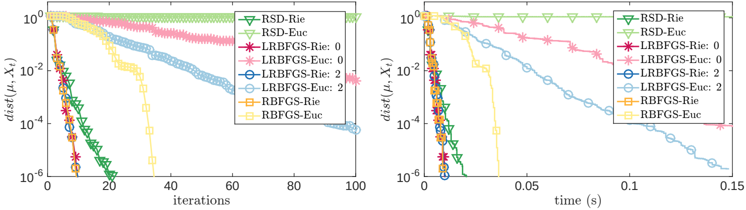


- $K = 30$, $n = 100$, and $1 \leq \kappa(A_i) \leq 10^5$.

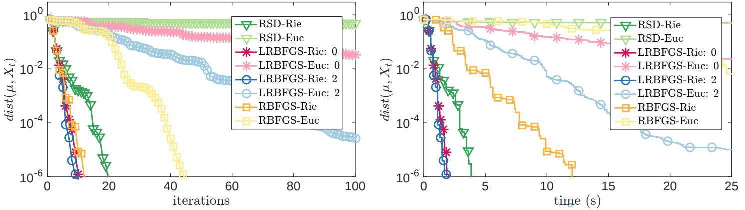

# Summary

- Introduced the framework of Riemannian optimization
- Used applications to show the importance of Riemannian optimization
- Briefly reviewed the history of Riemannian optimization
- Introduced the mean of SPD matrices
- Demonstrated the performance of the Riemannian methods

# Thank you

Thank you!

# References I

T. Ando, C. K. Li, and R. Mathias.
Geometric means.
*Linear Algebra and Its Applications*, 385:305–334, 2004.

Ognjen Arandjelovic, Gregory Shakhnarovich, John Fisher, Roberto Cipolla, and Trevor Darrell.
Face recognition with image sets using manifold density divergence.
In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 581–588. IEEE, 2005.

D. A. Bini and B. Iannazzo.
Computing the Karcher mean of symmetric positive definite matrices.
*Linear Algebra and its Applications*, 438(4):1700–1710, 2013.

J. F. Cardoso and A. Souloumiac.
Blind beamforming for non-gaussian signals.
*IEE Proceedings F Radar and Signal Processing*, 140(6):362, 1993.

A. Cherian and S. Sra.
Riemannian dictionary learning and sparse coding for positive definite matrices.
*CoRR*, abs/1507.02772, 2015.

Guang Cheng, Hesamoddin Salehian, and Baba Vemuri.
Efficient recursive algorithms for computing the mean diffusion tensor and applications to DTI segmentation.
*Computer Vision–ECCV 2012*, pages 390–401, 2012.

H. Drira, B. Ben Amor, A. Srivastava, M. Daoudi, and R. Slama.
3D face recognition under expressions, occlusions, and pose variations.
*Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(9):2270–2283, 2013.

# References II

P. T. Fletcher and S. Joshi.
Riemannian geometry for the statistical analysis of diffusion tensor data.
*Signal Processing*, 87(2):250–262, 2007.

W. Huang, P.-A. Absil, and K. A. Gallivan.
Intrinsic representation of tangent vectors and vector transport on matrix manifolds.
*Numerische Mathematik*, 136(2):523–543, 2016.

W. Huang, K. A. Gallivan, Anuj Srivastava, and P.-A. Absil.
Riemannian optimization for registration of curves in elastic shape analysis.
*Journal of Mathematical Imaging and Vision*, 54(3):320–343, 2015.
DOI:10.1007/s10851-015-0606-8.

Wen Huang, K. A. Gallivan, and Xiangxiong Zhang.
Solving PhaseLift by low rank Riemannian optimization methods for complex semidefinite constraints.
*SIAM Journal on Scientific Computing*, 39(5):B840–B859, 2017.

W. Huang and P. Hand.
Blind deconvolution by a steepest descent algorithm on a quotient manifold, 2017.
In preparation.

W. Huang.
*Optimization algorithms on Riemannian manifolds with applications*.
PhD thesis, Florida State University, Department of Mathematics, 2013.

Zhiwu Huang, Ruiping Wang, Shiguang Shan, and Xilin Chen.
Face recognition on large-scale video in the wild with hybrid Euclidean-and-Riemannian metric learning.
*Pattern Recognition*, 48(10):3113–3124, 2015.

# References III

Bruno Iannazzo and Margherita Porcelli.
The Riemannian Barzilai-Borwein method with nonmonotone line-search and the Karcher mean computation.
*Optimization online, December*, 2015.

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.

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.

H. Karcher.
Riemannian center of mass and mollifier smoothing.
*Communications on Pure and Applied Mathematics*, 1977.

H. Laga, S. Kurtek, A. Srivastava, M. Golzarian, and S. J. Miklavcic.
A Riemannian elastic metric for shape-based plant leaf classification.
*2012 International Conference on Digital Image Computing Techniques and Applications (DICTA)*, pages 1–7, December 2012. doi:10.1109/DICTA.2012.6411702.

J. Lawson and Y. Lim.
Monotonic properties of the least squares mean.
*Mathematische Annalen*, 351(2):267–279, 2011.

Jiwen Lu, Gang Wang, and Pierre Moulin.
Image set classification using holistic multiple order statistics features and localized multi-kernel metric learning.
In *Proceedings of the IEEE International Conference on Computer Vision*, pages 329–336, 2013.

# References IV

Melissa Marchand, Wen Huang, Arnaud Browet, Paul Van Dooren, and Kyle A. Gallivan.
A riemannian optimization approach for role model extraction.
In *Proceedings of the 22nd International Symposium on Mathematical Theory of Networks and Systems*, pages 58–64, 2016.

Q. Rentmeesters and P.-A. Absil.
Algorithm comparison for Karcher mean computation of rotation matrices and diffusion tensors.
In *19th European Signal Processing Conference*, pages 2229–2233, Aug 2011.

Q. Rentmeesters.
*Algorithms for data fitting on some common homogeneous spaces*.
PhD thesis, Universite catholiqué de Louvain, 2013.

Y. Rathi, A. Tannenbaum, and O. Michailovich.
Segmenting images on the tensor manifold.
In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2007.

Gregory Shakhnarovich, John W Fisher, and Trevor Darrell.
Face recognition from long-term observations.
In *European Conference on Computer Vision*, pages 851–865. Springer, 2002.

A. Srivastava, E. Klassen, S. H. Joshi, and I. H. Jermyn.
Shape analysis of elastic curves in Euclidean spaces.
*IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(7):1415–1428, September 2011.
doi:10.1109/TPAMI.2010.184.

F. J. Theis and Y. Inouye.
On the use of joint diagonalization in blind signal processing.
*2006 IEEE International Symposium on Circuits and Systems*, (2):7–10, 2006.

# References V

Oncel Tuzel, Fatih Porikli, and Peter Meer.
Region covariance: A fast descriptor for detection and classification.
In *European conference on computer vision*, pages 589–600. Springer, 2006.

B. Vandereycken.
Low-rank matrix completion by Riemannian optimization—extended version.
*SIAM Journal on Optimization*, 23(2):1214–1236, 2013.

Xinru Yuan, Wen Huang, P.-A. Absil, and Kyle A. Gallivan.
A Riemannian limited-memory BFGS algorithm for computing the matrix geometric mean.
*Procedia Computer Science*, 80:2147–2157, 2016.

Xinru Yuan, Wen Huang, P.-A. Absil, and K. A. Gallivan.
A Riemannian quasi-newton method for computing the Karcher mean of symmetric positive definite matrices.
Technical Report FSU17-02, Florida State University, 2017.
www.math.fsu.edu/ whuang2/papers/RMKMSPDM.htm.