

Riemannian Optimization for Elastic Shape Analysis

W. Huang², K. A. Gallivan¹, A. Srivastava¹ and P.-A. Absil²

¹Florida State University

²Catholic University of Louvain

Old Dominion University 23 January 2015

Riemannian Optimization

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

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

where \mathcal{M} is a Riemannian manifold.

Goal

Adapt unconstrained Euclidean algorithms to function on \mathcal{M} .

Motivation

Matrix optimization with constraints

- Procrustes: given $A \in \mathbb{R}^{m \times m}$, $B \in \mathbb{R}^{n \times n}$

$$\min_Q \|AQ - QB\|, \quad s.t. Q \in \mathbb{R}^{m \times n}, \quad Q^T Q = I_n$$

- $AV = BV\Lambda$

$$\min \text{trace}(V^T AV), \quad s.t. V \in \mathbb{R}^{n \times k}, \quad V^T V = I_k$$

- Given $x' = Ax + Bu$, $y = Cx$ find $W, V \in \mathbb{R}^{n \times k}$

$$\hat{x}' = \hat{A}\hat{x} + \hat{B}u, \quad \hat{y} = \hat{C}\hat{x}$$

$$\hat{A} = W^T AV, \quad \hat{B} = W^T B, \quad \hat{C} = CV$$

$$H(s) \approx \hat{H}(s)$$

Motivation

Reimposing constraints on each step:

- unit length

$$t_k = A^{-1}x_{k-1}, \quad x_k = t_k / \|t_k\|$$

- orthonormal basis


$$X_k = Q_{k-1} + \Delta_{k-1}$$
$$X_k \rightarrow Q_k, \quad QR \text{ factorization, SVD, etc.}$$

- orthogonal matrix

$$S_k = S_{k-1} + \Delta$$
$$Q_k = \text{Exp}(S_k)$$

Riemannian Manifolds

Roughly, a Riemannian manifold is a smooth set with a smoothly-varying inner product on the tangent spaces.



Noteworthy Manifolds and Applications

- Grassmann manifold

$\text{Gr}(p, n) = \text{Set of all } p\text{-dimensional subspaces of } \mathbb{R}^n \quad d = np - p^2$

- Stiefel manifold

$$\text{St}(p, n) = \{X \in \mathbb{R}^{n \times p} : X^T X = I_p\} \quad d = np - \frac{p(p+1)}{2}$$

- Oblique manifold

$$\mathbb{R}_*^{n \times p} / \mathcal{S}_{\text{diag}+} \simeq \{Y \in \mathbb{R}_*^{n \times p} : \text{diag}(Y^T Y) = I_p\} = S^{n-1} \times \dots \times S^{n-1}$$
$$d = (n-1)p^2$$

Noteworthy Manifolds and Applications

- Set of fixed-rank PSD matrices $S_+(p, n)$. A quotient representation:

$$X \sim Y \Leftrightarrow \exists Q \in O_p : Y = XQ$$

- Flag manifold

Elements of the flag manifold can be viewed as a p -tuple of linear subspaces $(\mathcal{V}_1, \dots, \mathcal{V}_p)$ such that $\dim(\mathcal{V}_i) = i$ and $\mathcal{V}_i \subset \mathcal{V}_{i+1}$.

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: $x_{k+1} = x_k - [\nabla^2 f(x_k)]^{-1} \nabla f(x_k)$
- Trust region method: Δx_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**



Tangent Vectors

- The concept of direction is provided by tangent vectors.
- **Intuitively**, tangent vectors are tangent to curves on the manifold.
- Tangent vectors are an **intrinsic** property of a differentiable manifold.




Definition

The tangent space $T_x M$ is the vector space comprised of the tangent vectors at $x \in M$. The Riemannian metric is an inner product on each tangent space.

Tangent Vectors

- The concept of direction is provided by tangent vectors.
- Intuitively, tangent vectors are tangent to curves on the manifold.
- Tangent vectors are an intrinsic property of a differentiable manifold.



Definition

The **tangent space** $T_x M$ is the vector space comprised of the tangent vectors at $x \in M$. The **Riemannian metric** is an inner product on each tangent space.

Example Tangent Spaces and Metrics

Stiefel manifold embedded in $\mathbb{R}^{n \times p}$

$$\text{St}(p, n) = \{X \in \mathbb{R}^{n \times p} : X^T X = I_p\}$$

$V \in T_X \text{St}(p, n)$ $X^T V + V^T X = 0$, i.e., $X^T V$ is skew symmetric

$V = XA + X_\perp B$ $-A = A^T$, i.e., A is skew symmetric

Euclidean metric: $g_e(V, V) = \text{tr} V^T V$

Canonical metric: $g_c(V, V) = \text{tr} V^T (I - \frac{1}{2} X X^T) V = \frac{1}{2} \text{tr} A^T A + \text{tr} B^T B$

Riemannian gradient and Riemannian Hessian

Definition

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

$$Df(x)[\eta] = \langle \mathbf{grad} f(x), \eta \rangle$$

and $\mathbf{grad} f(x)$ 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

$$\text{Hess} f(x) : T_x M \rightarrow T_x M : \eta \rightarrow \nabla_\eta \mathbf{grad} f$$

Retractions



Definition

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

- R is continuously differentiable
- $R_x(0) = x$
- $D R_x(0)[\eta] = \eta$
- maps tangent vectors back to the manifold
- lifts objective function f from M to $T_x M$, via the **pullback**

$$\hat{f}_x = f \circ R_x$$

- defines curves in a direction
- exponential map $\text{Exp}(t\eta)$ defines “straight lines” **geodesic**




Generic Riemannian Optimization Algorithm

1. At iterate $x \in M$, define $\hat{f}_x = f \circ R_x$.
2. Find $\eta \in T_x M$ which satisfies certain condition.
3. Choose new iterate $x_+ = R_x(\eta)$.
4. Goto step 1.

A suitable setting

This paradigm is sufficient for describing many optimization methods.



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: geodesic in the direction $-\text{grad } f(x)$
- Newton

Local model-based: series of flat space problems

- Riemannian Trust region (RTR)
- Riemannian Adaptive Cubic Overestimation (RACO)

Retraction and transport-based: information from multiple tangent spaces

- Conjugate gradient and accelerated iteration: multiple tangent vectors
- Quasi-Newton e.g. Riemannian BFGS: transport operators between tangent spaces

Basic principles

All/some elements required for optimizing a cost function (M, g) :

- an efficient numerical 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 $R_x : T_x M \rightarrow M$;
- formulas for $f(x)$, $\text{grad } f(x)$ and $\text{Hess } f(x)$ (or its action);
- formulas for combining information from multiple tangent spaces.

Parallel transport




Figure : Parallel transport

- Parallel transport one tangent vector along some curve $Y(t)$.
- It is often along the geodesic $\gamma_\eta(t) : \mathbb{R} \rightarrow M : t \rightarrow \text{Exp}_x(t\eta_x)$.
- In general, geodesics and parallel translation require solving an ordinary differential equation.

Vector transport

Definition

We define a **vector transport** on a manifold M to be a **smooth mapping**

$$TM \oplus TM \rightarrow TM : (\eta_x, \xi_x) \mapsto \mathcal{T}_{\eta_x}(\xi_x) \in TM$$

satisfying the following properties for all $x \in M$.

- (Consistency) $\mathcal{T}_{0_x} \xi_x = \xi_x$ for all $\xi_x \in T_x M$;
- (Linearity) $\mathcal{T}_{\eta_x}(a\xi_x + b\zeta_x) = a\mathcal{T}_{\eta_x}(\xi_x) + b\mathcal{T}_{\eta_x}(\zeta_x)$.
- There exists a retraction R that is associated with \mathcal{T} and satisfies the following illustrated relationship:

Vector transport




Figure : Vector transport.

Retraction/Transport-based Riemannian Optimization

Benefits

- Increased generality does not compromise the **important theory**
- Can easily employ classical optimization techniques
- Less expensive than or similar to previous approaches
- May provide theory to explain behavior of algorithms in a particular application – or closely related ones

Possible Problems

- May be inefficient compared to algorithms that exploit application details

Retraction-based Riemannian optimization

Equivalence of the pullback $\hat{f}_x = f \circ R_x$

	Exp_x	R_x
$\text{grad } f(x) = \text{grad } \hat{f}_x(0)$	yes	yes
$\text{Hess } f(x) = \text{Hess } \hat{f}_x(0)$	yes	no
$\text{Hess } f(x) = \text{Hess } \hat{f}_x(0)$ at critical points	yes	yes

Sufficient Optimality Conditions

If $\text{grad } \hat{f}_x(0) = 0$ and $\text{Hess } \hat{f}_x(0) > 0$,
 then $\text{grad } f(x) = 0$ and $\text{Hess } f(x) > 0$,
 so that x is a **local minimizer** of f .

Retraction-based Riemannian optimization

Equivalence of the pullback $\hat{f}_x = f \circ R_x$

	Exp_x	R_x
$\text{grad } f(x) = \text{grad } \hat{f}_x(0)$	yes	yes
$\text{Hess } f(x) = \text{Hess } \hat{f}_x(0)$	yes	no
$\text{Hess } f(x) = \text{Hess } \hat{f}_x(0)$ at critical points	yes	yes

Sufficient Optimality Conditions

If $\text{grad } \hat{f}_x(0) = 0$ and $\text{Hess } \hat{f}_x(0) > 0$,
 then $\text{grad } f(x) = 0$ and $\text{Hess } f(x) > 0$,
 so that x is a **local minimizer** of f .

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.

[Bertsekas \(1976, 1982\)](#) A different Gradient Projection Algorithm and Projected Newton Algorithm along with extensive convergence analysis on convex and bounds constraints. Kelley (1999) extended this to BFGS for bounds constraints.

[Gabay \(1982\)](#), *Minimizing a differentiable function over a differential manifold*. Stepest 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 ∇ ; Riemannian exponential; parallel translation.

[Helmke and Moore \(1994\)](#), *Optimization techniques on Riemannian manifolds*. dynamical systems, flows on manifolds, SVD, balancing, eigenvalues

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 $\text{Exp}_x \eta$ 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.

[Absil, Mahony, Sepulchre \(2007\)](#) Nonlinear CG using retractions.

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

[Dresigmeyer \(2007\)](#), Nelder-Mead using geodesics, derivative-free algorithm

[Absil, Gallivan, Qi \(2007-10\)](#), Basic theory and implementations of Riemannian BFGS and Riemannian Adaptive Cubic Overestimation. Parallel translation and Exponential map theory, Retraction and vector transport empirical evidence.

Some History of Optimization On Manifolds (III)

[Ring and With \(2012\)](#), combination of differentiated retraction and isometric vector transport for RBFGS

[Absil, Gallivan, Huang \(2009-2015\)](#), Complete theory of Riemannian Quasi-Newton and related transport/retraction conditions, Riemannian SR1 with trust-region, RBFGS on partly smooth problems, new library (soon to be public).

[Many people](#) Application interests start to increase noticeably

Other Applications

- Large-scale Generalized Symmetric Eigenvalue Problem and SVD (RTR, Absil, Baker, Gallivan 2004-08)
- Blind source separation on both Orthogonal group and Oblique manifold (RTR, Absil and Gallivan 2006, Selvan et al. 2012)
- Low-rank approximate solution symmetric positive definite Lyapunov $AXM + MXA = C$ (RTR, Vandereycken and Vanderwalle 2009)
- Best low-rank approximation to a tensor (RTR, Ishteva, Absil, Van Huffel, De Lathauwer 2010)
- rotation synchronization (RTR, Absil, Boumal, 2012)
- ICA (RTR, RQN, LM-RQN, Absil, Gallivan, Huang 2013)
- rotation synchronization (RTR, RQN, LM-RQN, Absil, Gallivan, Huang 2013)
- phase retrieval (RQN, LM-RQN, Gallivan, Huang, Zhang 2014)
- graph similarity and community detection (Gallivan, Huang, Van Dooren, Zhou 2015)

Current FSU/UCL Methods

- Riemannian Steepest Descent
- Riemannian Trust Region: global quadratic convergence
- Riemannian Broyden Family : local superlinear convergence using Riemannian Wolf conditions provably and others empirically (soon to be proven)
- Riemannian Trust Region SR-1: local $(d + 1)$ -superlinear convergence
- Riemannian SR-1 line search (more restrictive in use)
- For large problems
 - Limited memory RTRSR1
 - Limited memory RBFGS
- Riemannian CG (much more work to do on analysis)

Current FSU/UCL Joint Efforts

Gallivan, Huang, Van Dooren, Zhou (current), rank-related directions, rank-related retractions, and rigorous rank updating to handle matrix rank constraints on union of parameterized manifolds

Absil, Gallivan, Huang, Yuan (current), manifolds plus inequality constraints

Absil, Gallivan, Huang, You (current), further algorithms for shape analysis and discretization effects for infinite dimensional manifolds problems

Source Separation Joint diagonalization on Orthogonal Group



Source Separation Joint diagonalization on Orthogonal Group



Blind Source Separation – Source Density Estimation



Figure : source, mixed, RBF-GS, GEKD-ICA (Xue, Wang, Yang 09)

Source Separation – Source Density Estimation

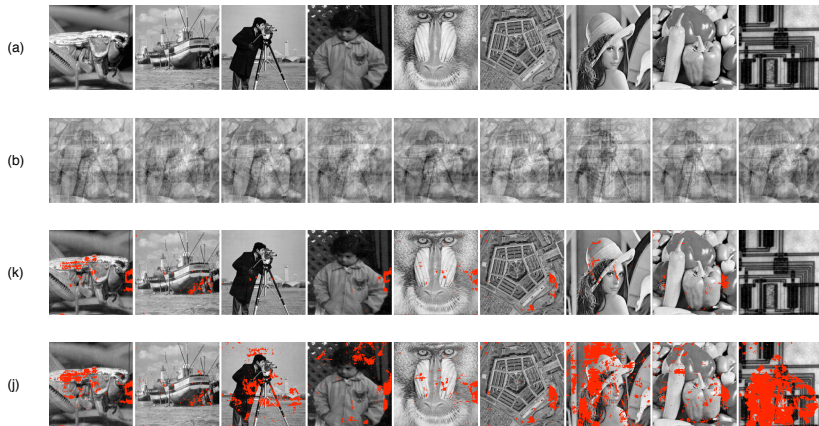


Figure : source, mixed, RBF-GS, GEKD-ICA

Future Riemannian Work

- relaxing conditions on vector transport, retractions, Wolfe conditions etc.
- constraints on union of parameterized manifolds
- manifold and inequality constraints
- discretization of infinite dimensional manifolds and the convergence/accuracy of the approximate minimizers – specific to a problem and extracting general conclusions
- Riemannian approaches to optimization problems arising out of convex relaxations of nonconvex problems

Elastic Shape Analysis

- Elastic shape analysis invariants:
 - Rescaling
 - Translation
 - Rotation
 - Reparametrization
- Younes 1998, and Younes, Michor, Shah, Mumford 2008 [11, 12]
- Square Root Velocity Function framework used (Srivastava, Klassen, Joshi, and Jermyn [8]).
- extensive analysis and application of elastic shape
- much less work on understanding efficient and robust algorithms




Figure : All are the same shape.

SRVF and Preshape Space

Preshape space, denoted \mathfrak{I}_n , removes translation and rescaling for \mathbb{L}_2 .

- A shape is represented by a function $\beta : \mathbb{D} \rightarrow \mathbb{R}^2$, where \mathbb{D} is $[0, 1]$ for open curves and unit circle \mathbb{S}^1 for closed curves.
- Square Root Velocity (SRV) function of the shape β is

$$q(t) = \begin{cases} \frac{\dot{\beta}(t)}{\sqrt{\|\dot{\beta}(t)\|_2}} & \text{if } \|\dot{\beta}(t)\|_2 \neq 0; \\ 0 & \text{if } \|\dot{\beta}(t)\|_2 = 0. \end{cases}$$

- Translation is removed by derivative.
- Rescaling is removed by forcing the length of the curve to be 1, i.e., $\int_0^1 \|\dot{\beta}(t)\|_2 dt = \int_0^1 \|q(t)\|_2^2 dt = 1$.
- Preshape spaces (closure condition added for closed curves)

$$\mathfrak{I}_n^o = \{q : [0, 1] \rightarrow \mathbb{R}^n \mid \int_0^1 \|q(t)\|_2^2 dt = 1\}$$

$$\mathfrak{I}_n^c = \{q : \mathbb{S}^1 \rightarrow \mathbb{R}^n \mid \int_{\mathbb{S}^1} \|q(t)\|_2^2 dt = 1, \int_{\mathbb{S}^1} q(t) \|q(t)\|_2 dt = 0\}$$

Shape Space

Shape space removes rotation and reparameterization. Inherits metric from \mathbb{L}_2

$$\begin{aligned} \text{SO}(n) &= \{O \in \mathbb{R}^{n \times n} \mid O^T O = I_n, \det(O) = 1\} \\ \text{SO}(n) \times \mathfrak{L}_n &\rightarrow \mathfrak{L}_n : (O, q) \rightarrow Oq \end{aligned}$$

$$\begin{aligned} \Gamma &= \{\gamma : \mathbb{D} \rightarrow \mathbb{D} \mid \gamma \text{ is a diffeomorphism.}\} \\ \mathfrak{L}_n \times \Gamma &\rightarrow \mathfrak{L}_n : (q, \gamma) \rightarrow (q \circ \gamma) \sqrt{\dot{\gamma}} \end{aligned}$$

$$[q] = \{(O, (q, \gamma)) \mid O \in \text{SO}(n), \gamma \in \Gamma\}$$

$$\mathfrak{L}_n = \mathfrak{L}_n / \text{SO}(n) \times \Gamma = \{[q] \mid q \in \mathfrak{L}_n\}.$$

both are isometric group actions. Note: diffeomorphism is $\gamma, \gamma^{-1} \in \mathcal{C}^\infty$

Best Rotation and Reparameterization

An important and basic task is to find the best rotation O_* and reparameterization γ_* .

- Geodesic
- Mean
- Other statistical analysis

Best Rotation and Reparameterization

$$(O_*, \gamma_*) = \operatorname{argmin}_{(O, \gamma) \in \operatorname{SO}(n) \times \Gamma} \operatorname{dist}_{l_n}(q_1, O\sqrt{\dot{\gamma}}q_2 \circ \gamma).$$




Figure : Align representation of $[q_2]$ with q_1 .

Cost Functions

- Minimization problem

$$\min_{O \in \text{SO}(n), \gamma \in \Gamma} \text{dist}_{I_n}(Oq_1, (q_2, \gamma)).$$

- Open curve

$$d_{I_n}(Oq_1, (q_2 \circ \gamma)\sqrt{\dot{\gamma}}) = \cos^{-1} \langle Oq_1, (q_2 \circ \gamma)\sqrt{\dot{\gamma}} \rangle_{\mathbb{L}^2}$$

$$H^o(O, \gamma(t)) = \int_0^1 \|Oq_1(t) - (q_2 \circ \gamma(t))\sqrt{\dot{\gamma}(t)}\|_2^2 dt \quad \text{same extreme points}$$

- Closed curve
 - Closed form of preshape space distance is unknown.
 - Preshape (Extrinsic) distance is used.


$$H^c(O, \gamma) = \int_{\mathbb{S}^1} \|Oq_1(t) - (q_2 \circ \gamma(t))\sqrt{\dot{\gamma}(t)}\|_2^2 dt$$

Coordinate Relaxation

Optimize rotation and reparameterization alternately.

- Open curves
 - Rotation: Procrustes problem solved using SVD
 - Reparameterization: Dynamic programming (DP) with slope constraints
- Closed curves
 - Choose a point on the closed curve and break it into an open curve
 - Apply coordinate relaxation method of open curves
 - Compare results for a sufficiently large number of break points

Two Shapes



Coordinate Relaxation Method

One iteration, denoted CR1, is used in [8].

- Complexity is $O(N^3)$, where N is the number of points in the curves.
- Note rotation and the correspondence of portions of the structures.
- Does iterating more improve results?




Figure : Results given by CR1


Representations and Implementation Difficulties

Representation Approach 1

- q_1 and q_2 are represented by points.
- Evaluation of (q_2, γ) over iterations on q
- $q_2^{(k+1)} = (q_2^{(k)}, \gamma^{(k)})$ computed on each iteration by evaluation of interpolating function of $q_2^{(k)}$.
- New interpolating function for $q_2^{(k+1)} \rightarrow$ Shape of q_2 changes.



!



Representations and Implementation Difficulties

- Representation Approach 2
 - q_1 is represented by points and q_2 is represented by an fixed interpolating curve.
- Difficulty: Lack of computational associativity may not reduce the cost function in practice
- Cost function evaluated in DP uses points $(q_2, \gamma^{(k)})$, and evaluates $((q_2, \gamma^{(k)}), \tilde{\gamma}^{(k+1)})$.
- Next q iterate is obtained using $(q_2, \gamma^{(k)} \circ \tilde{\gamma}^{(k+1)}) = (q_2, \gamma^{(k+1)})$ since fixed interpolation function for q_2
- Cost function values for the two forms of applying $\gamma^{(k+1)}$ can differ

iteration (k)	1	2	3
H^c iterate form	0.390583	0.378312	0.390114
H^c in DP	0.285534	0.248016	0.241679

Table : Computed cost function values. Difference continues growing with k .

Riemannian Approach

- Optimizing H is a Riemannian optimization problem on $\text{SO}(n) \times \Gamma$.
- Many Riemannian optimization algorithms have been systematically analyzed recently.
 - Riemannian trust-region Newton method (RTR-Newton) [2]
 - Riemannian Broyden family method including BFGS method and its limited-memory version (RBroyden family, RBFGS, LRBFGS) [7, 4, 6]
 - Riemannian trust-region symmetric rank-one update method and its limited-memory version (RTR-SR1, LRTR-SR1) [4, 5]
 - Riemannian Newton method (RNewton) [1]
- See W. Huang's thesis, Optimization algorithms on Riemannian manifolds with applications, FSU, Math Dept. [4] for details on analysis, applications and library design

Cost Function of Closed Curves

- Γ^c is represented by its covering space, i.e., $\tilde{\Gamma} \times \mathbb{R}$ where

$$\tilde{\Gamma} = \{\gamma : [0, 2\pi] \rightarrow [0, 2\pi] \mid \gamma \text{ is diffeomorphism}\}.$$

and the $\tilde{\Gamma} \times \mathbb{R}$ group action on q is defined by

$$(q, (\gamma, m)) = (q(\gamma + m \bmod 2\pi))\sqrt{\dot{\gamma}}, \quad (\gamma, m) \in \tilde{\Gamma} \times \mathbb{R}.$$

- The cost function on the Riemannian manifold $\text{SO}(n) \times \mathbb{R} \times \tilde{\Gamma}$ is

$$H^c(O, m, \gamma) = \int_0^{2\pi} \|Oq_1(t) - (q_2(\gamma(t) + m \bmod 2\pi))\sqrt{\dot{\gamma}(t)}\|_2^2 dt$$

where $\gamma(0) = 0$, $\int_0^{2\pi} \dot{\gamma}(t) dt = 2\pi$, $\dot{\gamma} > 0$.

2-norm Sphere

- Optimization on the manifold $\tilde{\Gamma}$ directly has some difficulties, e.g., step limits due to limited domains of the exponential map
 $Exp_{\gamma}(v) = \gamma + v$
- $\tilde{\Gamma}$ can be replaced with the 2-norm sphere
- Replace the term $\sqrt{\dot{\gamma}(t)}$ in H^c by a function ℓ .
- $\ell > 0$ and $\ell \in \mathbb{S}_{\mathbb{L}_2}$, where $\mathbb{S}_{\mathbb{L}_2} = \{\ell \in C^0 \mid \int_0^{2\pi} \ell^2(t) dt = 2\pi\}$.
- A constrained optimization is obtained

$$\min_{O \in \text{SO}(n), m \in \mathbb{R}, \ell \in \mathbb{S}_{\mathbb{L}_2}, \ell > 0} \int_0^{2\pi} \|Oq_1(t) - q_2(\int_0^t \ell^2(s) ds + m \pmod{2\pi})\ell(t)\|_2^2 dt.$$

4-norm Sphere

To avoid the constrained optimization, 4-norm sphere is used instead.

- 4-norm sphere
- Replace the term $\sqrt{\dot{\gamma}(t)}$ in H^c by a function ℓ^2 .
- $\ell \in \mathbb{S}_{\mathbb{L}_4}$, where $\mathbb{S}_{\mathbb{L}_4} = \{\ell \in C^0 \mid \int_0^{2\pi} \ell^4(t) dt = 2\pi\}$.
- A unconstrained optimization is obtained

$$\min_{O \in \text{SO}(n), m \in \mathbb{R}, \ell \in \mathbb{S}_{\mathbb{L}_4}} L(O, m, \ell)$$

where

$$L(O, m, \ell) = \int_0^{2\pi} \|Oq_1(t) - q_2(\int_0^t \ell^4(s) ds + m \bmod 2\pi)\ell^2(t)\|_2^2 dt.$$

Barrier Function

- A barrier function can be added to avoid the slope of γ being zero or going to ∞ :

$$B(\gamma) = \int_0^{2\pi} \left(\dot{\gamma}(t) + \frac{1}{\dot{\gamma}(t)} \right) \sqrt{1 + \dot{\gamma}^2(t)} dt = \int_0^{2\pi} \left(\ell^4(t) + \frac{1}{\ell^4(t)} \right) \sqrt{1 + \ell^8(t)} dt$$




which satisfies the symmetric property, i.e., $B(\gamma) = B(\gamma^{-1})$.

- The user can control the approach to a slope of 0 or ∞ by adding $\omega B(\gamma)$ to the cost function.



Riemmanian Algorithm

- q_1 is represented by points and q_2 is represented by an interpolating curve.
- Multiple values of m are used based on the variation of angle along the curve.
- Procrustes and DP on a coarse grid give initial ℓ_0 and O_0 for each m .
- Improvements
 - Keep the shape of q_2 constant
 - Avoid the problem with computational associativity of group action
 - Computational complexity reduces


Example




Known γ_T : rotation and γ off







LRBFGS matching curves, f:0.083826



CR1 matching curves, f:0.53435




Known γ_T : rotation and γ off significantly



Data Sets

Flavia leaf dataset [10]


- 1907 images of leaves
- 32 species



- Boundary curves: BWBOUNDARIES function in Matlab
- 100 points in \mathbb{R}^2 used for each boundary

MPEG-7 dataset [9]

- 1400 binary images
- 70 clusters



Representative of Riemannian Algorithm

- Five Riemannian methods are tested.
- 1000 pairs of shape in each data set are used.
- Based on the following table, LRBFGRS is chosen to be the representative one.

		RBFGRS	LRBFGRS	RTR-SR1	LRTR-SR1	RSD
Flavia dataset	L_{ave}	0.1727	0.1836	0.1772	0.1958	0.2079
	t_{ave}	0.4113	0.1525	0.4585	0.2052	0.2218
MPEG-7 dataset	L_{ave}	0.3639	0.3919	0.3735	0.4407	0.4798
	t_{ave}	1.2823	0.4370	1.3352	0.5572	0.7537

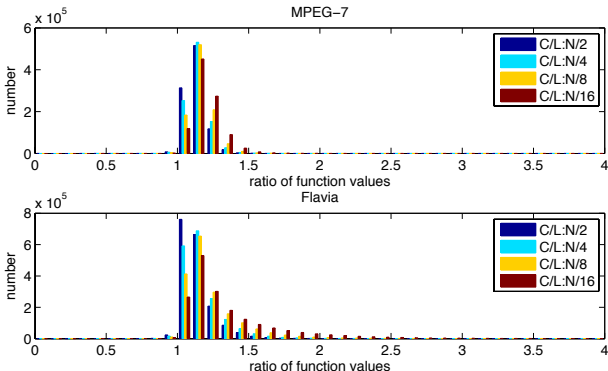
Table : Comparison of Riemannian Methods for representative sets from the Flavia and MPEG-7 datasets: average time per pair (t_{ave}) in seconds and average cost function per pair (L_{ave}).

Comparisons of LRBFGRS and CR1

Test Environment and Tests Performed



- Environment
 - All codes written in C and compiled with gcc
 - Performs on Florida State University HPC system using Quad-Core 2356 2.3 GHz Opteron [3]
- Experiments
 - Compute all pairwise distances in the Flavia and MPEG-7 respectively
 - For CR1 method, the results of the breaking points chosen to be every 2, 4, 8, 16 point are reported.

Cost Function Ratios



- Percent of Flavia pairs reduced 99.2%, 99.4%, 99.6% and 99.8% for $N/i, i = 2, 4, 8, 16$
- Percent of MPEG-7 pairs reduced 98.5%, 99.0%, 99.3% and 99.6% for $N/i, i = 2, 4, 8, 16$

Computational Time Ratios



- LRBFGS computation time adjusts based on the complexity of shape based on number of m points.
- CR1 is essentially constant due to simple choice of number of break points.
- LRBFGS generically faster even with same number of initial points.

One Nearest Neighbor Results

- The quality of the extrinsic distance computations is assessed by the one nearest neighbor (1NN) metric
- The 1NN metric, μ , computes the percentage of points whose nearest neighbor are in the same cluster, i.e.,

$$\mu = \frac{1}{n} \sum_{i=1}^n C(i), \quad C(i) = \begin{cases} 1 & \text{if point } i \text{ and its nearest neighbor} \\ & \text{are in the same cluster;} \\ 0 & \text{otherwise.} \end{cases}$$

One Nearest Neighbor Results

		LRBFGS	CR1			
			$N/16$	$N/8$	$N/4$	$N/2$
Flavia	ave. time (sec.)	0.37201	0.59379	1.1026	2.1203	4.2404
	1NN of 32 species	87%	76%	79%	81%	85%
MPEG-7	ave. time (sec.)	0.74442	0.59272	1.1006	2.1164	4.2327
	1NN of 70 clusters	98%	92%	95%	96%	97%

Table : The average computation time and 1NN of LRBFGS and CR1 with break points chosen to be every 2, 4, 8 and 16 points.

Conclusion and Future Shape Work

- Conclusion
 - CR with multiple iterations unreliable; composition unreliable
 - CR1 may not be able to find an accurate solution
 - Riemannian approach is faster, better results, and more robust for more complicated shapes than CR1
- Future work
 - Intrinsic optimization for closed curves
 - Analysis of effects of discretization on accuracy
 - Test the influence of the accuracy of distance in other shape analyses, e.g., geodesic, means
 - Combination with more robust global reparameterization optimization of Klassen et al.
 - optimization on shapes in applications

References I



P.-A. Absil, R. Mahony, and R. Sepulchre.
Optimization algorithms on matrix manifolds.
Princeton University Press, Princeton, NJ, 2008.



C. G. Baker.
Riemannian manifold trust-region methods with applications to eigenproblems.
PhD thesis, Florida State University, 2008.



Florida State University Research Computing Center.
FSU high performance computing system.



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



W. Huang, P.-A. Absil, and K. A. Gallivan.
A Riemannian symmetric rank-one trust-region method.
Mathematical Programming, 2014.

References II



W. Huang, K. A. Gallivan, and P.-A. Absil.

A Broyden class of quasi-Newton methods for Riemannian optimization.

Submitted for publication, 2014.



W. Ring and B. Wirth.

Optimization methods on Riemannian manifolds and their application to shape space.

SIAM Journal on Optimization, 22(2):596–627, January 2012.



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.



Temple University.

Shape similarity research project.

References III



S. G. Wu, F. S. Bao, E. Y. Xu, Y.-X. Wang, Y.-F. Chang, and Q.-L. Xiang.

A leaf recognition algorithm for plant classification using probabilistic neural network.

2007 IEEE International Symposium on Signal Processing and Information Technology, pages 11–16, 2007.



L. Younes.

Computable elastic distances between shapes.

SIAM Journal on Applied Mathematics, 58(2):565–586, April 1998.



L. Younes, P. Michor, J. Shah, and D. Mumford.

A metric on shape space with explicit geodesics.

Rendiconti Lincei - Matematica e Applicazioni, 9(1):25–57, 2008.