On the Smallest Enclosing Riemannian Balls

— On Approximating the Riemannian 1-Center —

http://www.sonycsl.co.jp/person/nielsen/infogeo/RiemannMinimax/

Marc Arnaudon[1] Frank Nielsen[2]

[1]Université de Bordeaux, France
[2]École Polytechnique & Sony CSL
e-mail: Frank.Nielsen@acm.org

Computational Geometry 46(1): 93-104 (2013)
arXiv 1101.4718

Introduction: Euclidean Smallest Enclosing Balls

Given d-dimensional $\mathcal{P} = \{p_1, ..., p_n\}$, find the "smallest" (with respect to the volume $\equiv$ radius $\equiv$ inclusion) ball $B = \mathrm{Ball}(c, r)$ fully covering $\mathcal{P}$:

$$c^* = \min_{c \in \mathbb{R}^d} \max_{i=1}^{n} \|c - p_i\|.$$

- unique Euclidean **circumcenter** c^*, SEB [19].
- optimization problem *non-differentiable* [10]
  c^* lie on the *farthest* Voronoi diagram

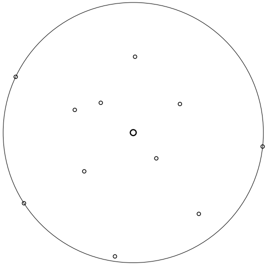

Euclidean smallest enclosing balls (SEBs)

- **1857**: $d = 2$, Smallest Enclosing Ball? of $P = \{p_1, ..., p_n\}$ (Sylvester [16])
- Randomized expected **linear** time algorithm [19, 5] in **fixed dimension** (but hidden constant *exponential* in d)
- **Core-set** [3] approximation: $(1+\epsilon)$-approximation in $O(\frac{dn}{\epsilon^2})$-time in arbitrary dimension, $O(\frac{dn}{\epsilon} + \frac{1}{\epsilon^{4.5}} \log \frac{1}{\epsilon})$ [7]
- Many other algorithms and heuristics [14, 9, 17], etc.

SEB also known as Minimum Enclosing Ball (MEB), minimax center, 1-center, bounding (hyper)sphere, etc.

$\rightarrow$ Applications in computer graphics (collision detection with ball cover proxies [15]), in machine learning (Core Vector Machines [18]), etc.

Optimization and core-sets [3]

Let $c(\mathcal{P})$ denote the circumcenter of the SEB and $r(\mathcal{P})$ its radius

Given $\epsilon > 0$, ϵ-**core-set** $\mathcal{C} \subset \mathcal{P}$, such that

$$\boxed{\mathcal{P} \subseteq \mathrm{Ball}(c(\mathcal{C}), (1+\epsilon)r(\mathcal{C}))}$$

$\Leftrightarrow$ Expanding $\mathrm{SEB}(\mathcal{C})$ by $1+\epsilon$ fully covers $\mathcal{P}$

Core-set of optimal size $\lceil \frac{1}{\epsilon} \rceil$, **independent of the dimension** d, and n. Note that **combinatorial basis** for SEB is from 2 to $d+1$ [19].

$\rightarrow$ Core-sets find many applications for problems in large-dimensions.

Euclidean SEBs from core-sets [2]

Bădoiu-Clarkson algorithm based on core-sets [2, 3]:

<u>`BCA`</u>:

- Initialize the center $c_1 \in \mathcal{P} = \{p_1, ..., p_n\}$, and
- Iteratively update the current center using the rule

$$\boxed{c_{i+1} \leftarrow c_i + \frac{f_i - c_i}{i+1}}$$

where f_i denotes the *farthest point* of $\mathcal{P}$ to c_i:

$$f_i = p_s, \quad s = \arg\max_{j=1}^{n} \|c_i - p_j\|$$

$\Rightarrow$ gradient-descent method
$\Rightarrow$ $(1+\epsilon)$-approximation after $\lceil \frac{1}{\epsilon^2} \rceil$ iterations: $O(\frac{dn}{\epsilon^2})$ time
$\Rightarrow$ Core-set: $f_1, ..., f_l$ with $l = \lceil \frac{1}{\epsilon^2} \rceil$

Euclidean SEBs from core-sets: Rewriting with $\#$

$a\#_t b$: point $(1-t)a + tb = a + t(b-a)$ on the line segment $[ab]$.
$D(x,y) = \|x-y\|^2$, $D(x,P) = \min_{y\in\mathcal{P}} D(x,y)$

Algorithm 1: $\mathtt{BCA}(\mathcal{P}, l)$.

$c_1 \leftarrow$ choose randomly a point in $\mathcal{P}$;

for $i = 2$ **to** $l-1$ **do**

```
// farthest point from c_i
```

$s_i \leftarrow \arg\max_{j=1}^n D(c_i, p_j)$;

```
// update the center: walk on the segment [c_i, p_{s_i}]
```

$\boxed{c_{i+1} \leftarrow c_i \#_{\frac{1}{i+1}} p_{s_i}}$;

end

```
// Return the SEB approximation
```

return $\mathrm{Ball}(c_l, r_l^2 = D(c_l, \mathcal{P}))$;

$\Rightarrow$ $(1+\epsilon)$-approximation after $l = \lceil \frac{1}{\epsilon^2} \rceil$ iterations.

Bregman divergences (incl. squared Euclidean distance)

SEB extended to Bregman divergences $B_F(\cdot : \cdot)$ [13]

$$B_F(c : x) = F(c) - F(x) - \langle c - x, \nabla F(x) \rangle,$$
$$B_F(c : X) = \min_{x \in X} B_F(c : x)$$

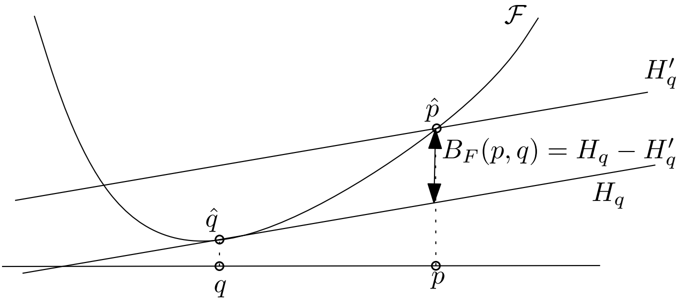

$\Rightarrow$ Bregman divergence = remainder of a first order Taylor expansion.

Smallest enclosing Bregman ball [13]

$F^* =$ *convex conjugate* of F with $(\nabla F)^{-1} = \nabla F^*$

Algorithm 2: $\mathtt{MBC}(\mathcal{P}, l)$.

```
// Create the gradient point set (η-coordinates)
```
$\mathcal{P}' \leftarrow \{\nabla F(p) : p \in \mathcal{P}\}$;

$g \leftarrow \mathrm{BCA}(\mathcal{P}', l)$;

return $\mathrm{Ball}(c_l = \nabla F^{-1}(c(g)), r_l = B_F(c_l : \mathcal{P}))$;

Guaranteed approximation algorithm with approximation factor depending on $\frac{1}{\min_{x\in\mathcal{X}} \|\nabla^2 F(x)\|}$, ... but **poor** in practice

$$\forall s,\ S_F(x; \nabla F^{-1}(c(g))) \leq \frac{(1+\epsilon)^2 r'^*}{\min_{x\in\mathcal{X}} \|\nabla^2 F(x)\|}$$

with $S_F(c; x) = B_F(c : x) + B_F(x : c)$

Smallest enclosing Bregman ball [13]

A better approximation in **practice**...

Algorithm 3: $\texttt{BBCA}(\mathcal{P}, l)$.

$c_1 \leftarrow$ choose randomly a point in $\mathcal{P}$;

for $i = 2$ **to** $l - 1$ **do**

```
// farthest point from c_i wrt.  B_F
```

$s_i \leftarrow \arg\max_{j=1}^{n} B_F(c_i : p_j)$;

```
// update the center: walk on the η-segment [c_i, p_{s_i}]_η
```

$\boxed{c_{i+1} \leftarrow \nabla F^{-1}(\nabla F(c_i) \#_{\frac{1}{i+1}} \nabla F(p_{s_i}))}$;

end

```
// Return the SEBB approximation
```

return $\mathrm{Ball}(c_l, r_l = B_F(c_l : X))$;

θ-, η-geodesic segments in dually flat geometry.

Basics of Riemannian geometry

- (M, g): Riemannian manifold
- $\langle\cdot,\cdot\rangle$, Riemannian *metric tensor* g: definite positive bilinear form on each tangent space $T_x M$ (depends smoothly on x)
- $\|\cdot\|_x$: $\|u\| = \langle u, u\rangle^{1/2}$: Associated norm in $T_x M$
- $\rho(x, y)$: *metric distance* between two points on the manifold M (length space)

$$\rho(x, y) = \inf \left\{ \int_0^1 \|\dot\varphi(t)\| \,\mathrm{d}t, \quad \varphi \in C^1([0,1], M), \quad \varphi(0) = x, \quad \varphi(1) = y \right\}$$

Parallel transport wrt. Levi-Civita metric connection ∇: $\nabla g = 0$.

Basics of Riemannian geometry: Exponential map

- Local map from the *tangent space* T_xM to the *manifold* defined with geodesics (wrt ∇).

$$\forall x \in M, D(x) \subset T_xM : D(x) = \{v \in T_xM : \gamma_v(1) \text{ is defined}\}$$

with γ_v maximal (i.e., largest domain) geodesic with $\gamma_v(0) = x$ and $\gamma_v'(0) = v$.

- **Exponential map**:

$$\begin{aligned} \exp_x(\cdot) \quad &: \quad D(x) \subseteq T_xM \to M \\ &\quad\ \exp_x(v) = \gamma_v(1) \end{aligned}$$

D is *star-shaped*.

Basics of Riemannian geometry: Geodesics

- *Geodesic*: smooth path which locally minimizes the distance between two points. (In general such a curve does not minimize it globally.)
- Given a vector $v \in T_x M$ with base point x, there is a unique geodesic started at x with speed v at time 0: $t \mapsto \exp_x(tv)$ or $t \mapsto \gamma_t(v)$.
- Geodesic on $[a, b]$ is *minimal* if its length is less or equal to others. For *complete* M (i.e., $\exp_x(v)$), taking $x, y \in M$, there exists a *minimal* geodesic from x to y in time 1. $\gamma_\cdot(x, y) : [0, 1] \to M$, $t \mapsto \gamma_t(x, y)$ with the conditions $\gamma_0(x, y) = x$ and $\gamma_1(x, y) = y$.
- $U \subseteq M$ is *convex* if for any $x, y \in U$, there exists a unique minimal geodesic $\gamma_\cdot(x, y)$ in M from x to y. Geodesic *fully lies* in U and depends smoothly on x, y, t.

Basics of Riemannian geometry: Geodesics

- Geodesic $\gamma(x, y)$: locally minimizing curves linking x to y
- Speed vector $\gamma'(t)$ **parallel** along γ:

$$\frac{D\gamma'(t)}{\mathrm{d}t} = \nabla_{\gamma'(t)}\gamma'(t) = 0$$

- When manifold M embedded in $\mathbb{R}^d$, acceleration is normal to tangent plane:

$$\gamma''(t) \perp T_{\gamma(t)}M$$

- $\|\gamma'(t)\| = c$, a constant (say, unit).

$\Rightarrow$ Parameterization of curves with constant speed...

Basics of Riemannian geometry: Geodesics

Constant speed geodesic $\gamma(t)$ so that $\gamma(0) = x$ and $\gamma(\rho(x,y)) = y$ (constant speed 1, the unit of length).

$$\boxed{x \#_t y = m = \gamma(t) : \rho(x,m) = t \times \rho(x,y)}$$

For example, in the Euclidean space:

$$x \#_t y = (1-t)x + ty = x + t(y-x) = m$$

$$\rho_E(x,m) = \|t(y-x)\| = t\|y-x\| = t \times \rho(x,y), t \in [0,1]$$

$\Rightarrow$ m interpreted as a **mean** (barycenter) between x and y.

Basics of Riemannian geometry: Injectivity radius

Diffeomorphism from the tangent space to the manifold

- *Injectivity radius* $\mathrm{inj}(M)$: largest $r > 0$ such that **for all** $x \in M$, the map $\exp_x(\cdot)$ restricted to the open ball in $T_x M$ with radius r is an embedding.
- *Global injectivity radius*: infimum of the injectivity radius over all points of the manifold.

Basics of Riemannian geometry: Sectional curvature

Given $x \in M$, u, v two non collinear vectors in $T_x M$, the *sectional curvature* $\mathrm{Sect}(u, v) = K$ is a number which gives information on how the geodesics issued from x behave near x.
More precisely, the image by $\exp_x(\cdot)$ of the circle centered at 0 of radius $r > 0$ in $\mathrm{Span}(u, v)$ has length

$$2\pi S_K(r) + o(r^3) \quad \text{as} \quad r \to 0$$

with

$$S_K(r) = \begin{cases} \frac{\sin(\sqrt{K}r)}{\sqrt{K}} & \text{if} \quad K > 0, \\ r & \text{if} \quad K = 0, \\ \frac{\sinh(\sqrt{-K}r)}{\sqrt{-K}} & \text{if} \quad K < 0. \end{cases}$$

positive, zero or negative curvatures...

Basics of Riemannian geometry: **Alexandrov's theorem**

Given an *upper bound* α^2 for sectional curvatures, compare **geodesic triangles** by *Alexandrov* theorem:
Let $x_1, x_2, x_3 \in M$ satisfy $x_1 \neq x_2$, $x_1 \neq x_3$ and

$$\rho(x_1, x_2) + \rho(x_2, x_3) + \rho(x_3, x_1) < 2 \min \left\{ \mathrm{inj}(M), \frac{\pi}{\alpha} \right\}$$

where $\alpha > 0$ is such that α^2 is an upper bound of sectional curvatures. Let the minimizing geodesic from x_1 to x_2 and the minimizing geodesic from x_1 to x_3 make an angle θ at x_1. Denoting by $S^2_{\alpha^2}$ the 2-**dimensional sphere** of constant curvature α^2 (hence of radius $1/\alpha$) and $\tilde{\rho}$ the distance in $S^2_{\alpha^2}$, we consider points $\tilde{x}_1, \tilde{x}_2, \tilde{x}_3 \in S^2_{\alpha^2}$ such that $\rho(x_1, x_2) = \tilde{\rho}(\tilde{x}_1, \tilde{x}_2)$, $\rho(x_1, x_3) = \tilde{\rho}(\tilde{x}_1, \tilde{x}_3)$. Assume that the minimizing geodesic from $\tilde{x}_1$ to $\tilde{x}_2$ and the minimizing geodesic from $\tilde{x}_1$ to $\tilde{x}_3$ also make an angle θ at $\tilde{x}_1$.
Then we have: $\boxed{\rho(x_2, x_3) \geq \tilde{\rho}(\tilde{x}_2, \tilde{x}_3)}$.

Basics of Riemannian geometry: **Topogonov's theorem**

Assume $\beta > 0$ is such that $-\beta^2$ is a lower bound for sectional curvatures in M. Let $x_1, x_2, x_3 \in M$ satisfy $x_1 \neq x_2$, $x_1 \neq x_3$. Let the minimizing geodesic from x_1 to x_2 and the minimizing geodesic from x_1 to x_3 make an angle θ at x_1. Denoting by $H^2_{-\beta^2}$ the **hyperbolic** 2-**dimensional space** of constant curvature $-\beta^2$ and $\tilde{\rho}$ the distance in $H^2_{-\beta^2}$, we consider points $\tilde{x}_1, \tilde{x}_2, \tilde{x}_3 \in H^2_{-\beta^2}$ such that $\rho(x_1, x_2) = \tilde{\rho}(\tilde{x}_1, \tilde{x}_2)$, $\rho(x_1, x_3) = \tilde{\rho}(\tilde{x}_1, \tilde{x}_3)$. Assume that the minimizing geodesic from $\tilde{x}_1$ to $\tilde{x}_2$ and the minimizing geodesic from $\tilde{x}_1$ to $\tilde{x}_3$ also make an angle θ at $\tilde{x}_1$.

Then we have: $\boxed{\rho(x_2, x_3) \leq \tilde{\rho}(\tilde{x}_2, \tilde{x}_3)}$.

Basics of Riemannian geometry: **First law of cosines**

In spherical/hyperbolic geometries:

- If $\theta_1, \theta_2, \theta_3$ are the angles of a triangle in $S^2_{\alpha^2}$ and l_1, l_2, l_3 are the lengths of the opposite sides, then

$$\boxed{\cos\theta_3 = \frac{\cos(\alpha l_3) - \cos(\alpha l_1)\cos(\alpha l_2)}{\sin(\alpha l_1)\sin(\alpha l_2)}}$$

- If $\theta_1, \theta_2, \theta_3$ are the angles of a triangle in $H^2_{-\beta^2}$ and l_1, l_2, l_3 are the lengths of the opposite sides, then

$$\boxed{\cos\theta_3 = \frac{\cosh(\beta l_1)\cosh(\beta l_2) - \cosh(\beta l_3)}{\sinh(\beta l_1)\sinh(\beta l_2)}}$$

Now ready for the "Smallest enclosing Riemannian ball"

(M, g): complete Riemannian manifold
ν: *probability measure* on M
$\rho(x, y)$: Riemannian metric distance

Assume the measure support $\mathrm{supp}(\nu) \subseteq$ in a **geodesic ball** $B(o, R)$.

$f : M \to \mathbb{R}$: measurable function

$$\|f\|_{L^\infty(\nu)} = \inf \{a > 0, \ \ \nu(\{y \in M, \ |f(y)| > a\}) = 0\}.$$

$\alpha > 0$ such that α^2 upper bounds the sectional curvatures in M.

$$\boxed{R_\alpha = \frac{1}{2} \min \left\{ \mathrm{inj}(M), \frac{\pi}{\alpha} \right\}}$$

$\mathrm{inj}(M)$: injectivity radius

Riemannian SEB: Existence and uniqueness [1]

Assume

$$R < R_\alpha$$

Consider **farthest point** map:

$$\begin{aligned} H \quad : \quad & M \to [0, \infty] \\ & x \mapsto \|\rho(x, \cdot)\|_{L^\infty(\nu)} \end{aligned} \tag{1}$$

$c \in B(o, R)$.
$\to c \subset \mathrm{CH}(\mathrm{supp}(\nu))$ [1] (convex hull)

$\Rightarrow$ center: notion of *centrality* of the measure
$\Rightarrow$ point set: discrete measure, center $\to$ circumcenter

Example of Riemannian manifold: SPD space

Space of **Symmetric Positive Definite** (SPD) matrices with

- Riemannian distance:

$$\boxed{\rho(P, Q) = \| \log(P^{-1}Q)\|_F = \sqrt{\sum_{i=1}^{d} \log^2 \lambda_i}}$$

  where λ_i are the eigenvalues of matrix $P^{-1}Q$.
- Non-compact Riemannian symmetric space of non-positive curvature (aka. Cartan-Hadamard manifold).
- *Any* measure ν with *bounded support* satisfies $R < R_\alpha$ (choose $\alpha > 0$).

$\Rightarrow$ Minimizer c of farthest point map H exists and is unique: 1-center or minimax center of ν.

Generalizing BCA to Riemannian manifolds

GeoA:

- Initialize the center with $c_1 \in \mathcal{P}$, and
- Iteratively update the current minimax center as

$$\boxed{c_{i+1} = \mathrm{Geodesic}\left(c_i, f_i, \frac{1}{i+1}\right)}$$

where f_i denotes the farthest point of $\mathcal{P}$ to c_i, and $\mathrm{Geodesic}(p, q, t)$ denotes the intermediate point m on the geodesic passing through p and q such that $\rho(p, m) = t \times \rho(p, q)$.

Generalizing BCA to Riemannian manifolds

$a\#_t^M b$: point $\gamma(t)$ on the geodesic line segment $[ab]$ wrt M.

Algorithm 4: `GeoA`

$c_1 \leftarrow$ choose randomly a point in $\mathcal{P}$;

for $i = 2$ **to** l **do**

```
// farthest point from c_i
```

$s_i \leftarrow \arg\max_{j=1}^{n} \rho(c_i, p_j)$;

```
// update the center: walk on the geodesic line
   segment [c_i, p_{s_i}]
```

$c_{i+1} \leftarrow c_i \#^M_{\frac{1}{i+1}} p_{s_i}$;

end

```
// Return the SEB approximation
```

return $\mathrm{Ball}(c_l, r_l = \rho(c_l, \mathcal{P}))$;

Proof sketch

Assume $\mathrm{supp}(\nu) \subset B(o, R)$ and

$$R < R_\alpha = \frac{1}{2} \min \left\{ \mathrm{inj}(M), \frac{\pi}{\alpha} \right\}$$

with $\alpha > 0$ such that α^2 is an upper bound for the sectional curvatures in M.

Lemma

There exists $\tau > 0$ such that for all $x \in B(o, R)$,

$$\boxed{H(x) - H(c) \geq \tau \rho^2(x, c)}$$

Stochastic approximation for measures

For $x \in B(o, R)$, $t \mapsto \gamma_t(v(x, \nu))$ a unit speed geodesic from $\gamma_0(v(x, \nu)) = x$ to one point $y = \gamma_{H(x)}(v(x, \nu))$ in $\mathrm{supp}(\nu)$ which realizes the maximum of the distance from x to $\mathrm{supp}(\nu)$.

$$\boxed{v = \frac{1}{H(x)} \exp_x^{-1}(y)}$$

`RieA`:

Fix some $\delta > 0$.

- **Step 1** Choose a starting point $x_0 \in \mathrm{supp}(\nu)$ and let $k = 0$
- **Step 2** Choose a step size $t_{k+1} \in (0, \delta]$ and let $x_{k+1} = \gamma_{t_{k+1}}(v(x_k, \nu))$, then do again step 2 with $k \leftarrow k + 1$.

Convergence theorem for `RieA`

$a \wedge b$: minimum operator $a \wedge b = \min(a, b)$.

$$R_0 = \frac{R_\alpha - R}{2} \wedge \frac{R}{2}.$$

Assume $\alpha, \beta > 0$ are such that $-\beta^2$ is a lower bound and α^2 an upper bound of the sectional curvatures in M. If the step sizes $(t_k)_{k \geq 1}$ satisfy

$$\delta \leq \frac{R_0}{2} \wedge \frac{2}{\beta} \mathrm{arctanh}\left(\tanh(\beta R_0/2) \cos(\alpha R) \tan(\alpha R_0/4)\right),$$

$$\lim_{k \to \infty} t_k = 0, \qquad \sum_{k=1}^{\infty} t_k = +\infty \quad \text{and} \quad \sum_{k=1}^{\infty} t_k^2 < \infty.$$

then the sequence $(x_k)_{k \geq 1}$ generated by the algorithm satisfies

$$\boxed{\lim_{k \to \infty} \rho(x_k, c) = 0}$$

Case study I: Hyperbolic planar manifold

In Klein disk (projective model), geodesics are straight (euclidean) lines [11].

$$\boxed{\rho(p,q) = \mathrm{arccosh} \frac{1 - p^\top q}{\sqrt{(1 - p^\top p)(1 - q^\top q)}}}$$

where $\mathrm{arccosh}(x) = \log(x + \sqrt{x^2 - 1})$.
Here, we choose non-constant speed curve parameterization (not constant-speed geodesic):

$$\tilde{\gamma}_t(p,q) = (1-t)p + tq, \;\; t \in [0,1].$$

$\Rightarrow$ Implement a dichotomy on $\tilde{\gamma}_t(p,q)$ to get $\#_t$.

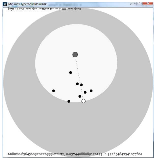

Initialization

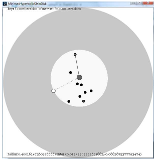

First iteration

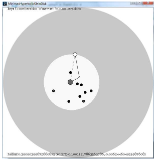

Second iteration

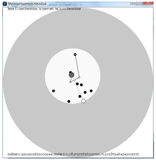

Third iteration

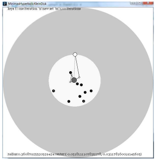

Fourth iteration

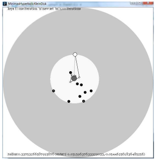

after 104 iterations

http://www.sonycsl.co.jp/person/nielsen/infogeo/RiemannMinimax/

Performance

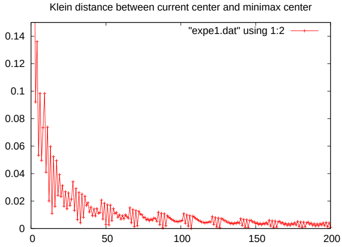

(a)

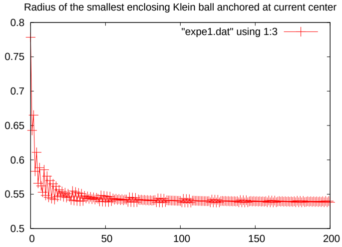

(b)

Convergence rate of the GeoA algorithm for the hyperbolic disk for the first 200 iterations. Horizontal axis: number of iterations Vertical axis: (a) the relative Klein distance between the current center and the optimal 1-center, (b) the radius of the smallest enclosing ball anchored at the current center.

Case study II: Space of SPD matrices

- $d \times d$ matrix M **Symmetric Positive Definite** (SPD) $\Leftrightarrow$ $M = M^\top$ and that for all $x \neq 0$, $x^\top M x > 0$.
- The set of $d \times d$ SPD matrices: manifold of dimension $\frac{d(d+1)}{2}$ [8]
- The geodesic linking (matrix) point P to point Q:

$$\boxed{\gamma_t(P, Q) = P^{\frac{1}{2}} \left(P^{-\frac{1}{2}} Q P^{-\frac{1}{2}} \right)^t P^{\frac{1}{2}}}$$

where the matrix function $h(M)$ is computed from the singular value decomposition $M = UDV^\top$ (with U and V unitary matrices and $D = \mathrm{diag}(\lambda_1, ..., \lambda_d)$ a diagonal matrix of eigenvalues) as $h(M) = U\mathrm{diag}(h(\lambda_1), ..., h(\lambda_d))V^\top$. For example, the square root function of a matrix is computed as $M^{\frac{1}{2}} = U\ \mathrm{diag}(\sqrt{\lambda_1}, ..., \sqrt{\lambda_d})\ V^\top$.

SPD space: Splitting the geodesic for operator $\#_t$

In this case, finding t such that

$$\|\log(P^{-1}Q)^t\|_F^2 = r\|\log P^{-1}Q\|_F^2, \qquad (2)$$

where $\|\cdot\|_F$ denotes the Fröbenius norm yields to $t = r$. Indeed, consider $\lambda_1, ..., \lambda_d$ the eigenvalues of $P^{-1}Q$, then $\rho(P, Q) = \|\log(P^{-1}Q)\|_F = \sqrt{\sum_i \log^2 \lambda_i}$ amounts to find

$$\sum_{i=1}^{d} \log^2 \lambda_i^t = t^2 \sum_{i=1}^{d} \log^2 \lambda_i = r^2 \sum_{i=1}^{d} \log^2 \lambda_i.$$

That is $t = r$.

Case study II: Performance

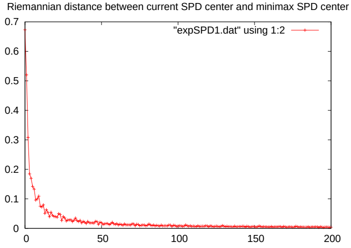

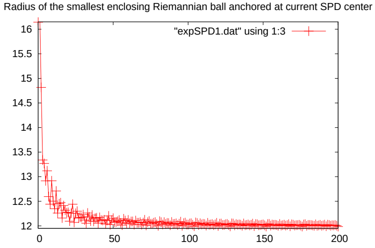

(a) (b)

Convergence rate of the `GeoA` algorithm for the SPD Riemannian manifold (dimension 5) for the first 200 iterations.
Horizontal axis: number of iterations i
Vertical axis:

- (a) the relative Riemannian distance between the current center c_i and the optimal 1-center c^* $\left(\frac{\rho(c*, c_i)}{r^*}\right)$
- (b) the radius r_i of the smallest enclosing SPD ball anchored at the current center.

Remark on SPD spaces and hyperbolic geometry

- 2D $\mathrm{SPD}(2)$ matrix space has dimension $d = 3$: A positive cone.

$$\left\{(a, b, c) : a > 0, \quad ab - c^2 > 0\right\}$$

- Can be *peeled into sheets* of dimension 2, each sheet corresponding to a constant value of the determinant of the elements [4]

$$\mathrm{SPD}(2) = \mathrm{SSPD}(2) \times \mathbb{R}^+,$$

where $\mathrm{SSPD}(2) = \{a, b, c = \sqrt{1 - ab}) : a > 0, ab - c^2 = 1\}$

- Map to $(x_0 = \frac{a+b}{2} \geq 1, x_1 = \frac{a-b}{2}, x_2 = c)$ in hyperboloid model [12], and $z = \frac{a-b+2ic}{2+a+b}$ in Poincaré disk [12].

Conclusion: Smallest Riemannian Enclosing Ball

- Generalize Euclidean 1-center algorithm of [2] to *Riemannian geometry*
- Proved the *convergence* under mild assumptions (for measures/point sets)
- Existence of *Riemannian core-sets* for optimization
- 1-center building block for k-**center clustering** [6]
- can be extended to sets of *Riemannian (geodesic) balls*

Reproducible research codes with interactive demos:

`http://www.sonycsl.co.jp/person/nielsen/infogeo/RiemannMinimax/`

Bibliographic references I

Bijan Afsari.

Riemannian L^p center of mass : existence, uniqueness, and convexity.

Proceedings of the American Mathematical Society, 139:655–674, February 2011.

Mihai Bădoiu and Kenneth L. Clarkson.

Smaller core-sets for balls.

In *SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 801–802, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.

Mihai Bădoiu and Kenneth L. Clarkson.

Optimal core-sets for balls.

Computational Geometry: Theory and Applications, 40:14–22, May 2008.

Pascal Chossat and Olivier P. Faugeras.

Hyperbolic planforms in relation to visual edges and textures perception.

PLoS Computational Biology, 5(12), 2009.

Kaspar Fischer and Bernd Gärtner.

The smallest enclosing ball of balls: combinatorial structure and algorithms.

Int. J. Comput. Geometry Appl., 14(4-5):341–378, 2004.

Teofilo F. Gonzalez.

Clustering to minimize the maximum intercluster distance.

Theoretical Computer Science, 38(0):293 – 306, 1985.

Bibliographic references II

Piyush Kumar, Joseph S. B. Mitchell, and E. Alper Yildirim.

Approximate minimum enclosing balls in high dimensions using core-sets.

ACM Journal of Experimental Algorithmics, 8, 2003.

Serge Lang.

Fundamentals of differential geometry, volume 191 of *Graduate Texts in Mathematics*.

Springer-Verlag, New York, 1999.

Thomas Larsson.

Fast and tight fitting bounding spheres.

In *Proceedings of the Annual SIGRAD Conference, Stockholm*, 2008.

Frank Nielsen and Richard Nock.

Approximating smallest enclosing balls with applications to machine learning.

Int. J. Comput. Geometry Appl., 19(5):389–414, 2009.

Frank Nielsen and Richard Nock.

Hyperbolic Voronoi diagrams made easy.

In *International Conference on Computational Science and its Applications (ICCSA)*, volume 1, pages 74–80, Los Alamitos, CA, USA, march 2010. IEEE Computer Society.

Bibliographic references III

Frank Nielsen and Richard Nock.

Visualizing hyperbolic Voronoi diagrams.

In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry*, SOCG'14, pages 90:90–90:91, New York, NY, USA, 2014. ACM.

Richard Nock and Frank Nielsen.

Fitting the smallest enclosing bregman balls.

In *16th European Conference on Machine Learning (ECML)*, pages 649–656, October 2005.

Jack Ritter.

An efficient bounding sphere.

In *Graphics gems*, pages 301–303. Academic Press Professional, Inc., 1990.

Jonas Spillmann, Markus Becker, and Matthias Teschner.

Efficient updates of bounding sphere hierarchies for geometrically deformable models.

Journal of Visual Communication and Image Representation, 18(2):101–108, 2007.

J. J. Sylvester.

A question in the geometry of situation.

Quarterly Journal of Pure and Applied Mathematics, 1:79, 1857.

Bo Tian.

Bouncing Bubble: A fast algorithm for Minimal Enclosing Ball problem.

2012.

Bibliographic references IV

Ivor W. Tsang, Andras Kocsor, and James T. Kwok.

Simpler core vector machines with enclosing balls.

In *Proceedings of the 24th International Conference on Machine Learning (ICML)*, pages 911–918, New York, NY, USA, 2007. ACM.

Emo Welzl.

Smallest enclosing disks (balls and ellipsoids).

In H. Maurer, editor, *New Results and New Trends in Computer Science*, LNCS. Springer, 1991.