Approximating Covering and Minimum Enclosing Balls in Hyperbolic Geometry

Frank Nielsen[1]   Gaëtan Hadjeres[2]

École Polytechnique [1]

Sony Computer Science Laboratories, Inc [1,2]

Conference on Geometric Science of Information

The Minimum Enclosing Ball problem

Finding the *Minimum Enclosing Ball* (or the *1-center*) of a finite point set $P = \{p_1, \ldots, p_n\}$ in the metric space $(X, d_X(.,.))$ consists in finding $c \in X$ such that

$$\boxed{c = \mathrm{argmin}_{c' \in X} \max_{p \in P} d_X(c', p)}$$

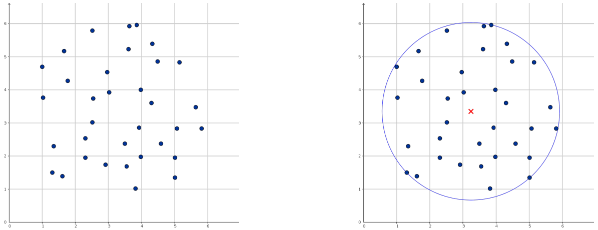

Figure : A finite point set P and its minimum enclosing ball $\mathrm{MEB}(P)$

The approximating minimum enclosing ball problem

In a euclidean setting, this problem is

- well-defined: uniqueness of the center c^* and radius R^* of the MEB
- computationally intractable in high dimensions.

We fix an $\epsilon > 0$ and focus on the *Approximate Minimum Enclosing Ball problem* of finding an ϵ-approximation $c \in X$ of $\mathrm{MEB}(P)$ such that

$$d_X(c, p) \leq (1 + \epsilon)R^* \quad \forall p \in P.$$

The approximating minimum enclosing ball problem: prior work

Approximate solution in the euclidean case are given by Badoiu and Clarkson's algorithm [Badoiu and Clarkson, 2008]:

- Initialize center $c_1 \in P$
- Repeat $\lfloor 1/\epsilon^2 \rfloor$ times the following update:

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

where $f_i \in P$ is the farthest point from c_i.

How to deal with point sets whose underlying geometry is not euclidean ?

The approximating minimum enclosing ball problem: prior work

This algorithm has been generalized to

- dually flat manifolds [Nock and Nielsen, 2005]
- Riemannian manifolds [Arnaudon and Nielsen, 2013]

Applying these results to hyperbolic geometry give the existence and uniqueness of $\mathrm{MEB}(P)$, but

- give no explicit bounds on the number of iterations
- assume that we are able to precisely cut geodesics.

The approximating minimum enclosing ball problem: our contribution

We analyze the case of point sets whose underlying geometry is hyperbolic.
Using a closed-form formula to compute geodesic α-midpoints, we obtain

- a intrinsic $(1+\epsilon)$-approximation algorithm to the approximate minimum enclosing ball problem
- a $O(1/\epsilon^2)$ convergence time guarantee
- a one-class clustering algorithm for specific subfamilies of normal distributions using their Fisher information metric

Model of d-dimensional hyperbolic geometry: The Poincaré ball model

The *Poincaré ball model* $(\mathbb{B}^d, \rho(.,.))$ consists in the open unit ball $\mathbb{B}^d = \{x \in \mathbb{R}^d : \|x\| < 1\}$ together with the hyperbolic distance

$$\rho(p, q) = \operatorname{arcosh}\left(1 + \frac{2\|p-q\|^2}{(1-\|p\|^2)(1-\|q\|^2)}\right), \quad \forall p, q \in \mathbb{B}^d.$$

This distance induces on the metric space $(\mathbb{B}^d, \rho)$ a Riemannian structure.

Geodesics in the Poincaré ball model

Shorter paths between two points (geodesics) are exactly

- straight (euclidean) lines passing through the origin
- circle arcs orthogonal to the unit sphere

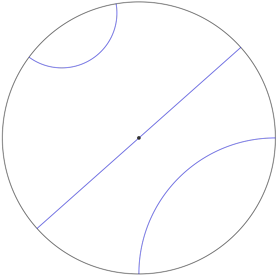

Figure : "Straight" lines in the Poincaré ball model

Circles in the Poincaré ball model

Circles in the Poincaré ball model

- look like euclidean circles
- but with different center

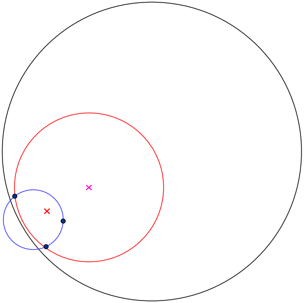

Figure : Difference between euclidean MEB (in blue) and hyperbolic MEB (in red) for the set of blue points in hyperbolic Poincaré disk (in black). The red cross is the hyperbolic center of the red circle while the pink one is its euclidean center.

Translations in the Poincaré ball model

$$T_p(x) = \frac{\left(1 - \|p\|^2\right) x + \left(\|x\|^2 + 2\langle x, p\rangle + 1\right) p}{\|p\|^2 \|x\|^2 + 2\langle x, p\rangle + 1}$$

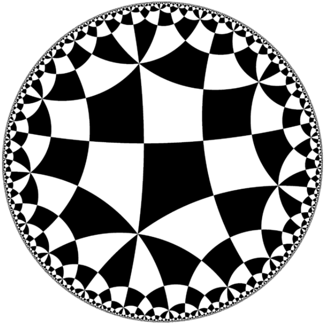

Figure : Tiling of the hyperbolic plane by squares

Closed-form formula for computing α-midpoints

A point m is the α-midpoint $p\#_\alpha q$ of two points p, q for $\alpha \in [0,1]$ if

- m belongs to the geodesic joining the two points p, q
- m verifies

$$\rho\left(p, m_\alpha\right) = \alpha \rho\left(p, q\right).$$

Closed-form formula for computing α-midpoints

A point m is the α-midpoint $p\#_\alpha q$ of two points p, q for $\alpha \in [0, 1]$ if

- m belongs to the geodesic joining the two points p, q
- m verifies

$$\rho(p, m_\alpha) = \alpha\rho(p, q).$$

For the special case $p = (0, \ldots, 0)$, $q = (x_q, 0, \ldots, 0)$, we have

$$p\#_\alpha q := (x_\alpha, 0, \ldots, 0)$$

with

$$x_\alpha = \frac{c_{\alpha,q} - 1}{c_{\alpha,q} + 1}, \quad \text{where} \quad c_{\alpha,q} := e^{\alpha\rho(p,q)} \left(= \left(\frac{1 + x_q}{1 - x_q} \right)^\alpha \right).$$

Closed-form formula for computing α-midpoints

Noting that

$$p \#_\alpha q = T_p \left(T_{-p}(p) \#_\alpha T_{-p}(q) \right) \quad \forall p, q \in \mathbb{B}^d$$

we obtain

- a closed-form formula for computing $p \#_\alpha q$
- how to compute $p \#_\alpha q$ in linear time $O(d)$
- that these transformations are exact.

$(1+\epsilon)$-approximation of an hyperbolic enclosing ball of fixed radius

For a fixed radius $r > R^*$, we can find $c \in \mathbb{B}^d$ such that

$$\rho\left(c, P\right) \leq (1+\epsilon) r \quad \forall p \in P$$

with

Algorithm 1: $(1+\epsilon)$-approximation of $\mathrm{EHB}(P, r)$

```
c_0 := p_1
t := 0
while ∃p ∈ P such that p ∉ B(c_t, (1+ε) r) do
    let p ∈ P be such a point
    α := (ρ(c_t, p) − r) / ρ(c_t, p)
    c_{t+1} := c_t #_α p
    t := t+1
end while
return c_t
```

Idea of the proof

By the hyperbolic law of cosines :

$$\mathrm{ch}\left(\rho_t\right) \geq \mathrm{ch}\left(h\right)\mathrm{ch}\left(\rho_{t+1}\right)$$

$$\mathrm{ch}\left(\rho_1\right) \geq \mathrm{ch}\left(h\right)^T \geq \mathrm{ch}\left(\epsilon r\right)^T .$$

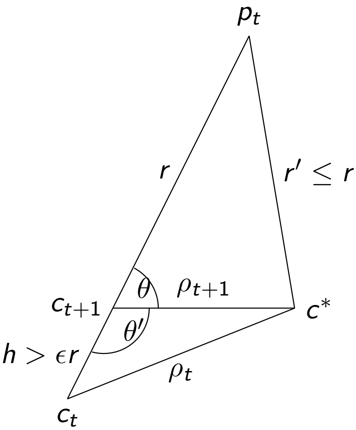

Figure : Update of c_t

$(1+\epsilon)$-approximation of an hyperbolic enclosing ball of fixed radius

The $\mathrm{EHB}(P, r)$ algorithm is a $O(1/\epsilon^2)$-time algorithm which returns

- the center of a hyperbolic enclosing ball with radius $(1+\epsilon)r$
- in less than $4/\epsilon^2$ iterations.

$(1+\epsilon)$-approximation of an hyperbolic enclosing ball of fixed radius

The $\mathrm{EHB}(P, r)$ algorithm is a $O(1/\epsilon^2)$-time algorithm which returns

- the center of a hyperbolic enclosing ball with radius $(1+\epsilon)r$
- in less than $4/\epsilon^2$ iterations.

Our error with the true MEHB center c^* verifies

$$\rho\left(c, c^*\right) \leq \operatorname{arcosh}\left(\frac{\operatorname{ch}\left((1+\epsilon)\, r\right)}{\operatorname{ch}\left(R^*\right)}\right)$$

$(1+\epsilon+\epsilon^2/4)$-approximation of $\mathrm{MEHB}(P)$

In fact, as R^* is unknown in general, the EHB algorithm returns for any r:

- an $(1+\epsilon)$-approximation of $\mathrm{EHB}(P)$ if $r \geq R^*$
- the fact that $r < R^*$ if the result obtained after more than $4/\epsilon^2$ iterations is not good enough.

$(1+\epsilon+\epsilon^2/4)$-approximation of $\mathrm{MEHB}(P)$

In fact, as R^* is unknown in general, the EHB algorithm returns for any r:

- an $(1+\epsilon)$-approximation of $\mathrm{EHB}(P)$ if $r \geq R^*$
- the fact that $r < R^*$ if the result obtained after more than $4/\epsilon^2$ iterations is not good enough.

This suggests to implement a dichotomic search in order to compute an approximation of the *minimal hyperbolic enclosing ball*. We obtain

- a $O(1+\epsilon+\epsilon^2/4)$-approximation of $\mathrm{MEHB}(P)$
- in $O\left(\frac{N}{\epsilon^2}\log\left(\frac{1}{\epsilon}\right)\right)$ iterations.

$(1+\epsilon+\epsilon^2/4)$-approximation of $\mathrm{MEHB}(P)$ algorithm

Algorithm 2: $(1+\epsilon)$-approximation of $\mathrm{MEHB}(P)$

$c := p_1$

$r_{\max} := \rho(c, P)$; $r_{\min} = \frac{r_{\max}}{2}$; $t_{\max} := +\infty$

$r := r_{\max}$;

repeat

  $c_{\text{temp}} := \text{Alg1}\left(P, r, \frac{\epsilon}{2}\right)$, interrupt if $t > t_{\max}$ in Alg1

  if call of Alg1 has been interrupted **then**

    $r_{\min} := r$

  else

    $r_{\max} := r$; $c := c_{\text{temp}}$

  end if

  $dr := \frac{r_{\max} - r_{\min}}{2}$; $r := r_{\min} + dr$;
  $t_{\max} := \frac{\log(\text{ch}(1+\epsilon/2)r) - \log(\text{ch}(r_{\min}))}{\log(\text{ch}(r\epsilon/2))}$

until $2dr < r_{\min}\frac{\epsilon}{2}$

return c

Experimental results

- The number of iterations does not depend on d.

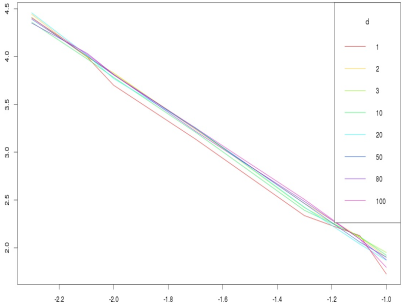

Figure : Number of α-midpoint calculations as a function of ϵ in logarithmic scale for different values of d.

Experimental results

- The running time is approximately $O(\frac{dn}{\epsilon^2})$ (vertical translation in logarithmic scale).

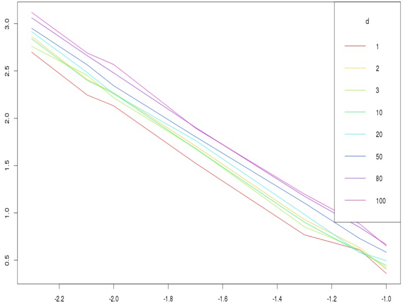

Figure : execution time as a function of ϵ in logarithmic scale for different values of d.

Applications

Hyperbolic geometry arises when considering certain subfamilies of multivariate normal distributions.
For instance, the following subfamilies

- $\mathcal{N}(\mu, \sigma^2 \mathrm{I}_n)$ of n-variate normal distributions with scalar covariance matrix (I_n is the $n \times n$ identity matrix),
- $\mathcal{N}(\mu, \mathrm{diag}(\sigma_1^2, \ldots, \sigma_n^2))$ of n-variate normal distributions with diagonal covariance matrix
- $\mathcal{N}(\mu_0, \Sigma)$ of d-variate normal distributions with fixed mean μ_0 and arbitrary positive definite covariance matrix Σ

are statistical manifolds whose Fisher information metric is hyperbolic.

Applications

In particular, our results apply to the two-dimensional location-scale subfamily:

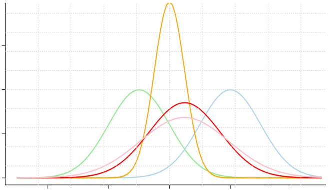

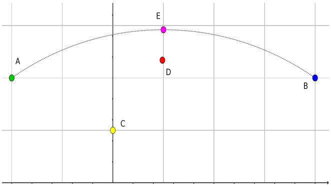

Figure : $\mathrm{MEHB}\,(\mathrm{D})$ of probability density functions (left) in the (μ, σ) superior half-plane (right). $P = \{A, B, C\}$.

Openings

Plugging the EHB and MEHB algorithms to compute clusters centers in the approximation algorithm by [Gonzalez, 1985], we obtain approximate algorithms for

- covering in hyperbolic spaces
- the k-center problem in $O\left(\frac{kNd}{\epsilon^2}\log\left(\frac{1}{\epsilon}\right)\right)$

Algorithm 3: Gonzalez farthest-first traversal approximation algorithm

$C_1 := P, \quad i = 0$
while $i \leq k$ **do**
    $\forall j \leq i$, compute $c_j := \mathrm{MEB}(C_j)$
    $\forall j \leq i$, set $f_j := \mathrm{argmax}_{p \in P}\, \rho(p, c_j)$
    find $f \in \{f_j\}$ whose distance to its cluster center is maximal
    create cluster C_i containing f
    add to C_i all points whose distance to f is inferior to the distance to their cluster center
    increment i
end while
return $\{C_i\}_i$

Openings

The computation of the minimum enclosing hyperbolic ball does not necessarily involve all points $p \in P$.

- Core-sets in hyperbolic geometry
  - the MEHB obtained by the algorithm is an ϵ-core-set
  - differences with the euclidean setting: core-sets are of size at most $\lfloor 1/\epsilon \rfloor$ [Badoiu and Clarkson, 2008]

Thank you!

Bibliography I

- Arnaudon, M. and Nielsen, F. (2013).
  On approximating the Riemannian 1-center.
  Computational Geometry, 46(1):93–104.
- Badoiu, M. and Clarkson, K. L. (2008).
  Optimal core-sets for balls.
  Comput. Geom., 40(1):14–22.
- Gonzalez, T. F. (1985).
  Clustering to minimize the maximum intercluster distance.
  Theoretical Computer Science, 38:293–306.
- Nock, R. and Nielsen, F. (2005).
  Fitting the smallest enclosing Bregman ball.
  In *Machine Learning: ECML 2005*, pages 649–656. Springer.