Computational Geometry 46 (2013) 93–104

Contents lists available at SciVerse ScienceDirect

Computational Geometry: Theory and Applications

www.elsevier.com/locate/comgeo

On approximating the Riemannian 1-center ☆

Marc Arnaudon [a], Frank Nielsen [b,c,*]

[a] Laboratoire de Mathématiques et Applications, CNRS: UMR 7348, Université de Poitiers, Téléport 2, BP 30179, F-86962 Futuroscope Chasseneuil Cedex, France
[b] Ecole Polytechnique, Computer Science Department (LIX), Palaiseau, France
[c] Sony Computer Science Laboratories, Inc. (FRL), 3-14-13 Higashi Gotanda 3F, Shinagawa-Ku, Tokyo 141-0022, Japan

ARTICLE INFO

Article history:
Received 25 January 2011
Accepted 30 April 2012
Available online 3 May 2012
Communicated by Clarkson

Keywords:
1-center
Minimax center
Riemannian geometry
Core-set
Approximation

ABSTRACT

We generalize the Euclidean 1-center approximation algorithm of Bădoiu and Clarkson (2003) [6] to arbitrary Riemannian geometries, and study the corresponding convergence rate. We then show how to instantiate this generic algorithm to two particular settings: (1) the hyperbolic geometry, and (2) the Riemannian manifold of symmetric positive definite matrices.

1. Introduction and prior work

Finding the unique smallest enclosing ball (SEB) of a finite Euclidean point set $P = \{p_1, \ldots, p_n\}$ is a fundamental problem that was first posed by Sylvester [19]. This problem has been thoroughly investigated in the computational geometry community by Welzl [21] and Nielsen and Nock [13], where it is also known as the minimum enclosing ball (MEB), the 1-center problem, or the minimax optimization problem in operations research. In practice, since computing the SEB exactly is intractable in high dimensions, efficient approximation algorithms have been proposed. An algorithmic breakthrough was achieved by Bădoiu and Clarkson [7] that proved the existence of a *core-set* $C \subseteq P$ of *optimal size* $|C| = \lceil \frac{1}{\epsilon} \rceil$ so that $r(C) \leqslant (1+\epsilon) r(P)$ (for any arbitrary $\epsilon > 0$), where $r(S)$ denotes the radius of the SEB of S. Let $c(S)$ denote the ball center, i.e. the minimax center. Since the size of the core-set depends *only* on the approximation precision ϵ and is *independent* of the dimension, core-sets have become widely popular in high-dimensional applications such as supervised classification in machine learning (see for example, the core vector machines of Tsang et al. [20]). In the work of Bădoiu and Clarkson [6], a fast and simple approximation algorithm is designed as algorithm BC-ALG.

It can be proved that a $(1+\epsilon)$-approximation of the SEB is obtained after $\lceil \frac{1}{\epsilon^2} \rceil$ iterations, thereby showing the existence of a core-set $C = \{f_1, f_2, \ldots\}$ of a size at most $\lceil \frac{1}{\epsilon^2} \rceil$: $r(C) \leqslant (1+\epsilon) r(P)$. This simple algorithm runs in time $O(\frac{dn}{\epsilon^2})$, and has been generalized to Bregman divergences by Nock and Nielsen [15] which include the (squared) Euclidean distance, and are the canonical distances of dually flat spaces, including the particular case of self-dual Euclidean geometry. (Note that if we start from the optimal center $c_1 = c(S)$, the first iteration yields a center c_2 away from $c(S)$ but it will converge in the long run to $c(S)$.) Bădoiu and Clarkson [7] proved the existence of optimal ϵ-core-set of size $\lceil \frac{1}{\epsilon} \rceil$. Since finding tight core-sets

☆ Second revision. Source codes for reproducible research available at http://www.informationgeometry.org/RiemannMinimax/.

* Corresponding author at: Ecole Polytechnique, Computer Science Department (LIX), Palaiseau, France.
E-mail addresses: Marc.Arnaudon@math.univ-poitiers.fr (M. Arnaudon), Frank.Nielsen@acm.org (F. Nielsen).

0925-7721/$ – see front matter
http://dx.doi.org/10.1016/j.comgeo.2012.04.007

BC-ALG:

- Initialize the center $c_1 \in P$, and
- Iteratively update the current center using the rule

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

where f_i denotes the farthest point of P to c_i.

requires as a black box primitive the computation of the exact smallest enclosing balls of small-size point sets, we rather consider the Riemmanian generalization of the BC-ALG, although that even in the Euclidean case it does not deliver optimal size core-sets.

Many data-sets arising in medical imaging (see [17]) or in computer vision (refer to [16]) cannot be considered as emanating from vectorial spaces but rather as lying on curved manifolds. For example, the space of rotations or the space of invertible matrices are not flat, as the arithmetic average of two elements does not necessarily lie inside the space.

In this work, we extend the Euclidean BC-ALG algorithm to Riemannian geometry. In the remainder, we assume the reader familiar with basic notions of Riemannian geometry (see [4] for an introductory textbook) in order not to burden the paper with technical Riemannian definitions. However in Appendix A, we recall some specific notions which play a key role in the paper, such as geodesics, sectional curvature, injectivity radius, Alexandrov and Toponogov theorems, and cosine laws for triangles. Furthermore, we consider probability measures instead of finite point sets[1] so as to study the most general setting.

Let M be a complete Riemannian manifold and ν a probability measure on M. Denote by $\rho(x,y)$ the Riemannian distance from x to y on M that satisfies the metric axioms. Assume the measure support $\operatorname{supp}(\nu)$ is included in a geodesic ball $B(o,R)$.

Recall that if $p \in [1,\infty)$ and $f : M \to \mathbb{R}$ is a measurable function then

$$\|f\|_{L^p(\nu)} = \left(\int_M \left|f(y)\right|^p \nu(dy)\right)^{1/p}$$

and

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

Let

$$R_{\alpha,p} = \begin{cases} \frac{1}{2}\min\{\operatorname{inj}(M), \frac{\pi}{2\alpha}\} & \text{if } 1 \leqslant p < 2, \\ \frac{1}{2}\min\{\operatorname{inj}(M), \frac{\pi}{\alpha}\} & \text{if } 2 \leqslant p \leqslant \infty \end{cases} \tag{1}$$

where $\operatorname{inj}(M)$ is the injectivity radius (see Appendix A) and $\alpha > 0$ is such that α^2 is an upper bound for the sectional curvatures in M (in fact replacing M by $B(o,2R)$ is sufficient, so that we can always assume that $\alpha > 0$). For $p \in [1,\infty]$, under the assumption that

$$R < R_{\alpha,p} \tag{2}$$

it has been proved by Afsari [2] that there exists a *unique* point c_p which minimizes the following cost function

$$\begin{aligned} H_p &: M \to [0,\infty] \\ &x \mapsto \left\|\rho(x,\cdot)\right\|_{L^p(\nu)} \end{aligned} \tag{3}$$

with $c_p \in B(o,R)$ (in fact, lying inside the closure of the convex hull of the support of ν).

For a discrete uniform measure viewed as a "point cloud" in a Euclidean space and $p \in [1,\infty)$, we have $H_p(x) = (\frac{1}{n}\sum_{i=1}^n \|p_i - x\|_p^p)^{1/p}$, with $\|\cdot\|_p$ denoting the L_p-norm, and $H_\infty(x)$ is the distance from x to its farthest point in the cloud.

In the general situation the point c_p that realizes the minimum represents a notion of centrality of the measure (e.g., median for $p=1$, mean for $p=2$, and minimax center for $p=\infty$). This center is a *global* minimizer (not only in $B(o,R)$), and this explains why a bound for the sectional curvature is required on the whole manifold M (in fact $B(o,2R)$ is sufficient, see [2]).

[1] We view finite point sets as discrete uniform probability measures.