

Algorithms on negatively curved spaces

[extended abstract]

Robert Krauthgamer*
IBM Almaden

James R. Lee †
Institute for Advanced Study

Abstract

We initiate the study of approximate algorithms on negatively curved spaces. These spaces have recently become of interest in various domains of computer science including networking and vision. The classical example of such a space is the real-hyperbolic space \mathbb{H}^d for $d \geq 2$, but our approach applies to a more general family of spaces characterized by Gromov’s (combinatorial) hyperbolic condition. We give efficient algorithms and data structures for problems like approximate nearest-neighbor search and compact, low-stretch routing on subsets of negatively curved spaces of fixed dimension (including \mathbb{H}^d as a special case). In a different direction, we show that there is a PTAS for the Traveling Salesman Problem when the set of cities lie, for example, in \mathbb{H}^d . This generalizes Arora’s results for \mathbb{R}^d .

Most of our algorithms use the intrinsic distance geometry of the data set, and only need the existence of an embedding into some negatively curved space in order to function properly. In other words, our algorithms regard the inter-point distance function as a black box, and are independent of the representation of the input points.

1 Introduction

The algorithmic and structural theory of finite metric spaces has been a very active and fruitful area of study, with a diverse range of applications in computer science. This connection is most straightforward when the input at hand is equipped with an explicit distance metric, for instance when the distance function represents a similarity measure on a set of data, or the hop-distance between nodes in a network. In these cases, a number of important computational tasks become apparent, e.g. nearest-neighbor search, clustering, routing, object location, finding optimal salesman tours, etc.

Often these tasks are prohibitively difficult on general

metric spaces, and thus one seeks natural and plausible restrictions on the spaces under consideration which could lead to the existence of efficient algorithms and data structures. A classical assumption might be that the data points lie in some finite-dimensional vector space \mathbb{R}^d , where distances are computed according to an ℓ_p norm. Indeed, this is the setting of classical computational geometry. On the other hand, such restrictions are very rigid, and often inappropriate as they not only make assumptions about the pairwise distances between points, but also about their representation.

Intrinsic metric properties. Instead, we would like to place conditions on the *intrinsic* geometry of the input metric space, and ideally these conditions are both independent of the input representation, and stable under small perturbations of the distances. Studies of this type have a number of benefits: They suggest new algorithmic approaches which would have been infeasible under general assumptions, they lend new insight into the understanding of well-known algorithms by exposing the simplest set of restrictions necessary for their analysis, and they can yield algorithms and data structures which are more amenable to efficient implementation.

As an example, consider the *doubling dimension* of a metric space, which was suggested in [16] (based on a classical notion of [4], and inspired by the approach in [18]) as an intrinsic notion of metric dimension that affects algorithmic tractability—we refer to [16] for the definition, and merely note that subsets of \mathbb{R}^d are a special case of spaces having small dimension in this sense. The paper [19] gave algorithms and data structures for nearest-neighbor search problems in such spaces whose efficiency depended on the doubling dimension. Beygelzimer, Kakade, and Langford [7] built upon these techniques to obtain efficient implementations that compare favorably with known approaches for general metric spaces. Finally, in a sequence of papers [20, 17, 12], culminating in the recent result of Cole and Gottlieb, it is shown that one can give data structures whose theoretical efficiency matches the well-known optimal algorithms for low-dimensional Euclidean spaces [3], giving a better understanding of even the classical case.

*IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120, USA. Email: robi@almaden.ibm.com

†Institute for Advanced Study, Einstein Dr., Princeton, NJ 08540, USA. Email: jrl@math.ias.edu. This research was supported by NSF grant DMS-011298.

Negatively curved spaces. In the present paper, we study approximate algorithms on *negatively curved* spaces and their metric generalizations. The classical example of such a space is the real-hyperbolic space \mathbb{H}^d for $d \geq 2$, but significant generalizations are made possible by the pioneering work of Gromov [15] (see Section 2). We are motivated in part by recent studies in networking and vision which suggest that there are a variety of interesting data sets which exhibit negatively curved properties.

Shavitt and Tankel [26] show empirically that the internet topology (namely, the AS graph) embeds with better accuracy (smaller average distortion) into a low-dimensional hyperbolic space than into a Euclidean space of comparable dimension. They also demonstrate the application of their embedding to delay estimation and server selection. Begelfor and Werman [6] obtain similar conclusions for several data sets including the internet (where distance is number of hops), actors database (distance based on number of joint movies), and genomic data (distance between potential genes computed using BLAST), using a different notion of accuracy—the average additive distortion. They also suggest that such embeddings may be useful for clustering problems. Finally, Sharon and Mumford [25] model the similarity between 2-D objects in the plane using conformal mappings (i.e. maps that locally preserve angles). The distance function they define turns out to represent the geodesic distance along a certain manifold of non-positive curvature.

Another motivation for studying negatively curved spaces concerns the overarching approach of all algorithms based on doubling dimension restrictions—hierarchical local search. These algorithms depend crucially on the fact that the “volume growth” of balls in such spaces is only polynomial in the radius of the ball. At the end of the day, all the algorithmic approaches on spaces of small doubling dimension boil down to (sometimes very clever) forms of “brute force” local search. For negatively curved spaces, this volume growth constraint no longer exists because—by their very nature—such spaces exhibit exponential volume growth. Thus in order to solve geometric problems on these spaces, we are forced to confront exponential growth in a tractable way.

To demonstrate our techniques, we consider three representative problems: Nearest-neighbor search, sparse spanners and compact routing tables, and approximation algorithms for the Traveling Salesman Problem (TSP). While based on similar principles, the solution to each requires overcoming distinct obstacles. An important point is that many of our algorithms depend only on the *intrinsic geometry* of the distance function $d(\cdot, \cdot)$, and not on any representation of the points in a specific negatively curved space.

In Section 1.2, we give a precise statement of our results, and an overview of the techniques involved. For readers

interested in a crash course in negatively curved geometry, we refer to the introduction in Section 2, where we discuss *Gromov hyperbolic spaces*—these spaces are defined by a *combinatorial* condition of Gromov, which is completely divorced from any continuous (e.g. manifold) structure.

1.1 Preliminaries

We use standard notions from geometric analysis. Let (X, d_X) be a metric space. For $x \in X, r \geq 0$, we write $B(x, r) = \{y \in X : d_X(x, y) \leq r\}$ for the *closed ball of radius r about x* . Given another metric space (Y, d_Y) and a map $f : X \rightarrow Y$, we say that f is a *distortion C embedding of X into Y* if there exists a number $c_0 > 0$ for which

$$d_X(x, y) \leq c_0 \cdot d_Y(f(x), f(y)) \leq C \cdot d_X(x, y),$$

holds for all $x, y \in X$. We refer to such a map as a *C -embedding of X into Y* .

The *doubling constant* of X , written $\lambda(X)$, is the smallest λ such that every ball in X can be covered by λ balls of half the radius. A subset $S \subseteq X$ is *ε -separated* if $d(x, y) \geq \varepsilon$ for every $x, y \in S$. A set S is an *ε -net* if it is ε -separated and $X \subseteq \bigcup_{x \in S} B(x, \varepsilon)$.

1.2 Results and techniques

We now describe our results, and give a brief overview of the techniques involved. In Section 2, we give a short introduction to Gromov hyperbolic metric spaces, along with some equivalent characterizations (via the shapes of triangles and divergence of geodesics) that will be useful later. In general, our results hold for arbitrary subsets of locally doubling, geodesic, δ -hyperbolic spaces (see Section 2 for the definitions—this includes, e.g. \mathbb{H}^d), and most of the algorithms and constructions only access the distance function as an oracle. To avoid messy quantitative statements at this point, we will let \mathcal{M} represent a space satisfying these properties (for some fixed setting of the other parameters), and $X \subseteq \mathcal{M}$ an arbitrary n -point subset.

Random decompositions, spanners, and routing. In Section 3, we turn to *random low-diameter decompositions* of negatively curved spaces. Such decompositions are a well-studied and essential tool in the field of discrete metric spaces (see, e.g. [5, 24, 16, 23, 21, 11]). Using these decompositions, we construct a $(1 + \varepsilon)$ -stretch routing scheme for X , where the average amount of memory needed per node is only $O(\log \Phi)^2$ —here, Φ is the ratio of the largest to smallest distance in X . This application follows almost immediately from efficiently constructible *hierarchical* decompositions with small support (see [11]). We also show how, for every $\varepsilon > 0$, one can construct a $(1 + \varepsilon)$ -spanner of X with only $(\frac{1}{\varepsilon})^{O(1)} \cdot n$ edges, but using Steiner nodes. If we are not allowed Steiner nodes, it is not difficult to

see that $(1 + \varepsilon)$ -stretch is impossible to achieve; without Steiner nodes, we can still construct constant-stretch linear-sized spanners.

The difficulty in constructing random “padded” decompositions (especially with small support) comes from the phenomenon of exponential volume growth in negatively curved spaces. But this growth comes at a price—geodesics in such spaces “diverge” at an exponential rate. This creates a subtle trade-off that is best viewed “at infinity.” In fact, such a structure at infinity exists, called the Gromov boundary of X , and denoted ∂X (see Section 2). If X is locally doubling (i.e. is “flat” at small scales), then this structure is manifested on ∂X which, by a result of Bonk and Schramm [8] will also be doubling. Because of exponential divergence of geodesics, there is a tight relationship between ∂X and X , and thus we can use known decompositions for doubling spaces [16] as a crucial step in our constructions for X (see Figure 2(b)).

While the previous description is the best way to understand our approach, it does not yield efficient algorithms (in particular, we cannot actually go to infinity), but there is a way of transferring such results to the intrinsic algorithmic setting by explicitly using exponential divergence of geodesics in place of the structure of ∂X .

Nearest-neighbor search. In Section 4, we design approximate nearest neighbor search algorithms for a data set $X \subseteq \mathcal{M}$ and queries $q \in \mathcal{M}$. Specifically, our main technical result is an efficient algorithm achieving $O(1)$ -additive approximation, with query time $O(\log^2 n)$ and with preprocessing storage $O(n^2)$. The algorithm is intrinsic (or black-box) in the sense that it only requires distance computations over the input points. Our algorithm easily extends to $(1 + \varepsilon)$ -approximations—since our space \mathcal{M} (and hence X) is assumed to be locally doubling, we can locally employ the search procedures of [19].

We employ a divide-and-conquer approach based on structural properties of negatively curved spaces. Whereas algorithms for Euclidean spaces might try to recursively confine the set of feasible answers to smaller and smaller balls or half-spaces, our approach uses a sequence of hierarchically structured “cones” (such sets are allowed to be significantly larger in the “vertical” direction than the “horizontal” direction). A significant challenge is posed by the need for the search structure to occupy only polynomial space, and for the algorithm to only ask distance queries $d(q, x)$ for $x \in X$ (as opposed to $x \in \mathcal{M}$).

Traveling salesman tours. Finally, in Section 5, we give an overview of our approach to obtaining a PTAS for TSP on n -point subsets $X \subseteq \mathcal{M}$. Actually, we are only able to obtain a quasi-PTAS (i.e. quasi-polynomial running time) for such X , but we obtain a true PTAS for, e.g. subsets of \mathbb{H}^d (or any space where the local doubling property of

\mathcal{M} is replaced by a “uniformly locally Euclidean” property, discussed in Section 2).

Essentially, we show that for every finite subset $X \subseteq \mathcal{M}$ and every $\varepsilon > 0$, X admits an embedding into a distribution over dominating “neighborhood-trees” with distortion $1 + \varepsilon$. A neighborhood-tree is basically a tree of neighborhoods in \mathcal{M} , where each neighborhood has diameter at most $\tau = \tau(\varepsilon)$. Since \mathcal{M} is locally doubling and geodesic, for every fixed value of τ , the local neighborhoods are doubling (or close to Euclidean if \mathcal{M} satisfies a stronger property), hence inside each neighborhood we can use known techniques for doubling [27] or finite-dimensional Euclidean [2] spaces to compute near-optimal salesman tours. In the latter case, we need to observe that Arora’s geometric TSP framework can be extended to subsets of Euclidean space that are distorted by an arbitrary constant.

2 Gromov hyperbolic metric spaces

In this section, we give the basic definitions and properties of spaces that one might consider “negatively curved.” Such spaces include classical examples like the real and complex-hyperbolic spaces \mathbb{H}^d and $\mathbb{C}\mathbb{H}^d$, and Riemannian manifolds of (strictly) negative sectional curvature, but also many discrete spaces like graph-theoretic trees and the Cayley graphs of “word-hyperbolic” groups. The ability to argue at this level of generality is due, in large part, to Gromov [15] who introduced the notion of a δ -hyperbolic metric space. We refer to [14] for a comprehensive treatment of such spaces.

Let (X, d) be any metric space. We fix a basepoint $r \in X$, which will sometimes be implicit in the definitions that follow. For any $x, y \in X$, we write $|x - y| = d(x, y)$ and $|x| = |x - r|$. One defines the *Gromov product* of $x, y \in X$ with respect to r by

$$(x|y)_r = \frac{1}{2} (|x - r| + |y - r| - |x - y|),$$

where we will sometimes omit the basepoint r and simply write $(x|y)$. One calls a space δ -hyperbolic if, for every $x, y, z, r \in X$, we have

$$(x|y)_r \geq \min\{(x|z)_r, (y|z)_r\} - \delta. \quad (1)$$

Observe that the case $\delta = 0$ coincides precisely with the family of metric trees, and if one thinks of r as the root of the tree, then $(x|y)_r$ corresponds exactly to the distance from r to the least common ancestor of x and y . We note that if the inequality (1) holds only for a fixed basepoint $r \in X$, then it holds with respect to any basepoint if δ is replaced by 2δ , i.e. X is 2δ -hyperbolic (see e.g. [1, Lem. 2.3]).




Figure 1. Divergence of geodesics: $u(t) \gtrsim \exp(Ct)$ for hyperbolic spaces.

Geodesics. A *geodesic segment* in X is a map $\gamma : [0, L] \rightarrow X$ for some $L > 0$ which is an isometric embedding of $[0, L]$ onto its image. We say that the space X is a *geodesic space* if every pair of points $x, y \in X$ can be connected by a geodesic. For two points x, y in a geodesic space X , we will sometimes write $[x, y]$ to represent some geodesic between them.

In geodesic spaces, there are equivalent notions of “hyperbolicity” that will be useful.

Thin triangles. The following notion is due to Rips. Let $x, y, z \in X$ be distinct points. A *geodesic triangle* in X is the set $[x, y] \cup [y, z] \cup [z, x] \subseteq X$ for some choice of geodesics. Such a triangle is called δ -*thin* if, for every $w \in [x, y]$, we have $d(w, [y, z] \cup [z, x]) \leq \delta$ (and similarly for the other sides $[y, z], [z, x]$). We will say that X is δ -*thin* if every geodesic triangle in X is δ -thin. It is an exercise to see that a geodesic metric space X is δ -thin if and only if X is $O(\delta)$ -hyperbolic.

Exponential divergence of geodesics. In general, we define a *continuous path* in X to be the image of a continuous map $P : [0, 1] \rightarrow X$. We can then define the length of P as

$$\text{length}(P) = \inf_{x_1, \dots, x_m} \sum_{i=1}^{m-1} |P(x_i) - P(x_{i+1})|,$$

where the infimum is over all subdivisions $0 = x_1 < x_2 < \dots < x_m = 1$ of $[0, 1]$.

Let $\gamma_1, \gamma_2 : [0, L] \rightarrow X$ be two geodesic segments with the same initial point $x_0 = \gamma_1(0) = \gamma_2(0)$. They are said to *diverge at rate* $u : [0, \infty) \rightarrow \mathbb{R}$ if, for every pair of numbers t and T , with $0 \leq T, t + T \leq L$, the following property holds: Whenever $|\gamma_1(T) - \gamma_2(T)| > u(0)$, if a continuous path P connects $\gamma_1(T+t)$ and $\gamma_2(T+t)$ in the closure of $X \setminus B(x_0, T+t)$, then the length of P is at least $u(t)$. (See Figure 1, where the thick dotted line represents the boundary of $B(x_0, T+t)$.)

We will say that *geodesics in X diverge at an exponential rate* if there exists a constant $C > 0$ and a map $u : [0, \infty) \rightarrow \mathbb{R}$ such that $u(t) \geq \Omega(\exp(Ct))$ as $t \rightarrow \infty$, and every pair of geodesics γ_1, γ_2 in X diverge at rate u . The following theorem relates exponential divergence of geodesics to hyperbolicity (see, e.g. [1, Th. 2.19]).

Theorem 2.1. A geodesic metric space X is δ -hyperbolic for some $\delta > 0$ if and only if geodesics in X diverge at an exponential rate. Furthermore, the rate of divergence and the value of δ depend solely on each other.

Local geometry. Clearly the condition in (1) is relevant only for the “large scale” geometry of X . In particular, if $(x|z), (y|z) \leq \delta$, then the condition becomes trivial. On the other hand, the interaction between the local geometry of a space (below the “hyperbolicity radius” δ) with the negative curvature manifests itself in the large-scale geometry of X . Thus we define three progressively stronger ways of bounding the local geometry of a space.

1. We will say that X has *local geometry of type* (λ_0, R_0) if every ball of radius R_0 in X can be covered by λ_0 balls of half the radius.
2. We will say that X is *locally doubling* if there exist constants $\lambda_0, R_0 > 0$ such that X has local geometry of type (λ_0, R) for every $R \leq R_0$.
3. Finally, we will say that X is *locally Euclidean* if there exist constants $C, R, k > 0$ such that every ball of radius R in X admits a C -embedding into \mathbb{R}^k .

The following well-known lemma shows that in a geodesic space, a uniform bound on the local geometry implies that the space has at most singly-exponential growth. We omit its proof from this extended abstract.

Lemma 2.2. If a geodesic space (X, d) has local geometry of type (λ_0, R_0) , then it also has local geometry of type $(\lambda_0^{2R/R_0+1}, R)$ for every $R \geq R_0$.

Examples. Recall that there exists a $\delta > 0$ such that for every $d \geq 1$, the real and complex hyperbolic spaces $\mathbb{H}^d, \mathbb{C}\mathbb{H}^d$ are δ -hyperbolic. Furthermore, these spaces are locally Euclidean. In fact, Topogonov’s comparison theorem implies that if X is any complete simply-connected Riemannian manifold of dimension n with “pinched” negative sectional curvature κ (see e.g. [22]), i.e. $b \leq \kappa \leq a < 0$, then X is δ -hyperbolic and locally Euclidean, with parameters depending only on a and b . By definition, all these spaces are geodesic. Another interesting family of δ -hyperbolic spaces is given by the *word-hyperbolic groups* (see [1, 14]). Clearly the word metric on finitely-generated groups is locally doubling (as is the shortest path metric on any graph with uniformly bounded degrees).

The Gromov boundary. One of the most important objects associated to a Gromov hyperbolic space is its *boundary at infinity*. We recall the standard construction. Let X be a δ -hyperbolic space, and fix a basepoint $r \in X$. A sequence $\{x_i\}_{i=1}^\infty \subseteq X$ is said to *converge at infinity* if

$$\lim_{i,j \rightarrow \infty} (x_i | x_j)_r = \infty.$$

We define the *Gromov boundary* of X , denoted ∂X , as the set of equivalence classes of convergent sequences with equivalence relation

$$\{x_i\} \equiv \{y_i\} \iff \lim_{i \rightarrow \infty} (x_i | y_i)_r = \infty.$$

This equivalence relation is easily seen to be independent of the choice of basepoint.

We now extend the product operation to the boundary in the standard way. Let $a, b \in \partial X$. We define

$$(a|b)_r = \sup \left\{ \liminf_{i \rightarrow \infty} (x_i | y_i)_r : \{x_i\} \in a, \{y_i\} \in b \right\}.$$

For $a \in \partial X, y \in X$, define

$$(a|y)_r = (y|a)_r = \sup \left\{ \liminf_{i \rightarrow \infty} (x_i | y)_r : \{x_i\} \in a \right\}.$$

Observe that for any $a, b, x \in X \cup \partial X$, we have

$$(a|b) \geq \min\{(a|x), (b|x)\} - 2\delta, \quad (2)$$

using (1) and the definitions above.

Visual spaces. We will say that a Gromov hyperbolic space X is *visual (with respect to the basepoint r)* if, for every $x \in X$, there exists an isometric map $\gamma : [0, \infty) \rightarrow X$ for which $\gamma(0) = r, \gamma(t_0) = x$ for some $t_0 \geq 0$, and such that $\{\gamma(t_i)\}$ converges at infinity for every sequence $\{t_i\}$ with $t_i \rightarrow \infty$. In words, there is a *geodesic ray* starting from r , traveling through x , and converging at infinity. We can identify γ with $\{\gamma(t_i)\}$ so that $\gamma \in \partial X$. We mention, for instance, that \mathbb{H}^d is visual with respect to any basepoint.

The canonical gauge on ∂X . There is a family of standard metrics on the boundary of a Gromov hyperbolic space X . If $x, y \in \partial X, r \in X$, and $\epsilon > 0$, one defines

$$d_{r,\epsilon}(x, y) = \inf \left\{ \sum_{i=1}^n e^{-\epsilon(x_{i-1} | x_i)_r} \right\},$$

where the infimum is taken over all finite sequences $x = x_0, x_1, \dots, x_n = y \in \partial X$, with the convention that $e^{-\infty} = 0$. The following lemma can be found in [14, Ch. 7]—in a sense, it captures succinctly the effect of exponential divergence at infinity.

Lemma 2.3. *There exists a constant $C_0 > 0$ such that if X is δ -hyperbolic and $\epsilon\delta \leq C_0$, then*

$$\frac{1}{2} e^{-\epsilon(x|y)_r} \leq d_{r,\epsilon}(x, y) \leq e^{-\epsilon(x|y)_r}, \quad \forall x, y \in \partial X.$$

Following [8], and in light of Lemma 2.3, we define the *canonical gauge* $\mathcal{G}(X)$ on ∂X as

$$\mathcal{G}(X) = \{d_{r,\epsilon} : 0 < \epsilon \leq C_0/\delta\},$$

We refer to Figure 2(a) for a graphical representation of ∂X in an Escher painting (of the Poincaré disc model of \mathbb{H}^2).

We now discuss briefly the proper setting for Gromov hyperbolic spaces—coarse geometry—and also recent results on embeddings of such spaces.

Coarse geometry. A map $f : X \rightarrow Y$ between two metric spaces X and Y is called a *quasi-isometry* if there exist constants $K \geq 1, c \geq 0$ such that

$$\frac{1}{K} |x - y| - c \leq |f(x) - f(y)| \leq K|x - y| + c$$

for every $x, y \in X$. In other words, we allow both multiplicative and additive distortion of distances. Such a map is called a (K, c) -*quasi-isometry*. Since the δ -hyperbolic condition is independent of the local geometry of a space, it turns out that “quasi-isometric” notions are the proper setting in which to discuss Gromov hyperbolic spaces.

The relevance of such notions to computer science is that they are significantly more flexible. For instance, one could define a c -rough geodesic $\gamma : [0, L] \rightarrow X$ as a map which is a $(1, c)$ -quasi-isometry onto its image. (The image of such a map might be a discrete set in X instead of a continuous path.) This approach is taken, for instance, in [8]. We remark that all our results hold in the corresponding “coarse” setting (so anywhere we assume a space to be geodesic, we might as well assume it is only roughly geodesic for some $c > 0$), but we avoid these generalizations for simplicity of presentation. Similar approaches have been taken in the computer science literature, e.g. the concept of (α, β) -spanners [13].

Embeddings of Gromov hyperbolic spaces. Finally, we mention a couple of results about embedding δ -hyperbolic spaces into other spaces, though we do not use these results in the present work. It is known, through work of Buyalo and Schroeder, that every geodesic δ -hyperbolic space with local geometry of type (λ_0, R_0) admits a quasi-isometric embedding into a product of a bounded number of infinite complete binary trees [9, 10]. Furthermore, Bonk and Schramm [8] showed that every geodesic δ -hyperbolic space with local geometry of type (λ_0, R_0) admits a quasi-isometric embedding into a real-hyperbolic space \mathbb{H}^d . In both cases, the quantities involved in the embeddings (e.g. dimension, quasi-isometric distortion, number of trees) depend only on δ, λ_0, R_0 . Products of trees are (obviously) not themselves Gromov hyperbolic, and it seems far more difficult to design algorithms for such spaces.

3 Probabilistic decompositions

Let (X, d) be a metric space. We begin by recalling the notion of a *padded (probabilistic) decomposition* of X . If P is a partition of X and $x \in X$, we use $P(x)$ to denote the unique set $P \in \mathcal{P}$ which contains x . Let μ be a distribution over partitions of X . We say that μ is τ -*bounded* if, for every partition $P \in \text{supp}(\mu)$ and every $S \in P$, we have $\text{diam}(S) \leq \tau$.

We say that a τ -bounded distribution is (α, β) -*padded* if the following holds for every $x \in X$,

$$\Pr_{P \in \mu} [B(x, \tau/\alpha) \subseteq P(x)] \geq \beta.$$

We will say that μ is α -*padded* if it is $(\alpha, \frac{1}{2})$ -padded. Let $\mathcal{P} = (P_1, P_2, \dots, P_k)$ be a finite sequence of partitions of X . The *decomposition induced by \mathcal{P}* is the distribution P_i where i is a random variable uniformly distributed over $\{1, \dots, k\}$.

We recall the following decomposition theorem from [16] (although the focus therein is on doubling metrics and thus it is stated slightly differently.)

Theorem 3.1 ([16]). *There exists a constant $C \geq 1$ such that for every $\tau > 0$ and $\lambda > 0$, if the metric space X is locally doubling with parameters (λ, τ) , then there exists a sequence of partitions, \mathcal{P} , that induces a τ -bounded $O(\log \lambda)$ -padded decomposition of X , and furthermore $|\mathcal{P}| \leq O(\log \lambda \log \log \lambda)$.*

The main result of this section is the following decomposition theorem for Gromov hyperbolic spaces. It is weaker than our most general decomposition theorem, and is not obviously algorithmic (in particular, we will make liberal use of ∂X), but it contains all of the essential ideas.

Theorem 3.2. *Let (X, d) be a visual, geodesic δ -hyperbolic space that is locally doubling with parameters (λ_0, δ) . Then there exists a threshold $T = \Theta(\delta \log \log \lambda_0)$ such that the following holds. For every $\tau > 0$, there is a sequence \mathcal{P} of τ -bounded partitions of X such that $|\mathcal{P}| \leq O(\log \lambda_0 (\log \log \lambda_0)^2)$, and the distribution induced by \mathcal{P} is: $O(\log \lambda_0 \log \log \lambda_0)$ -padded if $\tau \leq T$, and $(5, \frac{1}{4})$ -padded if $\tau \geq T$.*

Before proving Theorem 3.2, let us present one of its key ingredients. Fix some metric $d_{\partial X} \in \mathcal{G}(X)$. Bonk and Schramm [8] show that if X is δ -hyperbolic with bounded growth at some scale, then for every $d_{\partial X} \in \mathcal{G}(X)$, one has $\dim_A(\partial X, d_{\partial X}) < \infty$, where $\dim_A(\cdot)$ is the Assouad dimension [4]. Following their proof, we state below a quantitative version of this result, which is a (straightforward) generalization of [8]. We recall that $\lambda(X, d)$ is the (global) doubling constant of the metric space (X, d) .

Theorem 3.3. *Let X be a geodesic δ -hyperbolic metric space and suppose X has local geometry of type (λ_0, R_0) for some $R_0 \geq \delta$. Then $\lambda(\partial X, d_{\partial X}) \leq \lambda_0^{O(R_0/\delta)}$.*

We will also need the following simple proposition.

Proposition 3.4. *Let (X, d) be a δ -hyperbolic space and suppose $x, y \in X$ and $a, b \in \partial X$ satisfy $(a|x) \geq |x|$ and $(b|y) \geq |y|$. Then*

$$\min\{|x|, |y|, (a|b)\} - 4\delta \leq (x|y) \leq (a|b) + 4\delta.$$

Proof. Two applications of (2) imply that

$$\begin{aligned} (a|b) &\geq \min\{(a|x), (x|y), (y|b)\} - 4\delta \\ &\geq \min\{|x|, (x|y), |y|\} - 4\delta. \end{aligned}$$

By definition, $(x|y) \leq \min\{|x|, |y|\}$, and thus we get $(a|b) \geq (x|y) - 4\delta$. On the other hand, applying (2) twice again, we have

$$\begin{aligned} (x|y) &\geq \min\{(x|a), (a|b), (b|y)\} - 4\delta \\ &\geq \min\{|x|, (a|b), |y|\} - 4\delta. \end{aligned}$$

□

We are now ready to prove Theorem 3.2. The general idea is as follows: As in Figure 2(b), we will use a partition of ∂X to induce a partition of X itself, and since $\lambda(\partial X)$ is small by Theorem 3.3, using Theorem 3.1 we get good decompositions for the boundary. The key connection between ∂X and X is given by Lemma 2.3.


Proof of Theorem 3.2. Setting $T = c\delta \log \log \lambda_0$ for a sufficiently large constant $c > 0$, the following holds.

First, we observe the existence of decompositions for small scales τ . If $\tau \leq \delta$, we can simply use the locally doubling condition along with Theorem 3.1. If $\delta < \tau \leq T$, then by Lemma 2.2 we know that X has local geometry of type $(\lambda_0^{2\tau/\delta+1}, \tau)$, hence applying Theorem 3.1, there exists a sequence \mathcal{P} partitions of X such that the induced decomposition is τ -bounded and has $O(\frac{\tau}{\delta} \log \lambda_0) = O(\log \lambda_0 \log \log \lambda_0)$ padding, and furthermore $|\mathcal{P}| = O(\log \lambda_0 (\log \log \lambda_0)^2)$.


Next, we handle the more difficult case of large scales $\tau \geq T$. The decomposition is constructed in two stages as follows. The first step decomposes X into two sets of concentric annuli of width $\Delta = \tau/5$. For $u \in \{0, 1\}$ and $k \in \mathbb{N}$, define

$$A_k^u = \{x \in X : (k - \frac{u}{2}) \Delta \leq |x| < (k + 1 - \frac{u}{2}) \Delta\}.$$

It is clear that for each $u \in \{0, 1\}$, $X = \bigcup_{k \geq 0} A_k^u$ is a partition of X . The second step further decomposes each annulus A_k^u separately. To this end, fix $u \in \{0, 1\}$ and $k \geq 1$. (If $k = 0$ then there is no further decomposition.) For every



(a) The boundary at infinity in an Escher painting



(b) Partitioning annuli using a partition of ∂X

Figure 2. A decomposition of the boundary ∂X induces a partition of X .

$x \in A_k^u$, fix a geodesic ray γ_x with $\gamma_x(0) = r$, $\gamma_x(|x|) = x$, and such that γ_x converges to infinity. Such a ray exists since X is assumed to be visual with respect to r . Observe that $(x|\gamma_x) = |x|$. Recall that, for a some $\epsilon = \epsilon(\delta)$, we have by Lemma 2.3

$$\frac{1}{2}e^{-\epsilon(ab)} \leq d_{\partial X}(a, b) \leq e^{-\epsilon(ab)}, \quad \forall a, b \in \partial X. \quad (3)$$

Clearly via the correspondence $x \mapsto \gamma_x$, any partition of ∂X induces a partition of A_k^u (à la Figure 2(b)). Set $\Delta^* = \frac{1}{2}e^{-\epsilon(k-1-u/2)\Delta}$, and use Theorem 3.1 derive a sequence \mathcal{P}^* of Δ^* -bounded partitions of $(\partial X, d_{\partial X})$. By Theorem 3.3, we know that

$$\begin{aligned} |\mathcal{P}^*| &\leq O((\log \lambda(\partial X))(\log \log \lambda(\partial X))^2) \\ &\leq O(\log \lambda_0(\log \log \lambda_0)^2), \end{aligned}$$

and thus, for each $u \in \{0, 1\}$, the decomposition of ∂X induces a sequence of $O(\log \lambda_0(\log \log \lambda_0)^2)$ partitions of X . By going over the two values for u , we obtain a sequence of $O(\log \lambda_0(\log \log \lambda_0)^2)$ partitions of X .

Let us now show that this decomposition of X is τ -bounded. Consider a partition P of X constructed in this way, and two points $x, y \in X$ for which $P(x) = P(y)$. Let $u \in \{0, 1\}$ be the value used to construct P . Clearly, there exists $k \geq 0$ such that both $x, y \in A_k^u$. We may assume that $k \geq 1$, as otherwise we're done. Let γ_x and γ_y be the rays corresponding to x and y . Thus, $d_{\partial X}(\gamma_x, \gamma_y) \leq \Delta^*$. Applying (3), we get

$$(\gamma_x|\gamma_y) \geq (k-1-u/2)\Delta,$$

which implies, by Proposition 3.4, that

$$\begin{aligned} |x-y| &\leq |x|+|y|-2\min\{|x|, |y|, (\gamma_x|\gamma_y)\}+8\delta \\ &\leq 2(k+1-u/2)\Delta-2(k-1-u/2)\Delta+8\delta \\ &\leq 4\Delta+8\delta. \end{aligned}$$

If $c > 0$ is chosen to be a sufficiently large constant, we get $8\delta \leq T/5 \leq \tau/5$ and thus $|x-y| \leq \tau$.

We now analyze the padding of this decomposition of X . Fix $x \in X$, and consider $B(x, \Delta)$ where $\Delta = \tau/5$. Then there exists $u \in \{0, 1\}$ and $k \in \mathbb{N}$ such that $B(x, \Delta) \subseteq A_k^u$. If $k = 0$ then we are done (since A_0^u is not further decomposed). Otherwise, consider any $y \in B(x, \Delta)$. Applying Proposition 3.4 we get that

$$(\gamma_x|\gamma_y) \geq \frac{1}{2}(|x|+|y|-|x-y|)-4\delta \geq (k-\frac{u}{2}-\frac{1}{2})\Delta-4\delta.$$

which implies, by applying (3),

$$\begin{aligned} d_{\partial X}(\gamma_x, \gamma_y) &\leq e^{-\epsilon(\gamma_x|\gamma_y)} \\ &\leq 2\Delta^* \cdot e^{-\epsilon(\Delta/2-4\delta)} \leq 2\Delta^* \cdot e^{-\epsilon\Delta/3}. \end{aligned}$$

Notice that $\epsilon\Delta \geq C_0c \log \log \lambda_0$, thus if $c > 0$ is chosen to be a sufficiently large constant, then $d_{\partial X}(\gamma_x, \gamma_y) \leq \frac{\Delta^*}{O(\log \lambda_0)}$. In particular, using the padding property of \mathcal{P}^* , with probability $\frac{1}{2}$ over the choice of a random partition $P \in \mathcal{P}^*$, we have

$$\{\gamma_y : y \in B(x, \Delta)\} \subseteq B_{\partial X}\left(\gamma_x, \frac{\Delta^*}{O(\log \lambda_0)}\right) \subseteq P(\gamma_x),$$

which implies that $B(x, \Delta)$ is also contained completely within the corresponding induced partition. Thus, the $O(\log \lambda_0)$ -padding for ∂X implies $(5, \frac{1}{4})$ -padding for X . \square

For an intrinsic version of this decomposition, along with applications to spanners and compact routing tables, we refer to the full version of the paper.

4 Nearest-neighbor search

A nearest-neighbor search algorithm is called *intrinsic* (or *black-box*) if its only access to the metric space is through distance computations involving the input point set S and the query point q . It may be crucial to assume certain properties of the metric space (X, d) , e.g. that it is geodesic, but the algorithm has no access to such points in $X \setminus S$, even though they are known to exist.

In this section we give an efficient algorithm for intrinsic, approximate nearest-neighbor search in Gromov hyperbolic metric spaces. Technically, we require our point set to be a subset of some geodesic hyperbolic metric space with a bound on the local geometry. Our main technical result below is a scheme for nearest-neighbor with an $O(\delta)$ -additive approximation guarantee.

Theorem 4.1. *Let (X, d) be a δ -hyperbolic metric space that is geodesic and has local geometry of type $(\lambda, \delta/3)$. Then there exists an intrinsic nearest-neighbor search algorithm that preprocesses an n -point subset $S \subseteq X$ using $O(n^2)$ storage, and can find for every given query $q \in X$ a point $a \in S$ such that $d(q, a) \leq d(q, S) + O(\delta)$, where the query procedure has running-time $\lambda^{O(1)} \log^2 n$.*

The theorem immediately extends to local geometry of type (λ, R) as follows. If $R > \delta/3$ then setting $\delta' = 3R$ it is clear that (X, d) is δ' -hyperbolic and has local geometry of type $(\lambda, \delta'/3)$. If $R < \delta/3$ then Lemma 2.2 implies that (X, d) has local geometry $(\lambda', \delta/3)$ for a suitable $\lambda' = \lambda^{O(\delta/R)}$.

To simplify the exposition and demonstrate the basic technique, we first show in Section 4.1 a simpler scheme whose query time is similar to the above, but it is not intrinsic and does not achieve the desired storage requirement. In the full version, we build on these ideas, to arrive at a full scheme that proves Theorem 4.1. It is possible to extend the scheme to achieve $(1 + \varepsilon)$ -approximation in metric spaces are locally doubling by combining the full scheme with the techniques of [19].

Preliminaries and basic observations. Throughout this section, let (X, d) be a geodesic δ -hyperbolic metric space with local geometry of type $(\lambda, \delta/3)$, and fix a basepoint $r \in X$. For $z \in X$ and $t \geq 0$ define

$$X_z^t = \{x \in X : (x|z) \geq |z| - t\}.$$

For example, if (X, d) is a tree metric (i.e. 0-hyperbolic), then X_z^t corresponds to a subtree, and in particular X_z^0 is the subtree rooted under z .

Proposition 4.2. *Let $x, y, c \in X$ and $k > 0$ be such that $x \in X_c^{k\delta}$ and $y \notin X_c^{(k+1)\delta}$. Then*

$$|x - y| \leq |x - c| + |c - y| \leq |x - y| + (2k + 2)\delta.$$

Proof. The first inequality is just the triangle inequality, so we only need to prove the second inequality. By the δ -hyperbolicity (1), we have

$$(y|c) \geq \min\{(x|y), (x|c)\} - \delta.$$

The minimum in the RHS cannot be attained by $(x|c)$, because we know that $(y|c) < |c| - (k + 1)\delta$ and $(x|c) \geq |c| - k\delta$. Thus, $(y|c) \geq (x|y) - \delta$, which by rearranging yields

$$|x - y| \geq |y - c| + |x| - |c| - 2\delta.$$

Observing that $x \in X_c^{k\delta}$ implies $|x| - |c| \geq |x - c| - 2k\delta$, and we get the desired inequality. \square

We say that $S \subseteq X$ is t -separated for $t > 0$ if the minimum inter-point distance in S is at least t . For example, a t -net is clearly t -separated. The following proposition follows using argument similar to Lemma 2.2.

Proposition 4.3. *If $N \subseteq X$ is δ -separated and contained in a ball of radius $t \geq \delta$, then $|N| \leq \lambda^{O(t/\delta)}$.*

4.1 A simple scheme for fast querying

In this subsection we prove a weaker version of Theorem 4.1, in which the query time is similar to that stated in the theorem, but the nearest-neighbor search algorithm is not intrinsic (black-box), and it does not achieve the desired storage requirement. First, we present a basic structural fact.

Lemma 4.4 (Separator point in hyperbolic metric). *Let $S \subseteq X$ be 20δ -separated and suppose $1 < |S| < \infty$. Then there exists a point $c \in X$ such that*

$$|S|/\lambda^{O(\delta)} \leq |S \cap X_c^\delta| \leq |S \cap X_c^{3\delta}| \leq |S|/2.$$

This lemma can be seen to generalize the following simple fact about a vertex separator in trees: *Every rooted tree $T = (V, E)$ with maximum degree λ contains a vertex $z \in V$, such that the number of vertices in the subtree of T rooted at z is between $|V|/2\lambda - 1$ and $|V|/2$ (inclusive).* For our intended application it is crucial that the lemma produces bounds on both X_c^δ and $X_c^{3\delta}$. The distinction between these two sets is hidden in the preceding fact about trees, because in trees $\delta = 0$ and thus the two sets are identical.

Proof of Lemma 4.4. The following proposition will be key to proving the lemma. We will actually require only the case $k = 2$.

Proposition 4.5. *For every $z \in X$ and $k \geq 1$ there is $N \subseteq X$ of size $|N| \leq \lambda^{O(k)}$, such that $X_z^{k\delta} \setminus B_X(z, 3k\delta) \subseteq \bigcup_{y \in N} X_y^\delta$. Furthermore,*

$$N \subseteq \{x \in X : |x| \geq |z| + (k - 1)\delta\}.$$

Proof. Fix $z \in X$ and $k \geq 1$ and let N' be a δ -net in X . Now define

$$N = \{x \in N' : |x - z| \leq 3(k+1)\delta \text{ and } |x| \geq (k-1)\delta\}.$$

Since $N \subseteq N'$ is δ -separated and contained in $B(z, 3(k+1)\delta)$, by Proposition 4.3, $|N| \leq \lambda^{O(k)}$.

Consider a point $x \in X_z^{k\delta} \setminus B_X(z, 3k\delta)$. Thus, $(x|z) \geq |z| - k\delta$ and $|x - z| > 3k\delta$; using the last two inequalities we get $|x| - |z| \geq |x - z| - 2k\delta > k\delta$. It follows that a geodesic between r and x must contain a point y' such that $|y'| = |z| + k\delta$. By the δ -hyperbolicity (1),

$$(y'|z) \geq \min\{(y'|x), (z|x)\} - \delta.$$

By definition $(y'|x) = |y'| = |z| + k\delta$, and recalling again that $(z|x) \geq |z| - k\delta$, we have $(y'|z) \geq |z| - (k+1)\delta$. Rearranging the last inequality yields

$$|y' - z| \leq |y'| - |z| + 2(k+1)\delta = (3k+2)\delta.$$

Now, the net N' must contain a point y with $|y - y'| \leq \delta$. Using the triangle inequality, it is easy to verify that $y \in N$, and furthermore

$$\begin{aligned} (x|y) - |y| &= \frac{1}{2}(|x| - |y| - |x - y|) \\ &\geq \frac{1}{2}(|x| - |y'| - |x - y'|) - \delta = -\delta. \end{aligned}$$

We conclude that $x \in X_y^\delta$, which proves the proposition. \square

We can now complete the proof of Lemma 4.4. Let $p > 0$ be a constant to be defined later, and define

$$Z = \{z \in X : |S \cap X_z^\delta| \geq |S|/\lambda^p\}.$$

Let $M = \sup_{z \in Z} |z|$ and observe that $M \leq \max_{x \in S} |x| + \delta$ and thus finite. Let $c \in Z$ be a point with $|c| \geq M - \delta$ (e.g., a point $c \in Z$ with maximum $|c|$, if such a point exists).

It remains to show that $|S \cap X_c^{3\delta}| \leq |S|/2$. To see this, apply Proposition 4.5 to $c \in X$ and $k = 3$ and obtain $N \subseteq X$ with $|N| \leq \lambda^{O(\delta)}$, such that $X_c^{3\delta} \setminus B_X(c, 9\delta) \subseteq \bigcup_{y \in N} X_y^\delta$. The last inequality implies that

$$|S \cap X_c^{3\delta}| \leq |S \cap B_X(c, 9\delta)| + \sum_{y \in N} |S \cap X_y^\delta|.$$

Recalling that S is 20δ -separated, we have $|S \cap B_X(c, 9\delta)| = 1$. For every $y \in N$, we have $|y| \geq |c| + 2\delta > M$, and by the choice of c among Z we get that $|S \cap X_y^\delta| < |S|/\lambda^p$. We conclude that $|S \cap X_c^{3\delta}| \leq 1 + |N| \cdot |S|/\lambda^p$, and the RHS can be upper bounded by $|S|/2$ if the constant p is chosen appropriately. This completes the proof of Lemma 4.4. \square

The algorithm. Let $S \subseteq X$ be an n -point input data set. We may assume without loss of generality that S is (20δ) -separated, as otherwise we could compute a (20δ) -net \hat{S} in S , and execute the entire algorithm only on \hat{S} , which clearly increases the additive approximation by at most 20δ .

We describe a query procedure that solves a more general problem $P(q, S')$: Given a subset $S' \subseteq S$ and a query point $q \in X$, find the point in S' that is closest to q within additive $O(\delta)$. This more general problem will be useful because we will exhibit a recursive algorithm for it. Clearly, the desired query algorithm is just an instantiation of this procedure with $S' = S$.

The algorithm for $P(q, S')$ proceeds as follows. We may assume that $|S'| > 1$, as otherwise the algorithm computes $d(q, S')$ directly. Let $c \in X$ be a point such that

$$|S|/\lambda^{O(\delta)} \leq |S \cap X_c^\delta| \leq |S \cap X_c^{3\delta}| \leq |S|/2.$$

We assume that c was computed at the preprocessing stage by applying Lemma 4.4 to S' . We now have two cases.

Case 1: Suppose $q \in X_c^{2\delta}$. Then use divide and conquer given by:

$$d(q, S') = \min\{d(q, S' \cap X_c^{3\delta}), d(q, S' \setminus X_c^{3\delta})\}.$$

First, the algorithm estimates $d(q, S' \cap X_c^{3\delta})$ within additive error $O(\delta)$, by recursively solving $P(q, S' \cap X_c^{3\delta})$. Second, the algorithm estimates $d(q, S' \setminus X_c^{3\delta})$ by $d(q, c) + d(c, S' \setminus X_c^{3\delta})$. The error in the second estimate is at most 6δ , since for all $x \in S' \setminus X_c^{3\delta}$, Proposition 4.2 gives $|q - x| \leq |q - c| + |c - x| \leq |q - x| + 6\delta$. Note further that this second estimate can be computed in $O(1)$ time, since the summand $d(c, S' \setminus X_c^{3\delta})$ can be calculated already at the preprocessing stage. Finally, the algorithm reports the minimum of the two estimates.

Case 2: Suppose $q \notin X_c^{2\delta}$. Then use divide and conquer given by:

$$d(q, S') = \min\{d(q, S' \cap X_c^\delta), d(q, S' \setminus X_c^\delta)\}.$$

First, the algorithm estimates $d(q, S' \cap X_c^\delta)$ by $d(q, c) + d(c, S' \setminus X_c^\delta)$. The error in this first estimate is at most 4δ , since for all $x \in S' \setminus X_c^\delta$, Proposition 4.2 gives $|q - x| \leq |q - c| + |c - x| \leq |q - x| + 4\delta$. Note further that this second estimate can be computed in $O(1)$ time, since the summand $d(c, S' \cap X_c^\delta)$ can be calculated already at the preprocessing stage. Second, the algorithm estimates $d(q, S' \setminus X_c^\delta)$ within additive error $O(\delta)$, by recursively solving $P(q, S' \setminus X_c^\delta)$. Finally, the algorithm reports the minimum of the two estimates. The algorithm above computes $d(q, S)$ within additive error of 6δ , as can be easily verified from the discussion above by induction on $|S'|$.

Analysis of running time and storage for the simple scheme. The running time of this algorithm is $\lambda^{O(1)} \log n$.

To see this, observe that solving $P(q, S')$ (in either of the two cases) takes $O(1)$ time except for one recursive call in which at least $1/\lambda^{O(1)}$ fraction of the points in S' are eliminated. Furthermore, the algorithm easily extends to also produce a point $a \in S$ whose distance from the query is at most the reported distance. Note, however, that the storage requirement of this algorithm might be exponential in $n = |S|$, since we might have to store data (during the pre-processing stage) for every subset of S (including e.g. a separator point c). This non-trivial problem is overcome in the full version, where we give an algorithm that requires polynomial storage (independent of λ), and that is also intrinsic.

5 The Traveling Salesman Problem

We recall that, given an n -point metric space (X, d) , the *Traveling Salesman Problem (TSP)* on X is to find an ordering (x_1, x_2, \dots, x_n) of the points in X such that the cost $\sum_{i=1}^{n-1} d(x_i, x_{i+1}) + d(x_n, x_1)$ is minimized. Here we outline an approach for approximating the TSP arbitrarily well when X is a subset of some negatively curved space.

In what follows, we refer to a randomized algorithm that computes a $(1 + \varepsilon)$ -approximation to TSP as a $(1 + \varepsilon)$ -TSP algorithm. Unless otherwise stated, *efficient* means polynomial-time. We first state the main result of this section, which follows from Theorem 5.7 below.

Theorem 5.1 (TSP). *There is a PTAS for TSP on any complete, simply-connected Riemannian manifold of pinched negative sectional curvature, including e.g. \mathbb{H}^d for every $d \geq 1$. Furthermore, there is a PTAS for TSP on any bounded-degree δ -hyperbolic graph.*

We define a *tree of metric spaces* as a graph-theoretic tree $T = (V, E)$ together with a family of disjoint metric spaces $\{(X_v, d_v)\}_{v \in V}$, and for every edge $e = \{u, v\} \in E$, a pair of points $p_u^e \in X_u, p_v^e \in X_v$. We define the *induced metric space* (X_T, d_T) by first taking $X' = \bigsqcup_{v \in V} X_v$ to be the disjoint union of $\{X_v\}_{v \in V}$, and then letting X_T be the quotient metric on X' under the equivalence relation

$$x \sim y \iff x = p_u^e, y = p_v^e \text{ for some } e = \{u, v\} \in E.$$

In other words, for every $e = \{u, v\} \in E$, we identify (i.e. glue) the points $p_u^e \in X_u$ and $p_v^e \in X_v$, and then take shortest paths to get a metric d_T on the quotient X_T . We will refer to X_T as a *tree metric over the family* $\{X_v\}_{v \in V}$. We will say that X_T is a τ -neighborhood-tree over (X, d) if X_T is a tree metric over some family \mathcal{F} where each $S \in \mathcal{F}$ is actually a subset $S \subseteq X$ with $\text{diam}(S) \leq \tau$.

We now present a sequence of results whose consequence will be our algorithm for approximating TSP. The proofs are deferred to the full version.

Lemma 5.2. *Given a tree metric X_T over the family $\{X_1, \dots, X_n\}$, there exists a $(1 + \varepsilon)$ -TSP algorithm on X_T running in time $O(n + \sum_{i=1}^n t_i(\varepsilon))$, where $t_i(\varepsilon)$ is an upper bound on the running time of a $(1 + \varepsilon)$ -TSP algorithm for X_i .*

Let (X, d) be a metric space, and consider a countable family of metric spaces $\{(X_i, d_i)\}_{i \in I}$, along with a family of non-contracting maps $\{f_i : X \rightarrow X_i\}_{i \in I}$. If μ is a probability measure on I , then we think of μ as an embedding of X into a distribution over the X_i 's, with the induced distance function d_μ given by

$$d_\mu(x, y) = \mathbb{E}_{i \sim \mu} [d_i(f_i(x), f_i(y))]$$

for $x, y \in X$. If we have $d_\mu(x, y) \leq \alpha \cdot d(x, y)$ for every $x, y \in X$, then we say that μ is an *expanding α -embedding of X into a distribution over the family $\{X_i\}_{i \in I}$* . The next lemma is straightforward.

Lemma 5.3. *If (X, d) admits an efficiently samplable expanding α -embedding into a distribution over the family $\{X_i\}_{i \in I}$, and there exists an efficient β -TSP algorithm for every X_i , then there is an efficient $\alpha\beta$ -TSP algorithm for X .*

Now we come to our main technical theorem.

Theorem 5.4. *Let \mathcal{M} be a geodesic, visual, δ -hyperbolic metric space with local geometry of type (λ_0, δ) . Then for every $\varepsilon > 0$, there exists a value $\tau = \tau(\varepsilon)$, such that for every finite $X \subseteq \mathcal{M}$, there is an efficiently computable, expanding $(1 + \varepsilon)$ -embedding of X into a family $\{X_T\}$ of τ -neighborhood trees over \mathcal{M} .*

Finally, we need a way to solve $(1 + \varepsilon)$ -TSP in τ -neighborhoods of \mathcal{M} . First, there is a quasi-PTAS (i.e. quasi-polynomial running time) for doubling neighborhoods.

Theorem 5.5 ([27]). *If (X, d) is a finite doubling metric space, then there exists a quasi-PTAS for TSP on X (where the running time depends only on $|X|$ and the doubling constant of X).*

Secondly, it is not difficult to see that Arora's algorithm [2] can be made to work on distorted Euclidean spaces.

Theorem 5.6. *If (X, d) is a metric space, and we are given a C -embedding of X into some finite-dimensional Euclidean space \mathbb{R}^k , then there is a PTAS for TSP on X (where the running time depends only on k, C , and $|X|$).*

Combining the preceding collection of reductions yields an algorithm for approximate TSP.

Theorem 5.7. *If \mathcal{M} is as in Theorem 5.4, and in addition \mathcal{M} is locally doubling, then there is a quasi-PTAS for TSP on finite subsets of \mathcal{M} . Additionally, if there exists a value $k \in \mathbb{N}$ such that every ball of radius τ in \mathcal{M} admits a $C(\tau)$ -embedding into \mathbb{R}^k , then there is a PTAS for TSP on finite subsets of \mathcal{M} .*

Acknowledgements. This paper grew out of a seemingly innocent question of Nati Linial during a visit of the authors to the Hebrew University.

References

- [1] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short. Notes on word hyperbolic groups. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 3–63. World Sci. Publishing, River Edge, NJ, 1991. Edited by H. Short.
- [2] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.
- [3] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *J. ACM*, 45(6):891–923, 1998.
- [4] P. Assouad. Plongements lipschitziens dans \mathbb{R}^n . *Bull. Soc. Math. France*, 111(4):429–448, 1983.
- [5] Y. Bartal. Probabilistic approximations of metric space and its algorithmic application. In *37th Annual Symposium on Foundations of Computer Science*, pages 183–193, Oct. 1996.
- [6] E. Begelfor and M. Werman. The world is not always flat or learning curved manifolds. Technical Report 65, Leibniz Center, The Hebrew University of Jerusalem, 2005.
- [7] A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbor. To appear, ICML 2006. Source code at http://hunch.net/~jl/projects/cover_tree/cover_tree.html.
- [8] M. Bonk and O. Schramm. Embeddings of Gromov hyperbolic spaces. *Geom. Funct. Anal.*, 10(2):266–306, 2000.
- [9] S. Buyalo and V. Schroeder. Embedding of hyperbolic spaces in the product of trees. *Geom. Dedicata*, 113:75–93, 2005.
- [10] S. Buyalo and V. Schroeder. A product of trees as universal space for hyperbolic groups. Preprint, arxiv:math.GR/0509355, 2005.
- [11] H. T.-H. Chan, A. Gupta, B. Maggs, and S. Zhou. On hierarchical routing in doubling metrics. In *Proceedings of the 16th Annual Symposium on Discrete Algorithms*, 2005.
- [12] R. Cole and L. Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *Proceedings of the 38th Annual Symposium on the Theory of Computing*, 2006.
- [13] M. Elkin and D. Peleg. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. *SIAM J. Comput.*, 33(3):608–631, 2004.
- [14] É. Ghys and P. de la Harpe, editors. *Sur les groupes hyperboliques d’après Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1990. Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988.
- [15] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [16] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th Annual IEEE Symposium on Foundations of Computer Science*, pages 534–543, Oct. 2003.
- [17] S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics, and their applications. In *Proceedings of the 21st Annual Symposium on Computational Geometry*, 2005.
- [18] D. Karger and M. Ruhl. Finding nearest neighbors in growth-restricted metrics. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pages 63–66, 2002.
- [19] R. Krauthgamer and J. R. Lee. Navigating nets: Simple algorithms for proximity search. In *15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 791–801, Jan. 2004.
- [20] R. Krauthgamer and J. R. Lee. The black-box complexity of nearest-neighbor search. *Theoret. Comput. Sci.*, 348(2-3):262–276, 2005.
- [21] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics. *Geom. Funct. Anal.*, 15(4):839–858, 2005.
- [22] J. M. Lee. *Riemannian manifolds*, volume 176 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997. An introduction to curvature.
- [23] J. R. Lee and A. Naor. Extending Lipschitz functions via random metric partitions. *Invent. Math.*, 160(1):59–95, 2005.
- [24] S. Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proceedings of the 15th Annual Symposium on Computational Geometry*, pages 300–306. ACM, 1999.
- [25] E. Sharon and D. Mumford. 2d-shape analysis using conformal mapping. In *CVPR (2)*, pages 350–357, 2004.
- [26] Y. Shavitt and T. Tankel. On the curvature of the internet and its usage for overlay construction and distance estimation. In *Proceedings of the 23th Annual INFOCOM*, 2004.
- [27] K. Talwar. Bypassing the embedding: algorithms for low dimensional metrics. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 281–290, New York, 2004. ACM.