Inapproximability for metric embeddings into $\mathbb{R}^d$

Jiří Matoušek
Department of Applied Mathematics and
Institute of Theoretical Computer Science (ITI)
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic

Anastasios Sidiropoulos
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
32 Vassar St, Room 32G-608, Cambridge, MA, USA

Abstract

We consider the problem of computing the smallest possible distortion for embedding of a given n-point metric space into $\mathbb{R}^d$, where d is fixed *(and small). For $d = 1$, it was known that approximating the minimum distortion with a factor better than roughly $n^{1/12}$ is NP-hard. From this result we derive inapproximability with factor roughly $n^{1/(22d-10)}$ for every fixed $d \geq 2$, by a conceptually very simple reduction. However, the proof of correctness involves a nontrivial result in geometric topology (whose current proof is based on ideas due to Jussi Väisälä).*

For $d \geq 3$, we obtain a stronger inapproximability result by a different reduction: assuming P≠NP, no polynomial-time algorithm can distinguish between spaces embeddable in $\mathbb{R}^d$ with constant distortion from spaces requiring distortion at least $n^{c/d}$, for a constant $c > 0$. The exponent c/d has the correct order of magnitude, since every n-point metric space can be embedded in $\mathbb{R}^d$ with distortion $O(n^{2/d}\log^{3/2} n)$ and such an embedding can be constructed in polynomial time by random projection.

For $d = 2$, we give an example of a metric space that requires a large distortion for embedding in $\mathbb{R}^2$, while all not too large subspaces of it embed almost isometrically.

1 Introduction

Let $\mathbb{X} = (X, \rho_X)$ and $\mathbb{Y} = (Y, \rho_Y)$ be metric spaces and let $f: X \to Y$ be an injective mapping (embedding). The *distortion* of f, denoted by $\operatorname{dist}(f)$, is the smallest $D \geq 1$ such that there exists $\alpha > 0$ (a scaling factor) for which $\alpha\rho_X(x, y) \leq \rho_Y(f(x), f(y)) \leq D\alpha\rho_X(x, y)$ for all $x, y \in X$. An embedding with distortion at most D is also called a *D-embedding*. We let $c_{\mathbb{Y}}(\mathbb{X})$ denote the infimum of all $D \geq 1$ such that $\mathbb{X}$ admits a D-embedding into $\mathbb{Y}$.

We will also use the symbol $\Delta(\mathbb{X})$ for the *aspect ratio* of a finite metric space $\mathbb{X}$, which is defined as the largest distance in $\mathbb{X}$ divided by the smallest nonzero distance in $\mathbb{X}$.

Over the past few decades, metric embeddings have resulted in some of the most beautiful and powerful algorithmic techniques, with applications in many areas of computer science [19, 14]. In most of these results, low-distortion embeddings provide a way to simplify the data, without losing too much information.

Here we focus on embeddings of finite metric spaces $\mathbb{X}$ into $\mathbb{R}^d$ with the Euclidean metric $\|.\|$, where d is a *fixed* integer. More precisely, we mainly consider the algorithmic problem of computing or estimating $c_{\mathbb{R}^d}(\mathbb{X})$ for a given n-point metric space $\mathbb{X}$. For $d \leq 3$, this problem has an immediate application to visualizing finite metric spaces.

It is known that every n-point metric space $\mathbb{X}$ embeds in $\mathbb{R}^d$ with distortion at most $O(n^{2/d}\log^{3/2} n)$ [20]. The proof is constructive and it yields an $O(n^{2/d}\log^{3/2} n)$-approximation algorithm for $c_{\mathbb{R}^d}(\mathbb{X})$, and as far as we know, this is the best known approximation algorithm for this problem.

Here we will show that, assuming P≠NP, there is no polynomial-time algorithm with approximation ratio *much* better than $n^{2/d}$. Namely, we will prove that $c_{\mathbb{R}^d}(\mathbb{X})$ cannot be approximated with factor smaller than $n^{c/d}$ for a universal constant c (so at least the exponent c/d has the correct order of magnitude as a function of d). We now state the results more precisely.

All dimensions hard . . . Bădoiu et al. [6] proved that it is NP-hard to approximate the minimum distortion required to embed a given n-point metric space $\mathbb{X}$ into $\mathbb{R}^1$ with factor better than roughly $n^{1/12}$ (see Theorem 3.1 below for a precise formulation). Using their result as a black box, we obtain an analogous hardness result for embeddings in $\mathbb{R}^d$ for every fixed $d \geq 2$:

Theorem 1.1 *For every fixed $d \geq 2$, and for every fixed $\varepsilon > 0$, it is NP-hard to approximate the minimum distortion required for embedding of a given n-point metric space into $\mathbb{R}^d$ within a factor of $\Omega(n^{1/(22d-10)-\varepsilon})$.*

Our derivation of this theorem from the 1-dimensional result is conceptually very simple: Given an n-point metric space $\mathbb{X}$, we consider a $(d-1)$-dimensional sphere S in $\mathbb{R}^d$ of radius R much larger than the largest distance in $\mathbb{X}$, and in this S we pick an ε-dense[1] finite set V for a sufficiently

[1] A set V in a metric space (X, ρ_X) is called ε-dense if for each $x \in X$ there is $v \in V$ with $\rho_X(x, v) \leq \varepsilon$.

Figure 1: A D-embedding of $\mathbb{Y}$ (a schematic illustration for $d = 2$ and $|X| = 3$).

small $\varepsilon > 0$. Then we form another metric space $\mathbb{Y} = (Y, \rho_Y)$ as a suitable Cartesian product of $\mathbb{X}$ with $(V, \|.\|)$.

It is easy to show that $c_{\mathbb{R}^d}(\mathbb{Y}) = O(c_{\mathbb{R}^1}(\mathbb{X}))$. The harder part is extracting an $O(D)$-embedding of $\mathbb{X}$ into $\mathbb{R}^1$ from an arbitrary given D-embedding $g: Y \to \mathbb{R}^d$; see Fig. 1. Intuitively, the image of each copy of V in $\mathbb{Y}$ has to "look like" a deformed sphere, and these "deformed spheres" all have to be nested. Hence they are linearly ordered, and this provides an ordering of the points of $\mathbb{X}$ in a sequence, say $(a_1, a_2, \ldots, a_n)$. Then we define an embedding $f: X \to \mathbb{R}^1$ so that $f(a_1) < f(a_2) < \cdots < f(a_n)$, and the difference $f(a_{i+1}) - f(a_i)$ is the distance of the $(i+1)$st "deformed sphere" from the ith one.

The claim about the nesting of the "deformed spheres" may seem intuitively obvious, but apparently it is not entirely easy to prove, and for establishing it rigorously we will apply some tools from analysis and from algebraic topology (Section 2); part of the current proof is due to Väisälä [26]. This result and some by-products of the proof can be of independent interest. Then we prove Theorem 1.1 in Section 3 along the lines just indicated.

…3 and more dimensions harder? Theorem 1.1 shows that for embeddability in $\mathbb{R}^d$ it is hard to distinguish bad spaces from even much worse ones. However, for applications of low-distortion embeddings, one is usually most interested in efficiently distinguishing good spaces (embeddable with a constant distortion, say) from bad ones.

For example, Theorem 1.1 leaves open the possibility of a polynomial-time algorithm that, given a metric space $\mathbb{X}$, constructs an embedding of $\mathbb{X}$ into $\mathbb{R}^d$ with distortion bounded by a polynomial in the optimal distortion $c_{\mathbb{R}^d}(\mathbb{X})$. For $d = 1$, there are indeed partial results of this kind for restricted classes of metrics, namely, for weighted trees [6] and for unweighted graphs [8].[2] Thus, at least for these two classes, good and bad embeddability in $\mathbb{R}^1$ can be distinguished efficiently (although in a somewhat weak sense).

[2] By a *unweighted graph* we mean a metric space whose point set is the vertex set of a graph G and whose metric is the shortest-path metric of G (where each edge has length 1). Similarly, the metric of a *weighted tree* is the shortest-path metric of some tree, where the edges may have arbitrary nonnegative lengths.

Our next result shows that for $d \geq 3$, even this kind of distinguishing good from bad is hard in general:

Theorem 1.2 *For every fixed $d \geq 3$, it is NP-hard to distinguish between n-point metric spaces that embed in $\mathbb{R}^d$ with distortion at most D_0, and ones that require distortion at least $n^{c/d}$, where $c > 0$ is a universal constant and D_0 is a constant depending on d.*

Before proving this result, we first establish a weaker but simpler one in Section 5. The tools developed in this simpler proof also appear in the proof of Theorem 1.2 in Section 6.

The techniques used in the proof of Theorem 1.2 do not seem to be applicable for the case of embedding into $\mathbb{R}^1$ or $\mathbb{R}^2$. So for $d = 1$ or $d = 2$, it is still possible that there exists a polynomial-time algorithm that computes an embedding of a given metric space $\mathbb{X}$ into $\mathbb{R}^d$ with distortion at most $c_{\mathbb{R}^d}(\mathbb{X})^{O(1)}$.

No Menger-type condition for approximate embeddings into the plane. While we cannot exclude the existence of an efficient algorithm that distinguishes "good spaces from bad ones" for embeddings in $\mathbb{R}^2$, we provide some evidence that obtaining such an algorithm may not be easy, since there is no "local" characterization of good embeddability.

First we recall a well-known lemma of Menger [21], asserting that an n-point metric space $\mathbb{X}$ embeds *isometrically* in $\mathbb{R}^d$ if (and only if) every subspace of $\mathbb{X}$ on at most $d+3$ points so embeds. In contract to this, we have the following result:

Theorem 1.3 *Let $\varepsilon \in (0, 1)$ be given, let n be sufficiently large, and let $1/\sqrt{\varepsilon} \leq k \leq c\sqrt{\varepsilon}n$, where c is a sufficiently small constant. Then there exists an n-point metric space $\mathbb{X}$, whose embedding in $\mathbb{R}^2$ requires distortion $\Omega(\sqrt{\varepsilon}\, n/k)$, while every k-point subspace can be embedded in $\mathbb{R}^2$ with distortion at most $1 + \varepsilon$. (Proof omitted.)*

We remark that similar questions have been studied for embeddings in ℓ_1 by Arora et al. [1] and by Charikar et al. [9].

For additional references and summary of work related to our results, as well as for complete proofs, we refer to a full version of this paper, which can be accessed on-line at `http://theory.csail.mit.edu/~tasos/hardness-full.pdf`

2 Deformed spheres and nesting lemmas

As was outlined in the introduction, in the proof of Theorem 1.1 we will be confronted with the following setting: We have a finite set V in a $(d-1)$-dimensional sphere S; for the purposes of this section we may assume that $S = S^{d-1}$ is the unit sphere in $\mathbb{R}^d$. We assume that V is ε-dense in S^{d-1}, and we are given a D-embedding $g: V \to \mathbb{R}^d$. By re-scaling we may assume that g is noncontracting and D-Lipschitz.

In order to employ topological reasoning about the image of such g, we extend g to a continuous map $\overline{g}: S^{d-1} \to \mathbb{R}^d$ by a suitable interpolation (a tool for doing this will be mentioned later in this Section. We can make sure that $\overline{g}$ is still D-Lipschitz, but generally it won't be noncontracting, and it can even fail to be injective.

However, $\overline{g}$ satisfies the following weaker version of "noncontracting": For all $x, y \in S^{d-1}$ we have $\|\overline{g}(x) - \overline{g}(y)\| \geq \|x - y\| - \delta$, where $\delta = 2D\varepsilon$ (this is easy to check; see Lemma 3.3 below).

The main goal of this section is to show that the image of such $\overline{g}$ behaves, in a suitable sense, as an "approximate sphere". This is expressed in Theorem 2.1 below.

For the proof of Theorem 1.2 we will need a technical extension of these results; namely, instead of images of S^{d-1}, we need to deal with images of more general shapes, e.g., long tubes and punctured spheres. This is done in Section 4.

Big holes and nested spheres. For a compact set $K \subset \mathbb{R}^d$, let us call a bounded component of $\mathbb{R}^d \setminus K$ a *hole* of K.

Theorem 2.1

(i) (A big hole exists) *Let $\delta \in [0, \frac{1}{4})$, let $f: S^{d-1} \to \mathbb{R}^d$ be a continuous map that satisfies $\|f(x) - f(y)\| \geq \|x-y\| - \delta$ for all $x, y \in S^{d-1}$, and let $\Sigma := f(S^{d-1})$. Then Σ has a hole containing a ball of radius $\frac{1}{4}$.*

(ii) *Let $f_1, f_2 : S^{d-1} \to \mathbb{R}^d$ be maps satisfying the condition as in (i) for some $\delta < \frac{1}{4}$, and suppose that, moreover, $\|f_1(x) - f_2(x)\| < \frac{1}{4}$ for all $x \in S^{d-1}$. Then some hole of $\Sigma_1 := f_1(S^{d-1})$ intersects some hole of $\Sigma_2 := f_2(S^{d-1})$.*

(iii) (All holes but one are narrow) *Let δ, f, and Σ be as in (i), and let us assume that, moreover, f is D-Lipschitz for some $D \geq 1$. Then there is at most one hole of Σ containing a ball of radius $4D\delta$.*

Part (ii) is what we will need, part (i) can be regarded as a by-product of the proof, and part (iii) we do not need but it comes almost for free and it completes the picture. The main ideas of the proof of (i) and (iii) as given below were found by Väisälä [26] in an answer to a question of the first author, and here they are used with his kind permission (we have independently found another proof, but since it was much less elegant, we reproduce Väisälä's).

Here is the result we need for the proof of Theorem 1.1:

Corollary 2.2 (Nesting lemma) *Let $\delta < \frac{1}{4}$ and let $f_1, f_2, \ldots, f_n: S^{d-1} \to \mathbb{R}^d$ be continuous maps satisfying*

- *$\|f_i(x) - f_i(y)\| \geq \|x - y\| - \delta$ for all $x, y \in S^{d-1}$ and all i,*
- *$\|f_i(x) - f_j(x)\| \leq \frac{1}{4}$ for all i, j and all $x \in S^{d-1}$, and*
- *$\Sigma_i \cap \Sigma_j = \emptyset$ whenever $i \neq j$, where $\Sigma_i = f_i(S^{d-1})$.*

Let U_i denote the unbounded component of $\mathbb{R}^d \setminus \Sigma_i$, and let us define a relation $\preceq$ on $[n]$ by setting $i \preceq j$ if $U_j \subseteq U_i$. Then $\preceq$ is a linear ordering on $[n]$. (Proof omitted.)

A lemma on approximate inverse. The first main step in the proof of Theorem 2.1 is the next lemma, which says that f has an "approximate inverse" mapping h that extends to some neighborhood of Σ.

Lemma 2.3 *Let f, Σ, and $\delta < \frac{1}{4}$ be as in Theorem 2.1(i), and let Ω_r denote the closed r-neighborhood of Σ in $\mathbb{R}^d$. Then there is a continuous map $h: \Omega_{1/4} \to S^{d-1}$ such that $\|h(f(x)) - x\| \leq 8\delta$ for all $x \in S^{d-1}$, and (consequently) the composition $hf: S^{d-1} \to S^{d-1}$ is homotopic[3] to the identity map $\mathrm{id}_{S^{d-1}}$.*

In the proof we will use a basic result about Lipschitz maps: the Kirszbraun theorem [17], which asserts that every Lipschitz mapping from a subset of a Hilbert space H_1 into a Hilbert space H_2 can be extended to a Lipschitz map $H_1 \to H_2$, with the same Lipschitz constant.

Proof of Lemma 2.3. Let us put $\varepsilon := \delta$ and $r := \frac{1}{4}$. Let $N \subset \Sigma$ be an ε-net[4] in Σ. We choose a mapping $g: N \to S^{d-1}$ with $fg = \mathrm{id}_N$; in other words, for every $y \in N$ we arbitrarily choose $g(y) \in f^{-1}(y)$.

We claim that g is 2-Lipschitz. Indeed, if $y, y' \in N$ are distinct and $x = g(y)$, $x' = g(y')$, then the condition on f gives $\|x - x'\| \leq \|f(x) - f(x')\| + \delta = \|y - y'\| + \delta < (1 + \delta/\varepsilon)\|y - y'\| = 2\|y - y'\|$ since $\|y - y'\| > \varepsilon$.

[3] We recall that two continuous maps $f, g: X \to Y$ of topological spaces are *homotopic*, in symbols $f \sim g$, if there exists a continuous map $F: X \times [0,1] \to Y$ such that $F(x, 0) = f(x)$ and $F(x, 1) = g(x)$ for all $x \in X$.

[4] We recall that a subset $N \subseteq M$ in a metric space (M, ρ_M) is an *ε-net* if every two points of N have distance greater than ε and N is inclusion-maximal with respect to this property; that is, every point of M is at most ε far from some point of N.

Next, using the Kirszbraun theorem mentioned above, we extend g to a 2-Lipschitz map $\overline{g}: \Omega_r \to \mathbb{R}^d$. We check that 0 is not in the image of $\overline{g}$; indeed, if we had $\overline{g}(y) = 0$ for some $y \in \Omega_r$, we could find a point $z \in N$ at distance at most $r + \varepsilon$ from y, hence $\|\overline{g}(z) - \overline{g}(y)\| \leq 2(r + \varepsilon) < 1$ (using $r = \frac{1}{4}$, $\varepsilon = \delta < \frac{1}{4}$), but $\overline{g}(y) = 0$ while $\|\overline{g}(z)\| = 1$ since $\overline{g}(z) = g(z) \in S^{d-1}$.

We can now define the desired $h: \Omega \to S^{d-1}$ as in the lemma, by $h(y) = \overline{g}(y)/\|\overline{g}(y)\|$.

Given $x \in S^{d-1}$, we pick $z \in N$ at most ε away from $f(x)$, and we calculate $\|\overline{g}(f(x)) - x\| \leq \|\overline{g}(f(x)) - g(z)\| + \|g(z) - x\| \leq 2\|z - f(x)\| + \|z - f(x)\| + \delta \leq 3\varepsilon + \delta = 4\delta$. Since $\|\overline{g}(f(x)) - h(f(x))\| = 1 - \|\overline{g}(f(x))\| \leq \|x - \overline{g}(f(x))\|$, we obtain $\|h(f(x)) - x\| \leq 2\|\overline{g}(f(x)) - x\| \leq 8\delta$ as claimed.

For $\delta < \frac{1}{4}$, this implies that $h(f(x)) \neq -x$ for all $x \in S^{d-1}$, and consequently, $hf \sim \mathrm{id}_{S^{d-1}}$ (this is a standard and easy fact in topology; if x and $h(f(x))$ are not antipodal, they are connected by a unique shortest arc, and the homotopy moves along this arc). The lemma is proved. □

The Alexander duality. In the subsequent proof of Theorem 2.1, we will use cohomology groups. We will not need their definition, only few very simple properties, which we will explicitly state, plus one slightly deeper result of algebraic topology. These can be taken as purely formal rules, which we will apply in the proof. We consider $(d-1)$-dimensional cohomology, since it is closely related to the number of holes.

Each compact set $X \subset \mathbb{R}^d$ is assigned the $(d-1)$-dimensional Čech (or equivalently, Alexander–Spanier) cohomology group[5] $\check{H}^{d-1}(X)$; for definiteness we consider integer coefficients, although the coefficient ring doesn't matter in our considerations. This $\check{H}^{d-1}(X)$ is an Abelian group, and if it is finitely generated, then it is isomorphic to $\mathbb{Z}^b$ for an integer $b \geq 0$, called the *rank* of $\check{H}^{d-1}(X)$.

A very rough intuition is that the elements of $\check{H}^{d-1}(X)$ correspond to (equivalence classes of) $(d-1)$-dimensional "surfaces" inside X, with nonzero elements corresponding to "surfaces" that "enclose" one or more of the holes of X. (This is really closer to the idea of homology, rather than cohomology, but hopefully it is not totally misleading for our purposes.)

A special case of the *Alexander duality*, which we will state precisely in the proof of Lemma 2.4 below, tells us that the rank of $\check{H}^{d-1}(X)$ equals the number of holes of X. For example, S^{d-1} encloses a single hole, and we have $\check{H}^{d-1}(S^{d-1}) \cong \mathbb{Z}$.

[5]We need Čech cohomology so that our considerations are valid even for X with various local pathologies. In our application of Theorem 2.1 we can assume that the mapping f is "nice", e.g., that its image Σ is the union of finitely many simplices, and then we could work with the perhaps more familiar singular or simplicial cohomology.

A continuous map $f: X \to Y$ of compact sets induces a group homomorphism $f^*: \check{H}^{d-1}(Y) \to \check{H}^{d-1}(X)$; we should stress that f^* goes in opposite direction compared to f. For the composition of maps we then have $(fg)^* = g^* f^*$ (the last two properties are usually expressed by saying that cohomology is a contravariant functor). Moreover, if $f_1, f_2: X \to Y$ are homotopic maps, then $f_1^* = f_2^*$.

The following lemma encapsulates what we will need from the Alexander duality.

Lemma 2.4 (($d-1$)-dimensional cohomology and holes)

(i) *Let $d \geq 2$, let $X \subseteq Y$ be compact sets in $\mathbb{R}^d$, let $j: X \to Y$ denote the inclusion map, and let $j^*: \check{H}^{d-1}(Y) \to \check{H}^{d-1}(X)$ be the induced homomorphism in cohomology. Then the number of holes of X that contain at least one hole of Y equals the rank of the image $\operatorname{Im} j^*$.*

(ii) *Let $d \geq 2$, let X_1, X_2, Y be compact sets in $\mathbb{R}^d$, $X_1 \subseteq Y$, $X_2 \subseteq Y$, let j_1, j_2 be the inclusion maps and j_1^*, j_2^* the induced homomorphisms in cohomology. Suppose that $\operatorname{Ker}(j_1^*) \cup \operatorname{Ker}(j_2^*)$ does not generate all of $\check{H}^{d-1}(Y)$. Then there is a hole of Y contained both in a hole of X_1 and in a hole of X_2.*

(Proof omitted.)

Proof of Theorem 2.1. Let us consider the map f as in part (i), and Ω_r and h as in Lemma 2.3. Let $j: \Sigma \to \Omega_{1/4}$ denote the inclusion map. The composed map $hf = hjf: S^{d-1} \to S^{d-1}$ is homotopic to the identity, and so the induced map $f^* j^* h^*: \check{H}^{d-1}(S^{d-1}) \to \check{H}^{d-1}(S^{d-1})$ in cohomology is the identity as well. Since $\check{H}^{d-1}(S^{d-1}) \neq 0$, the homomorphism $j^*: \check{H}^{d-1}(\Omega_{1/4}) \to \check{H}^{d-1}(\Sigma)$ cannot be zero. By Lemma 2.4(i) this means that there is a hole of Σ that contains a hole of $\Omega_{1/4}$, and such a hole of Σ contains a $\frac{1}{4}$-ball.

In part (ii), let Ω be the $\frac{1}{4}$-neighborhood of Σ_1, let $j_1: \Sigma_1 \to \Omega$ be the inclusion map, and let $h_1: \Omega \to S^{d-1}$ be as in the proof of (i), i.e., with $h_1 f_1 \sim \mathrm{id}_{S^{d-1}}$. By the assumption $\Sigma_2 \subseteq \Omega$ as well (with inclusion map j_2), and f_1 and f_2 are homotopic as maps $S^{d-1} \to \Omega$, since the segment $f_1(x)f_2(x)$ is contained in Ω for every $x \in S^{d-1}$. So the homomorphisms $f_1^* j_1^*$ and $f_2^* j_2^*$ in cohomology are equal, and also nonzero, since $f_1^* j_1^* h_1^*$ is the identity in $\check{H}^{d-1}(S^{d-1})$.

The kernels of j_1^* and j_2^* are both contained in $\operatorname{Ker}(f_1^* j_1^*) = \operatorname{Ker}(f_2^* j_2^*)$, and the latter is not all of $\check{H}^{d-1}(\Omega)$. Thus, Lemma 2.4(ii) implies that there is a hole of Ω contained both in a hole of Σ_1 and in a hole of Σ_2, and part (ii) is proved. Proof of part (iii) is omitted. □

3 Hardness for $\mathbb{R}^1$ implies hardness for $\mathbb{R}^d$

As was mentioned in the introduction, we derive Theorem 1.1 from the result of Bădoiu et al. [6] on inapproximability for embeddings into $\mathbb{R}^1$. By inspecting the full version of that paper (available on-line), one can check that their proof yields the following:

Theorem 3.1 (Bădoiu et al. [6]) *Assuming P $\neq$ NP, there is no polynomial-time algorithm with the following three properties: (i) The input of the algorithm is an n-point metric space $\mathbb{X}$ with $\Delta(\mathbb{X}) = O(n)$. (ii) If $\mathbb{X}$ admits an $O(n^{4/12})$-embedding into $\mathbb{R}^1$, the algorithm answers* YES. *(iii) If $\mathbb{X}$ is not embeddable in $\mathbb{R}^1$ with distortion smaller than $\Omega(n^{5/12-\varepsilon})$, the algorithm answers* NO.

This theorem together with the next proposition imply Theorem 1.1 by a simple calculation.

Proposition 3.2 *Let $\mathbb{X} = (X, \rho_X)$ be an n-point metric space, let $d \geq 2$ be a fixed integer, and let $D_{\max} \geq 1$ be a parameter (specifying the maximum distortions we want to consider). There exists a metric space $\mathbb{Y} = (Y, \rho_Y)$, $|Y| = O(nD_{\max}^{2(d-1)}\Delta(\mathbb{X})^{d-1})$, which can be constructed in time polynomial in n, $\Delta(\mathbb{X})$, and $D_{\max}$ (the implicit constants depending on d), with the following properties:*

(i) *If $\mathbb{X}$ can be D-embedded in $\mathbb{R}^1$ for some $D \geq 1$, then $\mathbb{Y}$ can be $(1.1D)$-embedded[6] in $\mathbb{R}^d$.*

(ii) *Given a D-embedding of $\mathbb{Y}$ in $\mathbb{R}^d$ for some D, $1 \leq D \leq D_{\max}$, one can construct a $1.1D$-embedding of $\mathbb{X}$ in $\mathbb{R}^1$ in polynomial time.*

The construction. We follow the sketch given after Theorem 1.1. Let us assume that the smallest distance in $\mathbb{X}$ is 1 and the largest one is Δ. We let $C = C(d)$ be a sufficiently large constant, we set $R := CD_{\max}\Delta$, and we let S be the $(d-1)$-dimensional sphere in $\mathbb{R}^d$ centered at 0 of radius R. We set $\varepsilon := \frac{1}{CD_{\max}}$, and we choose V as an ε-dense subset of S (that is, each point of S has distance at most ε to some point of V); as is well known, we can assume that V has size $O((R/\varepsilon)^{d-1})$ and is computable in time polynomial in R/ε.

Then we let $\mathbb{Y} = (Y, \rho_Y) := \mathbb{X} \times_{L_2} V$; that is, we set $Y := X \times V$, and we define the metric ρ_Y by $\rho_Y((a, v), (a', v')) = \sqrt{\rho_X(a, a')^2 + \|v - v'\|^2}$.

Checking part(i) of Proposition 3.2 is easy and we omit it.

For part (ii), let $g\colon Y \to \mathbb{R}^d$ be a D-embedding; for convenience, we assume that it is noncontracting and D-Lipschitz. For each $a \in X$ we consider the "slice" of g, i.e., the mapping $g_a\colon V \to \mathbb{R}^d$ given by $g_a(v) = g(a, v)$.

Next, we extend each g_a to a D-Lipschitz map $\overline{g}_a\colon S \to \mathbb{R}^d$ using the Kirszbraun theorem (mentioned after Lemma 2.3).[7] Let $\Sigma_a := \overline{g}_a(S)$ be the image of $\overline{g}_a$.

Now we want to use the nesting lemma (Corollary 2.2) to show that the Σ_a have to be nested. More precisely, we want to check that by scaling both the domain and range of each $\overline{g}_a$ by $\frac{1}{R}$, we obtain maps $S^{d-1} \to \mathbb{R}^d$ as in Corollary 2.2. This is done using the next two lemmas.

Lemma 3.3 *For all $x, y \in S$ we have $\|\overline{g}_a(x) - \overline{g}_a(y)\| \geq \|x - y\| - \frac{3}{C}$. (Proof omitted.)*

Lemma 3.4

(i) *For $a \neq b$, the Euclidean distance of Σ_a and Σ_b is at least $\rho_X(a, b) - \frac{2}{C}$, and in particular, $\Sigma_a \cap \Sigma_b = \emptyset$.*

(ii) *For any $a, b \in X$ and $x \in S$, we have $\|\overline{g}_a(x) - \overline{g}_b(x)\| \leq 2D_{\max}\Delta$.*

(Proof omitted.)

Proof of Proposition 3.2(ii). As was announced above, we can now apply Corollary 2.2 to the maps $\overline{g}_a$ with domain and range rescaled by $\frac{1}{R}$ (using $\delta = 3/CR < \frac{1}{4}$ and $2D_{\max}\Delta/R \leq \frac{1}{4}$). Letting U_a denote the unbounded component of $\mathbb{R}^d \setminus \Sigma_a$, we can number the points of X as $a_1, \ldots, a_n$ so that for $i < j$ we have $U_{a_i} \supset U_{a_j}$.

For $i = 1, 2, \ldots, n-1$ we define δ_i as the (Euclidean) distance of Σ_{a_i} from $\Sigma_{a_{i+1}}$, and we define a mapping $f\colon X \to \mathbb{R}^1$ by $f(a_i) = \sum_{j=1}^{i-1} \delta_j$.

Assuming that the original mapping g has distortion at most D, we will prove that f has distortion at most $1.1D$. First we show that f contracts distances by a factor of at most 1.1. Lemma 3.4(i) gives $\delta_i \geq \rho_X(a_i, a_{i+1}) - 2/C \geq \rho_X(a_i, a_{i+1})/1.1$ (assuming C large). The triangle inequality then shows that $|f(a_i) - f(a_j)| \geq \rho_X(a_i, a_j)/1.1$ for all i, j.

Next, we want to bound the Lipschitz constant of f. Let us fix a point $v_0 \in V$ and let us abbreviate $x_i := g(a_i, v_0)$. Let us fix $i < j$ and let us consider the line segment $x_i x_j$. We note that whenever k lies between i and j, the segment $x_i x_j$ intersects Σ_{a_k}. This is because $\Sigma_{a_j} \subset U_{a_k}$, while $\Sigma_{a_i} \cap U_{a_k} = \emptyset$. Thus, for each k, $i \leq k \leq j$, we can fix a point $y_k \in \Sigma_{a_k}$ on $x_i x_j$, where $y_i = x_i$ and $y_j = x_j$ (we note that y_k also depends on i and j). Then

$$\begin{aligned} |f(a_j) - f(a_i)| &= \sum_{k=i}^{j-1} \delta_k \leq \sum_{k=i}^{j-1} \|y_{k+1} - y_k\| \\ &= \|x_i - x_j\| \leq D\rho_Y((a_i, v_0), (a_j, v_0)) \\ &= D\rho_X(a_i, a_j) \end{aligned}$$

[6] If needed, we could replace 1.1 by any other constant greater than 1, with appropriate adjustments in other constants. We use 1.1 so that we need not introduce an extra parameter.

[7] Another way of extending the g_a is to assume that V is a vertex set of some fine enough triangulation of S^{d-1}, and extend affinely on each simplex of the triangulation. In this way we have more control about the local properties of the image (which is piecewise linear), but we need to worry about the existence of a suitable triangulation.

since g is D-Lipschitz.

It remains to show how f can be found from g in polynomial time. First we need to sort the Σ_a. To compare Σ_a and Σ_b, we can compute a point with the minimum x_1-coordinate, say, of $\Sigma_a \cup \Sigma_b$ and see if it lies in Σ_a or Σ_b (here we can use a property which follows from the proof of the Kirszbraun theorem, namely, that we may assume $\overline{g}_a(S) \subseteq \mathrm{conv}(g_a(V))$, which implies that the smallest point of Σ_a lies in $g_a(V)$). Then we can approximate the distance of Σ_a to Σ_b by the distance of the finite sets $g_a(V)$ and $g_b(V)$; this causes a small additive error which can increase the distortion of f only negligibly. This concludes the proof of Proposition 3.2. □

4 Punctured pseudospheres

For the stronger inapproximability result for dimensions 3 and higher, Theorem 1.2, we will need a nesting property not only for images of dense sets in spheres, but also for images for dense sets in other shapes.

In this section we develop a version of the nesting lemma that covers all of our applications. The definitions are tailored to these applications. In order to reduce the number of parameters, we use the same bound ε for several independent small quantities; if we were aiming at tighter bounds in the inapproximability results, we could fine-tune each of these quantities independently.

Let $S \subseteq \mathbb{R}^d \setminus \{0\}$ be a set, and let $\varepsilon > 0$. We call a set $V \subseteq S$ *ε-angularly dense* in S if for every $x \in S$ there exists $v \in V$ with $\|x - v\| \le \varepsilon\|v\|$.

We call a set $P \subseteq \mathbb{R}^d$ *ε-angularly small with respect to* a set $V \subset \mathbb{R}^d \setminus \{0\}$ if there is a choice of a radius $r_v \ge 0$ for every $v \in V$ such that $P \subseteq \bigcup_{v\in V} B(v, r_v)$ (where $B(x, r)$ denotes the ball of radius r centered at x) and $\sum_{v\in V} \frac{r_v}{\|v\|} \le \varepsilon$ (this is a wasteful definition; aiming at more precise quantitative results, we would take $(r_v/\|v\|)^{d-1}$ instead of $r_v/\|v\|$, for example).

For our purposes, a *pseudosphere* is a set $S \subset \mathbb{R}^d$ homeomorphic to an S^{d-1} such that

- the hole of S contains 0, and
- there is a retraction r_S of $\mathbb{R}^d\setminus\{0\}$ onto S (i.e., $r_S\colon \mathbb{R}^d\setminus\{0\} \to S$ is a continuous map whose restriction on S is the identity map).

A *punctured pseudosphere* is a pair (S, P), where S is a pseudosphere and $P \subseteq S$, the "punctures" of the pseudosphere, is a subset of S, which we will assume to be small in a suitable sense.

Proposition 4.1 (Nesting lemma for punctured pseudospheres) *Let $d \ge 2$, let $D \ge 1$ and let $\varepsilon := \frac{1}{16D}$, let (S, P) be a punctured pseudosphere in $\mathbb{R}^d$, let $V \subset S$ be an ε-angularly dense set in S, let us assume that $P \subseteq S$ is ε-angularly small w.r.t. V, that $P \cap V = \emptyset$, and that $S \setminus P$ is path-connected, and let $f_1, f_2, \ldots, f_n\colon S \to \mathbb{R}^d$ be maps such that:*

- *Each f_i is D-Lipschitz.*
- *Each f_i restricted to V is noncontracting.*
- *We have $\|f_i(v) - f_j(v)\| \le \frac{1}{4}\|v\|$ for all $v \in V$ and all i, j.*
- *Setting $\Sigma_i := f_i(S)$ and $\Sigma_i^* := f_i(S \setminus P)$, we have $\Sigma_i \cap \Sigma_j^* = \emptyset$ for all $i \ne j$.*

Let U_i denote the unbounded component of $\mathbb{R}^d \setminus \Sigma_i$, and let us define a relation $\preceq$ on $[n]$ by setting $i \preceq j$ if either $i = j$ or $\Sigma_j^* \subset U_i$. Then $\preceq$ is a linear ordering on $[n]$.

Moreover, $\preceq$ is independent of the behavior of the f_i on the punctures, in the following sense: If $\tilde{f}_1, \ldots, \tilde{f}_i$ are D-Lipschitz mappings $S \to \mathbb{R}^d$ such that $f_i(x) = \tilde{f}_i(x)$ for all $x \in S \setminus P$ and all i (in particular, $\Sigma_i^* = f_i(S \setminus P) = \tilde{f}_i(S \setminus P)$), and $\Sigma_i^* \cap \tilde{f}_j(S) = \emptyset$ for all $i \ne j$, then the linear ordering induced by the $\tilde{f}_i$ is the same as $\preceq$. (Proof omitted.)

A basic example. Since this proposition is rather technical, let us present a basic example of a setting in which it will be applied. Let C be a long cylinder in $\mathbb{R}^d$ of a large radius R, and let V be a set that is ε-dense in the lateral surface of C. With this V we make a construction similar to the one in the proof of Theorem 1.1 above. We set $Y = [n] \times V$, and we define a metric on Y by $\rho_Y((i, v), (i', v')) = \|v - v'\| + \delta_{ii'}$, where $\delta_{ii'}$ is the Kronecker delta (equal to 0 for $i = i'$ and to 1 otherwise).

We assume $\varepsilon \ll 1 \ll R$, and so we expect that if $g\colon Y \to \mathbb{R}^d$ is a D-embedding with D not too large, the images of the n copies of V in Y have to look like "nested cylinders". Let $g_i\colon V \to \mathbb{R}^d$ be the slice of g corresponding to i.

In order to speak of "inside and outside" of these images, we let S to be the whole surface of the cylinder C, including the top and the bottom, and we extend each g_i to a D-Lipschitz $\overline{g}_i\colon S \to \mathbb{R}^d$. Now the images of the lateral surface L of the cylinder under the $\overline{g}_i$ are disjoint (with an appropriate setting of r, ε, D), but we don't have much control over the images of the top and bottom. However, if we define P as a suitable neighborhood, of radius about DR, of the top and bottom of C, then it can be checked that $\overline{g}_i(S \setminus P)$ avoids $\overline{g}_j(S)$ for $i \ne j$. In this situation, if C is sufficiently long, Proposition 4.1 allows us to conclude that the images of the $\overline{g}_i$ are nested (in the sense defined in the proposition).

5 Stronger inapproximability for dimension 3: a warm-up

In this section we present a simple reduction, which provides an inapproximability result weaker than Theorem 1.2:

in that theorem, we claim the hardness of distinguishing between $O(1)$-embeddability and $n^{\mathrm{const}/d}$-embeddability, while here we show hardness of distinguishing between n^{ε}-embeddability ($\varepsilon > 0$ arbitrary but fixed) and $n^{\mathrm{const}/d}$-embeddability.

We will use an algorithmic problem called BETWEENNESS, which is NP-complete according to Opatrny [24] (the beautiful reduction of 3-SAT to this problem is also reproduced in [10]). An instance of BETWEENNESS is a set T of triples of the form (i,j,k), $i,j,k \in [n]$, and the problem is to decide whether T is *consistent*, i.e., whether there exists a linear ordering $\preceq$ of $[n]$ for which i is between j and k for every $(i,j,k) \in T$ (that is, either $j \preceq i \preceq k$ or $k \preceq i \preceq j$).

It will be more convenient to reduce to NON-BETWEENNESS, whose instance has the same form as for BETWEENNESS but the meaning of (i,j,k) is now "i must *not* be between j and k". Each constraint (i,j,k) in BETWEENNESS can be equivalently replaced by the two constraints (j,i,k) and (k,i,j) in NON-BETWEENNESS, and so NON-BETWEENNESS is NP-complete as well.

The reduction. Let $d \geq 3$ be fixed. Given an instance T of NON-BETWEENNESS for n elements and a bound D for distortion, we construct a metric space $\mathbb{Y} = (Y, \rho_Y)$, with $|Y| \leq (nD)^{O(d)}$, such that:

- If T is consistent, then $\mathbb{Y}$ is $O(n)$-embeddable in $\mathbb{R}^d$.
- If T is not consistent, then $\mathbb{Y}$ is not D-embeddable in $\mathbb{R}^d$.

Setting $D = n^C$ for a large constant C, we get that it is NP-hard to distinguish between $O(n)$-embeddability and n^C-embeddability of $\mathbb{Y}$ (and the size of $\mathbb{Y}$ is of order $n^{C_0 C d}$ for an absolute constant C_0).

We fix suitable parameters $\varepsilon \ll 1 \ll R$, with ε sufficiently small and R sufficiently large in terms of n and D, and we let S be a $(d-1)$-dimensional sphere of radius R. We let V be an ε-dense set in S, and similar to the example following Proposition 4.1, we set $Y_0 := [n] \times V$ and $\rho_{Y_0}((i,v),(i',v')) = \|v - v'\| + \delta_{ii'}$. We will refer to the set $\{i\} \times V$ as the *ith layer*. Next, we will modify (Y_0, ρ_{Y_0}) to obtain $\mathbb{Y}$; this modification reflects the structure of T.

We choose $|T|$ distinct points on V, sufficiently far from one another, corresponding to the triples in T. We will call these points the *loci*.

Let $(i,j,k) \in T$ and let $v = v_{(i,j,k)} \in V$ be the corresponding locus. We modify the metric space (Y_0, ρ_{Y_0}) near v as follows:

- We make a puncture of radius 1 in each of the layers except for the ith, jth, and kth. That is, we remove from Y_0 all points (ℓ, u) with $\ell \notin \{i,j,k\}$ and $\|u - v\| \leq 1$.
- We connect the jth and kth layers by a (discrete) path $\pi_{v,j,k}$ of length 1. Namely, we set $t = \lfloor 1/\varepsilon \rfloor$, we consider a graph-theoretic path on vertices $p_0, p_1, \ldots, p_t$ with edges of length $1/t$, and we glue this path to Y_0 by identifying p_0 with (j,v) and p_t with (k,v) (while $p_1, \ldots, p_{t-1}$ are new points).

Having made this modification for every triple of T, we call the resulting metric space $\mathbb{Y} = (Y, \rho_Y)$.

Now it is straightforward to check that for consistent instances, $\mathbb{Y}$ embeds with $O(n)$ distortion. Using Proposition 4.1, it is not hard to show that for inconsistent instances any embedding of $\mathbb{Y}$ incurs distortion at least D. We omit the details.

6 Proof of Theorem 1.2

Here we present a different reduction of NON-BETWEENNESS to approximate embeddability in $\mathbb{R}^d$, in which consistent instances yield $O(1)$-embeddability. The main idea is similar to the previous reduction: the linear ordering in NON-BETWEENNESS is encoded in nesting of suitable "discretized surfaces". The source of the $\Omega(n)$ distortion in the previous reduction was the nesting of all the surfaces at the same time.

Here we will allow simultaneous nesting of only at most 3 surfaces at a time. The surfaces won't be simply spheres, though, but rather each of them will resemble a network of branching pipes. We begin with a simple graph-theoretic lemma.

Lemma 6.1 *For every natural number n there is a graph G of size polynomial in n and subgraphs $G_1, G_2, \ldots, G_n$ of G such that (i) Each G_i, as well as each $G_i \cap G_j$, is a connected subgraph of G. (ii) No vertex of G belongs to more than 3 of the G_i. (iii) For every unordered triple $\{i,j,k\}$, there is a vertex $a_{ijk} \in V(G_i) \cap V(G_j) \cap V(G_k)$. (Proof omitted.)*

The construction. Let $d \geq 3$ be fixed. Given an instance T of NON-BETWEENNESS for n elements and a parameter D representing maximum distortion, we first construct an initial metric space $\mathbb{Y}_0$ that depends only on n and D.

We choose parameters $\varepsilon \ll 1 \ll R_{\mathrm{edge}} \ll R_{\mathrm{vert}}$ (polynomially depending on n and D, with the degree of the polynomial independent of d). We fix an embedding of the graph G as in Lemma 6.1 into $\mathbb{R}^d$, where vertices are represented by points and edges by straight segments. We assume that the minimum edge length is sufficiently large compared to R_{vert}, the maximum edge length is bounded by R_{vert} times a polynomial in n, the minimum distance of every two vertex-disjoint edges is much larger than R_{edge}, and that the minimum angle of two edges sharing a vertex is bounded below by an inverse polynomial in n.

We now "fatten" the embedded G: We replace each vertex $a \in V(G)$ by a ball B_a of radius R_{vert} and each edge e

by a cylinder C_e of radius $R_{\rm edge}$. We choose an ε-dense set V in the boundary of the resulting solid (the union of all B_a and all C_e). We let $V_a := V \cap \partial B_a$ and $V_e := V \cap \partial C_e$, and for $i \in [n]$ $V_i := \Big(\bigcup_{e\in E(G_i)} V_e\Big) \cup \Big(\bigcup_{a\in V(G_i)} V_a\Big)$, where the G_i are the subgraphs as in Lemma 6.1. The metric space $\mathbb{Y}_0 = (Y_0, \rho_{Y_0})$ is given by $Y_0 = \{(i,v) : i \in [n], v \in V_i\}$, $\rho_{Y_0}((i,v),(i',v')) = \|v - v'\| + \delta_{ii'}$.The ith layer of Y_0 is $\{i\} \times V_i$.

Now for every triple $(i,j,k) \in T$, we choose a point $v \in V_{a_{ijk}}$, not too close to any V_e, and we connect the points (j,v) and (k,v) by a discrete path of length 1 with spacing ε. (Since the vertices a_{ijk} are indexed by *unordered* triples, while the triples in T are ordered, we may need several such paths for a single vertex.) Adding such paths for all triples in T yields the metric space $\mathbb{Y}$.

If T is consistent, it is easy to embed $\mathbb{Y}$ in $\mathbb{R}^d$ with distortion $O(1)$. Rather than trying to formalize this, we refer to Fig. 2 for a (misleadingly planar) sketch for $n = 4$, with $T = \{(3,1,2), (4,1,2), (4,1,3), (2,3,4), (1,3,4)\}$ (for $n = 4$, the graph G can be taken very simple, as a K_4, with each G_i a triangle).

To show that D-embeddability implies consistency, we again apply Proposition 4.1. An additional issue, compared to the simpler reduction from the previous section, is showing that the orderings of the layers at different vertices are consistent. In this extended abstract we omit the proof.

Acknowledgment We thank Jussi Väisälä for providing us with a much better proof of Theorem 2.1(i) and (iii). We would like to thank Marianna Csörnyei for kindly answering a question of J. M. and suggesting a beautiful proof of the case $\delta = 0$ of Theorem 2.1(i).

References

[1] S. Arora, L. Lovász, I. Newman, Y. Rabani, Yu. Rabinovich, and S. Vempala. Local versus global properties of metric spaces. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[2] R. Babilon, J. Matoušek, J. Maxová, and P. Valtr. Low-distortion embeddings of trees. *J. Graph Algorithms Appl.*, 7:399–409, 2003.

[3] M. Bateni, M. Hajiaghayi, , E. D. Demaine, and M. Moharrami. Plane embeddings of planar graph metrics. *Discr. Comput. Geometry*, 38:615–637, 2007.

[4] J. Bourgain. On Lipschitz embedding of finite metric spaces into Hilbert space. *Isreal Journal of Mathematics*, 52:46–52, 1985.

[5] M. Bădoiu, , P. Indyk, and A. Sidiropoulos. Approximation algorithms for embedding general metrics into trees. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms*, 2007.

[6] M. Bădoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Low-distortion embeddings of general metrics into the line. In *Proc. 37th ACM Symposium on Theory of Computing*, 2005. Full version available at `http://www.mit.edu/~tasos/papers.html`.

[7] M. Bădoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Embedding ultrametrics into low-dimensional spaces. In *22nd Annual ACM Symposium on Computational Geometry*, 2006.

[8] M. Bădoiu, K. Dhamdhere, A. Gupta, Y. Rabinovich, H. Raecke, R. Ravi, and A. Sidiropoulos. Approximation algorithms for low-distortion embeddings into low-dimensional spaces. In *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms*, 2005.

[9] M. Charikar, K. Makarychev, and Y. Makarychev. Local global tradeoffs in metric embeddings. In *Proceedings of 48th Annual IEEE Symposium on Foundations of Computer Science*, 2007.

[10] B. Chor and M. Sudan. A geometric approach to betweenness. *SIAM J. Discrete Math.*, 11(4):511–523, 1998.

[11] M. M. Deza and M. Laurent. *Geometry of Cuts and Metrics*. Algorithms and Combinatorics 15. Springer-Verlag, Berlin etc., 1997.

[12] A. Gupta. Embedding tree metrics into low dimensional Euclidean spaces. *Discrete Comput. Geom.*, 24:105–116, 2000.

[13] A. Hall and C. H. Papadimitriou. Approximating the distortion. In *APPROX-RANDOM*, 2005.

[14] P. Indyk. Algorithmic applications of low-distortion embeddings. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[15] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mapping into Hilbert space. *Contemporary Mathematics*, 26:189–206, 1984.

[16] C. Kenyon, Y. Rabani, and A. Sinclair. Low distortion maps between point sets. *Proceedings of the Symposium on Theory of Computing*, 2004.

[17] M. D. Kirszbraun. Uber die zusammenziehenden und lipschitzchen Transformationen. *Fund. Math.*, 22:77–108, 1934.

[18] J. R. Lee, A. Naor, and Y. Peres. Trees and Markov convexity. In *Proc. 17th ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[19] N. Linial, E. London, and Yu. Rabinovich. The geometry of graphs and some its algorithmic applications. *Combinatorica*, 15:215–245, 1995.

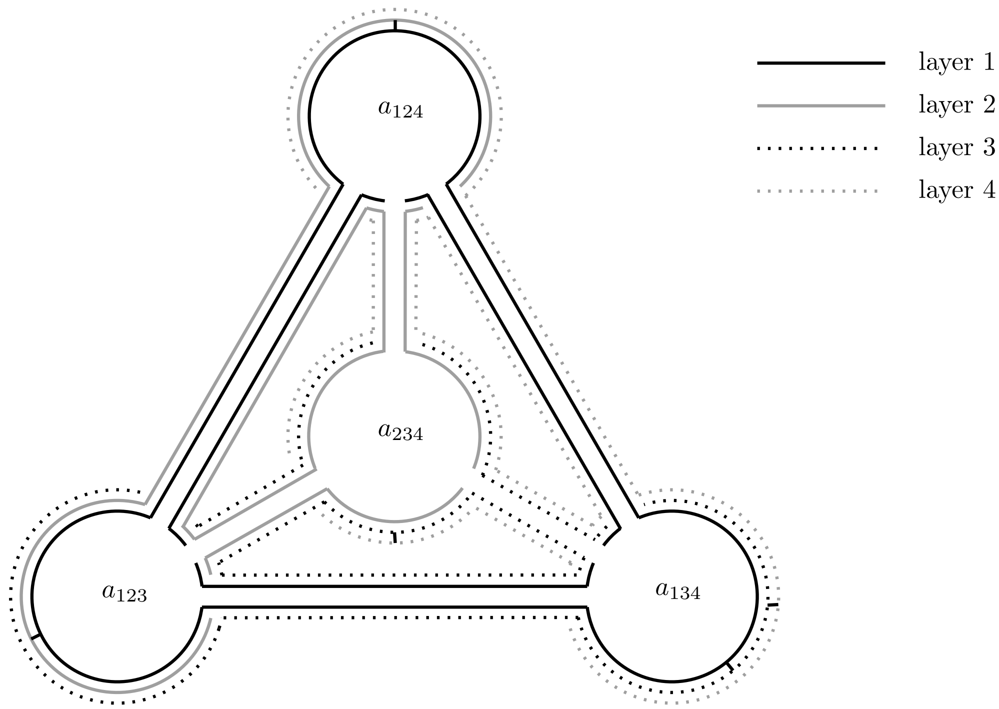

Figure 2: An embedding for a consistent instance.

[20] J. Matoušek. Bi-Lipschitz embeddings into low-dimensional Euclidean spaces. *Comment. Math. Univ. Carolinae*, 31:589–600, 1990.

[21] K. Menger. New foundation of Euclidean geometry. *American Journal of Mathematics*, 53(4):721–745, 1931.

[22] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, 1984.

[23] K. Onak and A. Sidiropoulos. Circular partitions with applications to visualization and embeddings. In *Proceedings of the 24th ACM Symposium on Computational Geometry, to appear*, 2008.

[24] J. Opatrny. Total ordering problem. *SIAM J. Comput.*, 8(1):111–114, 1979.

[25] C. Papadimitriou and S. Safra. The complexity of low-distortion embeddings between point sets. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 112–118, 2005.

[26] J. Väisälä. Holes and weakly bilipschitz maps. Manuscript, 3 pages, 2008.