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

Jiří Matoušek<br>Department of Applied Mathematics and Institute of Theoretical Computer Science (ITI)<br>Charles University, Malostranské nám. 25 11800 Praha 1, Czech Republic

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

Rev. July/14/2008 JM


#### 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 $n^{1 /(22 d-10)}$ 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 $\mathrm{P} \neq \mathrm{NP}$, 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 $O\left(n^{2 / d} \log ^{3 / 2} n\right)$ 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.


[^0]
## 1 Introduction

Let $\mathbb{X}=\left(X, \rho_{X}\right)$ and $\mathbb{Y}=\left(Y, \rho_{Y}\right)$ be metric spaces and let $f: X \rightarrow Y$ 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 $O\left(n^{2 / d} \log ^{3 / 2} n\right)$ [20]. The proof first embeds $\mathbb{X}$ into a high-dimensional Euclidean space using a well-known result of Bourgain [4, and then projects on a random $d$-dimensional subspace, following a method of Johnson and Lindenstrauss [15]. It provides a randomized polynomial-time algorithm for constructing an embedding with the distortion mentioned above. In particular, it yields an $O\left(n^{2 / d} \log ^{3 / 2} n\right)$ 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.

There exist $n$-point metric spaces $\mathbb{X}$ with $c_{\mathbb{R}^{d}}(\mathbb{X})=\Omega\left(n^{1 /\lfloor(d+1) / 2\rfloor}\right)$ [20], and thus the above worst-case upper bound cannot be much improved (in particular, for every even dimension it is tight up to the logarithmic factor).

Here we will show that, assuming $\mathrm{P} \neq \mathrm{NP}$, 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 $N P$-hard to approximate the minimum distortion required for embedding of a given $n$-point metric space into $\mathbb{R}^{d}$ within a factor of $\Omega\left(n^{1 /(22 d-10)-\varepsilon}\right)$.

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 $\varepsilon$-dens 1 finite set $V$ for a sufficiently small $\varepsilon>0$. Then we form another metric space $\mathbb{Y}=\left(Y, \rho_{Y}\right)$ as a suitable Cartesian product of $\mathbb{X}$ with $(V,\|\cdot\|)$.

[^1]

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

It is easy to show that $c_{\mathbb{R}^{d}}(\mathbb{Y})=O\left(c_{\mathbb{R}^{1}}(\mathbb{X})\right)$. The harder part is extracting an $O(D)$-embedding of $\mathbb{X}$ into $\mathbb{R}^{1}$ from an arbitrary given $D$-embedding $g: Y \rightarrow \mathbb{R}^{d}$; 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 $\left(a_{1}, a_{2}, \ldots, a_{n}\right)$. Then we define an embedding $f: X \rightarrow \mathbb{R}^{1}$ so that $f\left(a_{1}\right)<f\left(a_{2}\right)<$ $\cdots<f\left(a_{n}\right)$, and the difference $f\left(a_{i+1}\right)-f\left(a_{i}\right)$ is the distance of the $(i+1)$ st "deformed sphere" from the $i$ th 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.
$\ldots 3$ 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).

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 .

[^2]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 polynomialtime algorithm that computes an embedding of a given metric space $\mathbb{X}$ into $\mathbb{R}^{d}$ with distortion at $\operatorname{most} c_{\mathbb{R}^{d}}(\mathbb{X})^{O(1)}$.

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. Thus, isometric embeddability into $\mathbb{R}^{d}$ can be decided locally, by inspecting all $(d+3)$-tuples of points (we should remark that much better algorithms can be obtained by other methods).

In contract to this, we prove the following in Section 7:
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 $1+\varepsilon$.

Related work. Worst-case bounds for embedding in $\mathbb{R}^{d}$. Some special classes of metrics are known to embed in $\mathbb{R}^{d}$ with distortion better than the roughly $n^{2 / d}$ upper bound for general metrics: weighted trees with distortion $O\left(n^{1 /(d-1)}\right)$ [12] and ultrametrics with distortion $O\left(n^{1 / d}\right)$ [7]. Unweighted trees [2] and, more generally, unweighted outerplanar graphs [3] embed in $\mathbb{R}^{2}$ with distortion $O\left(n^{1 / 2}\right)$. On the other hand, there exist unweighted planar graphs for which $c_{\mathbb{R}^{2}}$ is $\Omega\left(n^{2 / 3}\right)$ 3].
Computing low-distortion embeddings in $\mathbb{R}^{1}$. We have already mentioned the inapproximability result of Bădoiu et al. [6] concerning embeddings of general metric spaces in $\mathbb{R}^{1}$. They complemented this result by an $O\left(n^{\beta}\right)$ approximation algorithm for embedding weighted trees in $\mathbb{R}^{1}$ for some constant $\beta<1$. For embedding an arbitrary metric space $\mathbb{X}$, they obtained an algorithm with approximation ratio depending on $\Delta(\mathbb{X})$; this was further improved in [5]. For embeddings in $\mathbb{R}^{1}$ there is also an $O\left(n^{1 / 3}\right)$ approximation algorithm for unweighted trees and an $O\left(n^{1 / 2}\right)$ approximation algorithm for unweighted graphs [8].
Embeddings in $\mathbb{R}^{2}$ and $\mathbb{R}^{3}$. In [7] it is shown that it is NP-hard to compute a minimum-distortion embedding into $\mathbb{R}^{2}$ with the $\ell_{\infty}$ metric. The same paper gives an $O\left(n^{1 / 3}\right)$-approximation algorithm for embedding an ultrametric in $\mathbb{R}^{2}$, and an $O\left(\log ^{O(1)} \Delta(\mathbb{X})\right)$-approximation algorithm for this case was given in [23].
Unlimited dimension and other cases. Linial et al. 19] observed that an embedding with the smallest possible distortion into $\ell_{2}$ (or equivalently, into a Euclidean space of an arbitrary dimension) can be computed in polynomial time via semidefinite programming. In contrast, it is well known that deciding isometric embeddability in $\ell_{1}$ is NP-hard (see [11]). Lee et al. [18] obtained an $O(1)$-approximation algorithm for embedding weighted trees into $\ell_{p}$.

The problem of approximating the minimum distortion embedding has also been studied for the case where we are given two metric spaces $\mathbb{X}, \mathbb{Y}$ of the same cardinality and we want to know $c_{\mathbb{Y}}(\mathbb{X})$; we refer to [13, 25, 16].


Figure 2: An extension of a $D$-embedding.

Menger-type questions. The kind of question which we addressed in Theorem 1.3 for embeddings in $\mathbb{R}^{2}$ was studied for embeddings in $\ell_{1}$ by Arora et al. [1] and their results were further strengthened by Charikar et al. [9]. The latter authors proved that if every $k$-point subspace of a an $n$-point metric space $\mathbb{X}$ embeds into $\ell_{1}$ with distortion $D$, then $\mathbb{X}$ embeds into $\ell_{1}$ with distortion $O(D \log (n / k))$; moreover, this result is nearly the best possible.

## 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 $\varepsilon$-dense in $S^{d-1}$, and we are given a $D$-embedding $g: V \rightarrow \mathbb{R}^{d}$. 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 $\operatorname{map} \bar{g}: S^{d-1} \rightarrow \mathbb{R}^{d}$ by a suitable interpolation (a tool for doing this will be mentioned in Section (2.2); see Fig. 2. We can make sure that $\bar{g}$ is still $D$-Lipschitz, but generally it won't be noncontracting, and it can even fail to be injective.

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

The main goal of this section is to show that the image of such $\bar{g}$ 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 .

### 2.1 Big holes and nested spheres

For a compact set $K \subset \mathbb{R}^{d}$, let us call a bounded component of $\mathbb{R}^{d} \backslash K$ a hole of $K$.
Theorem 2.1
(i) (A big hole exists) Let $\delta \in\left[0, \frac{1}{4}\right.$ ), let $f: S^{d-1} \rightarrow \mathbb{R}^{d}$ be a continuous map that satisfies $\| f(x)-$ $f(y)\|\geq\| x-y \|-\delta$ for all $x, y \in S^{d-1}$, and let $\Sigma:=f\left(S^{d-1}\right)$. Then $\Sigma$ has a hole containing
a ball of radius $\frac{1}{4}$.
(ii) Let $f_{1}, f_{2}: S^{d-1} \rightarrow \mathbb{R}^{d}$ be maps satisfying the condition as in (i) for some $\delta<\frac{1}{4}$, and suppose that, moreover, $\left\|f_{1}(x)-f_{2}(x)\right\|<\frac{1}{4}$ for all $x \in S^{d-1}$. Then some hole of $\Sigma_{1}:=f_{1}\left(S^{d-1}\right)$ intersects some hole of $\Sigma_{2}:=f_{2}\left(S^{d-1}\right)$.
(iii) (All holes but one are narrow) Let $\delta, f$, and $\Sigma$ 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 $\Sigma$ containing a ball of radius $4 D \delta$.

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).
Remark. The constants in the theorem are generally not optimal. However, in part (i), our method also provides an exact result "in the limit": By modifying the parameters in the proof appropriately (namely, setting $r=r(\delta)=1-2 \sqrt{\delta}$ and $\varepsilon=\sqrt{\delta}$ in the proof of Lemma 2.3 below), we obtain the existence of a hole containing a ball of radius $r(\delta)$ that tends to 1 as $\delta \rightarrow 0$.
Remark. In part (iii), for dimension $d \geq 3$ we cannot claim that all holes of $\Sigma$ but one have small diameter, since there can be many thin but very long holes. For $d=2$ all holes but one must have a small diameter [26].

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 $f_{1}, f_{2}, \ldots, f_{n}: S^{d-1} \rightarrow \mathbb{R}^{d}$ be continuous maps satisfying

- $\left\|f_{i}(x)-f_{i}(y)\right\| \geq\|x-y\|-\delta$ for all $x, y \in S^{d-1}$ and all $i$,
- $\left\|f_{i}(x)-f_{j}(x)\right\| \leq \frac{1}{4}$ for all $i, j$ and all $x \in S^{d-1}$, and
- $\Sigma_{i} \cap \Sigma_{j}=\emptyset$ whenever $i \neq j$, where $\Sigma_{i}=f_{i}\left(S^{d-1}\right)$.

Let $U_{i}$ denote the unbounded component of $\mathbb{R}^{d} \backslash \Sigma_{i}$, 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. It is clear that $\preceq$ is a partial ordering, so it suffices to check that for every $i \neq j$ we have either $U_{i} \subset U_{j}$ or $U_{j} \subset U_{i}$.

Theorem 2.1(ii) shows that some hole of $\Sigma_{i}$ intersects some hole of $\Sigma_{j}$. Since $\Sigma_{i} \cap \Sigma_{j}=\emptyset$ and $\Sigma_{i}, \Sigma_{j}$ are path-connected, $\Sigma_{i}$ and all of its holes are contained in some component of $\mathbb{R}^{d} \backslash \Sigma_{j}$, and vice versa, and the nesting lemma follows.

### 2.2 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 $\Sigma$.

Lemma 2.3 Let $f, \Sigma$, and $\delta<\frac{1}{4}$ be as in Theorem 2.1(i), and let $\Omega_{r}$ denote the closed $r$ neighborhood of $\Sigma$ in $\mathbb{R}^{d}$. Then there is a continuous map $h: \Omega_{1 / 4} \rightarrow S^{d-1}$ such that $\|h(f(x))-x\| \leq$ $8 \delta$ for all $x \in S^{d-1}$, and (consequently) the composition $h f: S^{d-1} \rightarrow S^{d-1}$ is homotopid ${ }^{3}$ 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 $H_{1} \rightarrow H_{2}$, 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 $\varepsilon$-net 4 in $\Sigma$. We choose a mapping $g: N \rightarrow S^{d-1}$ with $f g=\operatorname{id}_{N}$; 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 $y, y^{\prime} \in N$ are distinct and $x=g(y), x^{\prime}=g\left(y^{\prime}\right)$, then the condition on $f$ gives $\left\|x-x^{\prime}\right\| \leq\left\|f(x)-f\left(x^{\prime}\right)\right\|+\delta=\left\|y-y^{\prime}\right\|+\delta<(1+\delta / \varepsilon)\left\|y-y^{\prime}\right\|=2\left\|y-y^{\prime}\right\|$ since $\left\|y-y^{\prime}\right\|>\varepsilon$.

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

We can now define the desired $h: \Omega \rightarrow S^{d-1}$ as in the lemma, by $h(y)=\bar{g}(y) /\|\bar{g}(y)\|$.
Given $x \in S^{d-1}$, we pick $z \in N$ at most $\varepsilon$ away from $f(x)$, and we calculate $\|\bar{g}(f(x))-x\| \leq$ $\|\bar{g}(f(x))-g(z)\|+\|g(z)-x\| \leq 2\|z-f(x)\|+\|z-f(x)\|+\delta \leq 3 \varepsilon+\delta=4 \delta$. Since $\|\bar{g}(f(x))-h(f(x))\|=$ $1-\|\bar{g}(f(x))\| \leq\|x-\bar{g}(f(x))\|$, we obtain $\|h(f(x))-x\| \leq 2\|\bar{g}(f(x))-x\| \leq 8 \delta$ as claimed.

For $\delta<\frac{1}{4}$, this implies that $h(f(x)) \neq-x$ for all $x \in S^{d-1}$, and consequently, $h f \sim \operatorname{id}_{S^{d-1}}$ (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.

### 2.3 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, AlexanderSpanier) cohomology groun $5^{5} \breve{H}^{d-1}(X)$; 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 $\check{H}^{d-1}(X)$.

[^3]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 $\check{H}^{d-1}\left(S^{d-1}\right) \cong \mathbb{Z}$.

A continuous map $f: X \rightarrow Y$ of compact sets induces a group homomorphism $f^{*}: \check{H}^{d-1}(Y) \rightarrow$ $\check{H}^{d-1}(X)$; we should stress that $f^{*}$ goes in opposite direction compared to $f$. For the composition of maps we then have $(f g)^{*}=g^{*} f^{*}$ (the last two properties are usually expressed by saying that cohomology is a contravariant functor). Moreover, if $f_{1}, f_{2}: X \rightarrow Y$ 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 $j: X \rightarrow Y$ denote the inclusion map, and let $j^{*}: \check{H}^{d-1}(Y) \rightarrow \check{H}^{d-1}(X)$ 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 $i m a g e \operatorname{Im} j^{*}$.
(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 $\operatorname{Ker}\left(j_{1}^{*}\right) \cup \operatorname{Ker}\left(j_{2}^{*}\right)$ 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. The usual (modern) statement of Alexander duality, for the particular dimensions we are interested in, tells us that for every compact $X \subset \mathbb{R}^{d}$, there is an isomorphism

$$
\alpha_{X}: \check{H}^{d-1}(X) \rightarrow \tilde{H}_{0}\left(\mathbb{R}^{d} \backslash X\right)
$$

Here $\tilde{H}_{0}($.$) is the 0$-dimensional reduced singular homology group, which in our case can be concretely represented as follows: Each element $\gamma \in \tilde{H}_{0}\left(\mathbb{R}^{d} \backslash X\right)$ can be regarded as a function that assigns to every hole of $X$ an integer number (and with only finitely many nonzero values, in case that there are infinitely many holes). The group operation is componentwise addition; thus, $\tilde{H}_{0}\left(\mathbb{R}^{d} \backslash X\right)$ is a free Abelian group whose rank is the number of holes.

Now let $X \subseteq Y$, let $j: X \rightarrow Y$ be the inclusion map, and let $i: \mathbb{R}^{d} \backslash Y \rightarrow \mathbb{R}^{d} \backslash X$ be the inclusion map of the complements. We need a property of the Alexander duality called naturality with respect to inclusion $\sqrt[6]{6}$ which means that

$$
\begin{equation*}
\alpha_{X} j^{*}=i_{*} \alpha_{Y} \tag{1}
\end{equation*}
$$

where $i_{*}: \tilde{H}_{0}\left(\mathbb{R}^{d} \backslash Y\right) \rightarrow \tilde{H}_{0}\left(\mathbb{R}^{d} \backslash X\right)$ is the homomorphism in homology induced by the inclusion $i$.
If we represent $\tilde{H}_{0}\left(\mathbb{R}^{d} \backslash Y\right)$ and $\tilde{H}_{0}\left(\mathbb{R}^{d} \backslash X\right)$ as above, $i_{*}$ acts as follows: Given an integer function $\gamma$ on the holes of $Y$, the value of the function $i_{*}(\gamma)$ on a given hole $U$ of $X$ is the sum of $\gamma(V)$ over all holes $V$ of $Y$ contained in $U$. Thus, the rank of $\operatorname{Im} i_{*}$ is the number of holes of $X$ that contain some hole of $Y$, and since by (11), $\operatorname{Im} i_{*}$ and $\operatorname{Im} j^{*}$ are isomorphic, part (i) of the lemma follows.

[^4]For part (ii), from the assumption and (1) we get that $\operatorname{Ker}\left(i_{1 *}\right) \cup \operatorname{Ker}\left(i_{2 *}\right)$ do not generate all of $\tilde{H}_{0}\left(\mathbb{R}^{d} \backslash Y\right)$. Now $\operatorname{Ker}\left(i_{1 *}\right)$ contains all functions that are nonzero only on the holes of $Y$ not contained in any hole of $X_{1}$, and similarly for $\operatorname{Ker}\left(i_{2 *}\right)$. There is an element (function) $\gamma \in \tilde{H}_{0}\left(\mathbb{R}^{d} \backslash Y\right)$ that is not a linear combination of such functions, and such a $\gamma$ has to be nonzero on a hole that is contained both in a hole of $X_{1}$ and in a hole of $X_{2}$.

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

In part (ii), let $\Omega$ be the $\frac{1}{4}$-neighborhood of $\Sigma_{1}$, let $j_{1}: \Sigma_{1} \rightarrow \Omega$ be the inclusion map, and let $h_{1}: \Omega \rightarrow S^{d-1}$ 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 $S^{d-1} \rightarrow \Omega$, since the segment $f_{1}(x) f_{2}(x)$ is contained in $\Omega$ 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 $\check{H}^{d-1}\left(S^{d-1}\right)$.

The kernels of $j_{1}^{*}$ and $j_{2}^{*}$ are both contained in $\operatorname{Ker}\left(f_{1}^{*} j_{1}^{*}\right)=\operatorname{Ker}\left(f_{2}^{*} j_{2}^{*}\right)$, and the latter is not all of $\check{H}^{d-1}(\Omega)$. Thus, Lemma[2.4(ii) implies that there is a hole of $\Omega$ contained both in a hole of $\Sigma_{1}$ and in a hole of $\Sigma_{2}$, and part (ii) is proved.

For part (iii), let us consider the map $h$ as in Lemma 2.3, but restricted to the domain $\Sigma$. This time we set $r=4 D \delta$, and we let $j: \Sigma \rightarrow \Omega_{r}$ be the inclusion map. In part (i), we considered the composition $h f$; here we look at $f h$. This is a map $\Sigma \rightarrow \Sigma$, but we regard it as a map $\Sigma \rightarrow \Omega_{r}$, and we want to show that it is homotopic to $j$.

To this end, it suffices to check that $\|f(h(y))-y\| \leq 2 r$ for all $y \in \Sigma$, since then the segment connecting $f(h(y))$ to $y$ lies completely in $\Omega_{r}$ and it defines the desired homotopy.

Choosing $x \in S^{d-1}$ with $y=f(x)$, Lemma 2.3 gives $\|h(y)-x\| \leq 8 \delta$, and thus $\|f(h(y))-y\| \leq$ $8 D \delta=2 r$ as needed.

By the homotopy just established, we get that $h^{*} f^{*}=j^{*}$ as homomorphisms $\check{H}^{d-1}\left(\Omega_{r}\right) \rightarrow$ $\check{H}^{d-1}(\Sigma)$. But since $\operatorname{Im} f^{*} \subseteq H^{d-1}\left(S^{d-1}\right)$, which has rank $1, \operatorname{Im} j^{*}$ has rank at most 1 , and thus there is at most one hole of $\Sigma$ that contains a hole of $\Omega_{r}$.

## 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 $P \neq N P$, there is no polynomial-time algorithm with the following properties:

- The input of the algorithm is an n-point metric space $\mathbb{X}$ with $\Delta(\mathbb{X})=O(n)$.
- If $\mathbb{X}$ admits an $O\left(n^{4 / 12}\right)$-embedding into $\mathbb{R}^{1}$, the algorithm answers YES.
- If $\mathbb{X}$ is not embeddable in $\mathbb{R}^{1}$ with distortion smaller than $\Omega\left(n^{5 / 12-\varepsilon}\right)$, the algorithm answers NO.

This theorem together with the next proposition imply Theorem 1.1 by a simple calculation.
Proposition 3.2 Let $\mathbb{X}=\left(X, \rho_{X}\right)$ 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 $\mathbb{Y}=\left(Y, \rho_{Y}\right),|Y|=O\left(n D_{\max }^{2(d-1)} \Delta(\mathbb{X})^{d-1}\right)$, 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 ${ }^{7}$ 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 $1.1 D$ 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 $\Delta$. We let $C=C(d)$ be a sufficiently large constant, we set $R:=C D_{\max } \Delta$, and we let $S$ be the $(d-1)$-dimensional sphere in $\mathbb{R}^{d}$ centered at 0 of radius $R$. We set $\varepsilon:=\frac{1}{C D_{\max }}$, and we choose $V$ as an $\varepsilon$-dense subset of $S$ (that is, each point of $S$ has distance at most $\varepsilon$ to some point of $V$ ); as is well known, we can assume that $V$ has size $O\left((R / \varepsilon)^{d-1}\right)$ and is computable in time polynomial in $R / \varepsilon$.

Then we let $\mathbb{Y}=\left(Y, \rho_{Y}\right):=\mathbb{X} \times_{L_{2}} V$; that is, we set $Y:=X \times V$, and we define the metric $\rho_{Y}$ by

$$
\rho_{Y}\left((a, v),\left(a^{\prime}, v^{\prime}\right)\right)=\sqrt{\rho_{X}\left(a, a^{\prime}\right)^{2}+\left\|v-v^{\prime}\right\|^{2}} .
$$

Proof of Proposition 3.2(i). Let $f: X \rightarrow \mathbb{R}^{1}$ be a $D$-embedding; we assume that it is noncontracting and $D$-Lipschitz, and satisfies $\min f(X)=0$. We define $g: Y \rightarrow \mathbb{R}^{d}$ by $g(a, v):=$ $(R+f(a)) v$, so that the image of $(a, v)$ lies on the sphere of radius $R+f(a)$ concentric with $S$. This embedding is illustrated in Fig. 圂 (the picture is not realistic, though, since the radii should be larger and the points denser than shown).

Bounding the distortion of $g$ can be divided into two steps. Formally we "factor" $g$ through an auxiliary metric space $Z:=f(X) \times_{L_{2}} V$ (constructed from $f(X)$ with the metric of $\mathbb{R}^{1}$ in the same way as $\mathbb{Y}$ was constructed from $\mathbb{X})$. So we define $g_{1}: Y \rightarrow Z$ by $g_{1}(a, v)=(f(a), v)$ and $g_{2}: Z \rightarrow \mathbb{R}^{d}$ by $g_{2}(f(a), v)=g(a, v)$. Now $g=g_{2} g_{1}$, so $\operatorname{dist}(g) \leq \operatorname{dist}\left(g_{2}\right) \operatorname{dist}\left(g_{1}\right)$. It is almost obvious (and easy to check) that $\operatorname{dist}\left(g_{1}\right) \leq \operatorname{dist}(f)=D$. Checking $\operatorname{dist}\left(g_{2}\right) \leq 1.1$ (for $C$ sufficiently large) is a simple exercise; geometrically, it amounts to bounding the metric difference between a narrow annulus in $\mathbb{R}^{d}$ and a narrow cylindric band in $\mathbb{R}^{d+1}$, and we omit it.

For part (ii), let $g: Y \rightarrow \mathbb{R}^{d}$ 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 $g_{a}: V \rightarrow \mathbb{R}^{d}$ given by $g_{a}(v)=g(a, v)$.

Next, we extend each $g_{a}$ to a $D$-Lipschitz map $\bar{g}_{a}: S \rightarrow \mathbb{R}^{d}$ using the Kirszbraun theorem (mentioned after Lemma 2.3) Let $\Sigma_{a}:=\bar{g}_{a}(S)$ be the image of $\bar{g}_{a}$.

[^5]

Figure 3: Embedding of $Y$ from an embedding of $X$.

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

Lemma 3.3 For all $x, y \in S$ we have $\left\|\bar{g}_{a}(x)-\bar{g}_{a}(y)\right\| \geq\|x-y\|-\frac{3}{C}$.
Proof. (Routine.) We find $u \in V$ with $\|x-u\| \leq \varepsilon$, and similarly $v \in V$ for $y$. Then $\left\|\bar{g}_{a}(x)-\bar{g}_{a}(y)\right\| \geq\left\|g_{a}(u)-g_{a}(v)\right\|-\left\|\bar{g}_{a}(x)-g_{a}(u)\right\|-\left\|\bar{g}_{a}(y)-g_{a}(v)\right\| \geq\|u-v\|-2 D \varepsilon \geq \| x-$ $y\|-2 \varepsilon-2 D \varepsilon \geq\| x-y \|-\frac{3}{C}$.

The next lemma follows by very similar considerations and we omit the proof.

## Lemma 3.4

(i) For $a \neq b$, the Euclidean distance of $\Sigma_{a}$ and $\Sigma_{b}$ 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 $\left\|\bar{g}_{a}(x)-\bar{g}_{b}(x)\right\| \leq 2 D_{\max } \Delta$.

Proof of Proposition 3.2(ii). As was announced above, we can now apply Corollary 2.2 to the maps $\bar{g}_{a}$ with domain and range rescaled by $\frac{1}{R}$ (using $\delta=3 / C R<\frac{1}{4}$ and $2 D_{\max } \Delta / R \leq \frac{1}{4}$ ). Letting $U_{a}$ denote the unbounded component of $\mathbb{R}^{d} \backslash \Sigma_{a}$, 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 $\delta_{i}$ as the (Euclidean) distance of $\Sigma_{a_{i}}$ from $\Sigma_{a_{i+1}}$, and we define a mapping $f: X \rightarrow \mathbb{R}^{1}$ by $f\left(a_{i}\right)=\sum_{j=1}^{i-1} \delta_{j}$.

Assuming that the original mapping $g$ has distortion at most $D$, we will prove that $f$ has distortion at most $1.1 D$. First we show that $f$ contracts distances by a factor of at most 1.1. Lemma 3.4(i) gives $\delta_{i} \geq \rho_{X}\left(a_{i}, a_{i+1}\right)-2 / C \geq \rho_{X}\left(a_{i}, a_{i+1}\right) / 1.1$ (assuming $C$ large). The triangle inequality then shows that $\left|f\left(a_{i}\right)-f\left(a_{j}\right)\right| \geq \rho_{X}\left(a_{i}, a_{j}\right) / 1.1$ 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 $x_{i}:=g\left(a_{i}, v_{0}\right)$. 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 $\Sigma_{a_{k}}$. 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

$$
\begin{aligned}
\left|f\left(a_{j}\right)-f\left(a_{i}\right)\right| & =\sum_{k=i}^{j-1} \delta_{k} \leq \sum_{k=i}^{j-1}\left\|y_{k+1}-y_{k}\right\|=\left\|x_{i}-x_{j}\right\| \\
& \leq D \rho_{Y}\left(\left(a_{i}, v_{0}\right),\left(a_{j}, v_{0}\right)\right)=D \rho_{X}\left(a_{i}, a_{j}\right)
\end{aligned}
$$

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 $\Sigma_{a}$. To compare $\Sigma_{a}$ and $\Sigma_{b}$, we can compute a point with the minimum $x_{1}$-coordinate, say, of $\Sigma_{a} \cup \Sigma_{b}$ and see if it lies in $\Sigma_{a}$ or $\Sigma_{b}$ (here we can use a property which follows from the proof of the Kirszbraun theorem, namely, that we may assume $\bar{g}_{a}(S) \subseteq \operatorname{conv}\left(g_{a}(V)\right)$, which implies that the smallest point of $\Sigma_{a}$ lies in $g_{a}(V)$ ). Then we can approximate the distance of $\Sigma_{a}$ to $\Sigma_{b}$ 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 $\varepsilon$ 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 $S \subseteq \mathbb{R}^{d} \backslash\{0\}$ be a set, and let $\varepsilon>0$. We call a set $V \subseteq S \varepsilon$-angularly dense in $S$ if for every $x \in S$ there exists $v \in V$ with $\|x-v\| \leq \varepsilon\|v\|$.

We call a set $P \subseteq \mathbb{R}^{d} \varepsilon$-angularly small with respect to a set $V \subset \mathbb{R}^{d} \backslash\{0\}$ if there is a choice of a radius $r_{v} \geq 0$ for every $v \in V$ such that $P \subseteq \bigcup_{v \in V} B\left(v, r_{v}\right)$ (where $B(x, r)$ denotes the ball of radius $r$ centered at $x$ ) and $\sum_{v \in V} \frac{r_{v}}{\|v\|} \leq \varepsilon$ (this is a wasteful definition; aiming at more precise quantitative results, we would take $\left(r_{v} /\|v\|\right)^{d-1}$ 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 $\mathbb{R}^{d} \backslash\{0\}$ onto $S$ (i.e., $r_{S}: \mathbb{R}^{d} \backslash\{0\} \rightarrow S$ 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 $d \geq 2$, let $D \geq 1$ and let $\varepsilon:=\frac{1}{16 D}$, let $(S, P)$ be a punctured pseudosphere in $\mathbb{R}^{d}$, let $V \subset S$ be an $\varepsilon$-angularly dense set in $S$, let us assume that $P \subseteq S$ is $\varepsilon$-angularly small w.r.t. $V$, that $P \cap V=\emptyset$, and that $S \backslash P$ is path-connected, and let $f_{1}, f_{2}, \ldots, f_{n}: S \rightarrow \mathbb{R}^{d}$ be maps such that:

- Each $f_{i}$ is D-Lipschitz.
- Each $f_{i}$ restricted to $V$ is noncontracting.
- We have $\left\|f_{i}(v)-f_{j}(v)\right\| \leq \frac{1}{4}\|v\|$ for all $v \in V$ and all $i, j$.
- Setting $\Sigma_{i}:=f_{i}(S)$ and $\Sigma_{i}^{*}:=f_{i}(S \backslash P)$, we have $\Sigma_{i} \cap \Sigma_{j}^{*}=\emptyset$ for all $i \neq j$.

Let $U_{i}$ denote the unbounded component of $\mathbb{R}^{d} \backslash \Sigma_{i}$, 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 $S \rightarrow \mathbb{R}^{d}$ such that $f_{i}(x)=\tilde{f}_{i}(x)$ for all $x \in S \backslash P$ and all $i$ (in particular, $\Sigma_{i}^{*}=f_{i}(S \backslash P)=\tilde{f}_{i}(S \backslash P)$ ), and $\Sigma_{i}^{*} \cap \tilde{f}_{j}(S)=\emptyset$ for all $i \neq j$, then the linear ordering induced by the $\tilde{f}_{i}$ is the same as $\preceq$.

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 $R$ (see Fig. 4 left), and let $V$ be a set that is $\varepsilon$-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 $\rho_{Y}\left((i, v),\left(i^{\prime}, v^{\prime}\right)\right)=\left\|v-v^{\prime}\right\|+\delta_{i i^{\prime}}$, where $\delta_{i i^{\prime}}$ is the Kronecker delta (equal to 0 for $i=i^{\prime}$ and to 1 otherwise).

We assume $\varepsilon \ll 1 \ll R$, and so we expect that if $g: Y \rightarrow \mathbb{R}^{d}$ 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 $g_{i}: V \rightarrow \mathbb{R}^{d}$ 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 $\bar{g}_{i}: S \rightarrow \mathbb{R}^{d}$. Now the images of the lateral surface $L$ of the cylinder under the $\bar{g}_{i}$ are disjoint (with an appropriate setting of $r, \varepsilon, D)$, but we don't have much control over the images of the top and bottom, as is schematically indicated in Fig. 4 on the right (in a 2-dimensional cross-section). However, if we define $P$ as a suitable neighborhood, of radius about $D R$, of the top and bottom of $C$, then it can be checked that $\bar{g}_{i}(S \backslash P)$ avoids $\bar{g}_{j}(S)$ for $i \neq j$. In this situation, if $C$ is sufficiently long, Proposition 4.1 allows us to conclude that the images of the $\bar{g}_{i}$ are nested (in the sense defined in the proposition).
Towards the proof of Proposition 4.1. Let us define

$$
\Omega:=\bigcup_{v \in V} B\left(f_{1}(v), \frac{1}{2}\|v\|\right)
$$

As the next lemma shows, $\Omega$ is a neighborhood of all the $\Sigma_{i}$, and to some extent it will play the role of $\Omega_{\frac{1}{4}}$ from the proof of Theorem 2.1(i)-(ii).

Lemma 4.2 We have $\Sigma_{i} \subseteq \Omega$ for all $i$, and all the $f_{i}$ are homotopic as maps $S \rightarrow \Omega$.

Proof. (A routine " $\varepsilon$-density" argument.) Given a point $f_{i}(x) \in \Sigma_{i}$, we choose $v \in V$ with $\|x-v\| \leq \varepsilon\|v\|$. Then $\left\|f_{i}(x)-f_{1}(v)\right\| \leq\left\|f_{i}(x)-f_{i}(v)\right\|+\left\|f_{i}(v)-f_{1}(v)\right\| \leq D \varepsilon\|v\|+\frac{1}{4}\|v\|<\frac{1}{2}\|v\|$, and hence $f_{i}(x) \in B\left(f_{1}(v), \frac{1}{2}\|v\|\right)$. Similarly $f_{1}(x) \in B\left(f_{1}(v), \frac{1}{2}\|v\|\right)$, so the segment $f_{1}(x) f_{i}(x)$ is contained in $\Omega$ for every $x$, and thus $f_{i} \sim f_{1}$ as claimed.


Figure 4: The set $V$ (left), and an embedding of $Y$ and its extension (right).

Lemma 4.3 There is a point $c \notin \Omega$ that lies in a hole $H_{i}$ of $\Sigma_{i}$ for each $i$.

Proof. We imitate the proof of Theorem [2.1(ii). We set $N:=f_{1}(V)$, for every $z \in N$ we fix $g(z) \in f_{1}^{-1}(z)$, and we fix a 1-Lipschitz extension $\bar{g}: \Omega \rightarrow \mathbb{R}^{d}$ of $g$.

We check that $0 \notin \operatorname{Im} \bar{g}$ : Assuming $0=\bar{g}(y), y \in \Omega$, we choose $v \in V$ with $y \in B\left(f_{1}(v), \frac{1}{2}\|v\|\right)$. Then $\|v\|=\left\|\bar{g}(y)-\bar{g}\left(f_{1}(v)\right)\right\| \leq\left\|y-f_{1}(v)\right\| \leq \frac{1}{2}\|v\|$, a contradiction since $v \neq 0$.

Next, we claim that $\bar{g} f_{1} \sim \operatorname{id}_{S}$ as maps $S \rightarrow \mathbb{R}^{d} \backslash\{0\}$. It suffices to verify that the segment connecting $x$ to $\bar{g}\left(f_{1}(x)\right)$ doesn't contain 0 , for all $x \in S$. As usual, we fix $v \in V$ with $\|x-v\| \leq \varepsilon\|v\|$; then we also have $\|v\| \leq\|x\| /(1-\varepsilon) \leq 2\|x\|$. We estimate $\left\|\bar{g}\left(f_{1}(x)\right)-x\right\| \leq\left\|\bar{g}\left(f_{1}(x)\right)-\bar{g}\left(f_{1}(v)\right)\right\|+$ $\left\|\bar{g}\left(f_{1}(v)\right)-x\right\|=\left\|\bar{g}\left(f_{1}(x)\right)-\bar{g}\left(f_{1}(v)\right)\right\|+\|v-x\| \leq D\|v-x\|+\varepsilon\|v\| \leq(D+1) \varepsilon\|v\| \leq 4 D \varepsilon\|x\|<\frac{1}{2}\|x\|$, and $\bar{g} f_{1} \sim \mathrm{id}_{S}$ follows.

Next, we set $h:=r_{S} \bar{g}$, where $r_{S}$ is the retraction of $\mathbb{R}^{d} \backslash\{0\}$ onto $S$. It follows that $h f_{1} \sim \operatorname{id}_{S}$ as well (compose the homotopy witnessing $\bar{g} f_{1} \sim \operatorname{id}_{S}$ with $r_{S}$ ). With such an $h$ at our disposal, an argument very similar to the proof of Theorem 2.1(ii) finishes the proof of the lemma.

Proof of Proposition 4.1. Let $c$ be a point as in Lemma 4.3, for notational convenience, we will assume that $c=0$. Let $H_{i}$ denote the hole of $\Sigma_{i}$ that contains 0 .

For a set $X \subseteq \mathbb{R}^{d} \backslash\{0\}$, let $\mu(X)$ denote the "spatial angle" of $X$ as seen from 0 , i.e., the fraction of $S^{d-1}$ occupied by the central projection of $X$. We claim that if $x \neq 0$ and $r \leq \frac{1}{2}\|x\|$, then we have $\mu(B(x, r)) \leq r /\|x\|$. Indeed, a simple projection argument shows that it is enough to deal with the case $d=2$, and this case follows easily with a bit of trigonometry.

We claim that $\mu\left(f_{i}(P)\right) \leq \frac{1}{4}$ for all $i$. To check this, we use the assumption that the "puncture set" $P$ is $\varepsilon$-angularly small w.r.t. $V$ : If $r_{v}$ are the radii as in the definition of " $\varepsilon$-angularly small", then $f_{i}(P) \subseteq \bigcup_{v \in V} B\left(f_{i}(v), D r_{v}\right)$, and thus

$$
\mu\left(f_{i}(P)\right) \leq \sum_{v \in V} \frac{D r_{v}}{\left\|f_{i}(v)\right\|}
$$

Since $0 \notin \Omega$, the definition of $\Omega$ gives $\left\|f_{1}(v)\right\|>\frac{1}{2}\|v\|$, and so $\left\|f_{i}(v)\right\| \geq\left\|f_{1}(v)\right\|-\left\|f_{1}(v)-f_{i}(v)\right\| \geq$ $\frac{1}{2}\|v\|-\frac{1}{4}\|v\|=\frac{1}{4}\|v\|$. Then $\mu\left(f_{i}(P)\right) \leq 4 D \sum_{v \in V} r_{v} /\|v\| \leq 4 D \varepsilon=\frac{1}{4}$.

Now we are ready to prove that $\preceq$ is a linear ordering. We first look at two of the indices, say $i=1$ and $j=2$, and show that either $1 \preceq 2$ or $2 \preceq 1$. Let $U_{12} \subseteq U_{1} \cap U_{2}$ denote the unbounded component of $\mathbb{R}^{d} \backslash\left(\Sigma_{1} \cup \Sigma_{2}\right)$, and let $H_{12} \subseteq H_{1} \cap H_{2}$ be the hole of $\Sigma_{1} \cup \Sigma_{2}$ containing 0 .

The boundary $\partial U_{12}$ is contained in $\Sigma_{1} \cup \Sigma_{2}$, and we have $\mu\left(\partial U_{12}\right)=1$. It follows that $\partial U$ cannot be covered by $f_{1}(P) \cup f_{2}(P)$, and thus at least one of $\Sigma_{1}^{*}=f_{1}(S \backslash P)$ and $\Sigma_{2}^{*}=f_{2}(S \backslash P)$ is incident to $\partial U_{12}$; let us suppose $x \in \Sigma_{2}^{*} \cap \partial U_{12}$, for example. Then, since $\Sigma_{2}^{*} \cap \Sigma_{1}=\emptyset$, we have $x \in U_{1}$, and since $\Sigma_{2}^{*}$ is path-connected, we obtain $\Sigma_{2}^{*} \subset U_{1}$. Hence $1 \preceq 2$.

By a similar argument, looking at $\partial H_{12}$, we get that one of $\Sigma_{1}^{*} \subset H_{2}$ or $\Sigma_{2}^{*} \subset H_{1}$ holds, and this shows that $1 \preceq 2$ and $2 \preceq 1$ cannot hold simultaneously.

It remains to verify transitivity of $\preceq$. To this end, we consider three indices, say $1,2,3$, and we suppose $1 \preceq 2 \preceq 3$. This means that $\Sigma_{2}^{*} \subset U_{1}$ and $\Sigma_{3}^{*} \subset U_{2}$. If we had $3 \preceq 1$ as well, then $\Sigma_{1}^{*} \subset U_{3}$. So $\Sigma_{1}^{*} \cup \Sigma_{2}^{*} \cup \Sigma_{3}^{*} \subset U_{1} \cup U_{2} \cup U_{3}$. But a measure argument as above shows that at least one of $\Sigma_{1}^{*}$, $\Sigma_{2}^{*}, \Sigma_{3}^{*}$ must appear on the boundary of the hole $H_{123}$ of $\Sigma_{1} \cup \Sigma_{2} \cup \Sigma_{3}$ that contains 0 , but this is impossible since $H_{123} \subseteq H_{1} \cap H_{2} \cap H_{3}$, and the latter is disjoint from $U_{1} \cup U_{2} \cup U_{3}$. This shows that $\preceq$ is indeed a linear ordering on $[n]$.

It remains to verify the assertion that $\preceq$ doesn't depend on the behavior of the $f_{i}$ on the puncture set $P$. Thus, let the $\tilde{f}_{i}$ be as in the proposition. We can replace the $f_{i}$ by the $\tilde{f}_{i}$ gradually one by one (since the already proved part of the proposition applies to $f_{1}, f_{2}, \ldots, f_{i}, \tilde{f}_{i+1}, \ldots, \tilde{f}_{n}$ and shows that they induce some linear ordering), and thus it suffices to consider $f_{1}, f_{2}$, and $\tilde{f}_{2}$ and show that $1 \preceq 2$ iff $1 \preceq^{\prime} 2$, where $\preceq^{\prime}$ is the relation induced by $f_{1}$ and $\tilde{f}_{2}$.

First, if $1 \preceq 2$, then $\Sigma_{2}^{*}$ lies in the unbounded component of $\mathbb{R}^{d} \backslash \Sigma_{1}$, and hence $1 \preceq^{\prime} 2$ by definition. Second, if $2 \preceq 1$, then as shown above, $\Sigma_{2}^{*}$ is contained in a hole of $\Sigma_{1}$, and this means $2 \preceq^{\prime} 1$. The proposition is proved.

## 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 $n^{\text {const/d }}$-embeddability, while here we show hardness of distinguishing between $n^{\varepsilon}$-embeddability ( $\varepsilon>0$ arbitrary but fixed) and $n^{\text {const } / d}$-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 $\mathbb{Y}=\left(Y, \rho_{Y}\right)$, with $|Y| \leq$ $(n D)^{O(d)}$, 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 $n^{C C_{0} C d}$ for an absolute constant $C_{0}$ ).

We fix suitable parameters $\varepsilon \ll 1 \ll R$, with $\varepsilon$ 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 $\varepsilon$-dense set in $S$, and similar to the example following Proposition 4.1, we set $Y_{0}:=[n] \times V$ and $\rho_{Y_{0}}\left((i, v),\left(i^{\prime}, v^{\prime}\right)\right)=\left\|v-v^{\prime}\right\|+\delta_{i i^{\prime}}$. We will refer to the set $\{i\} \times V$ as the $i$ th layer. Next, we will modify $\left(Y_{0}, \rho_{Y_{0}}\right)$ 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 $\left(Y_{0}, \rho_{Y_{0}}\right)$ near $v$ as follows:

- We make a puncture of radius 1 in each of the layers except for the $i$ th, $j$ th, and $k$ th. That is, we remove from $Y_{0}$ all points $(\ell, u)$ with $\ell \notin\{i, j, k\}$ and $\|u-v\| \leq 1$.
- We connect the $j$ th and $k$ th 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 $\mathbb{Y}=\left(Y, \rho_{Y}\right)$.
Embeddability for consistent instances. We assume that $T$ is consistent and, for notational convenience, that a linear ordering obeying all the constraints is the natural ordering $\leq$. Then we embed $\mathbb{Y}$ in $\mathbb{R}^{d}$ as sketched in Fig. 5 (the picture is 2 -dimensional although the actual embedding is in dimension 3 or higher). The $i$ th layer is mapped isometrically to the sphere of radius $R+i$ centered at 0 , and the connecting paths $\pi_{v, j, k}$ are embedded on straight segments that pass through the punctures in the other layers. Clearly, this embedding incurs distortion $O(n)$.
$\boldsymbol{D}$-embeddability implies consistency. Next, let $g$ be an arbitrary $D$-embedding of $\mathbb{Y}$ into $\mathbb{R}^{d}$, which we assume to be noncontracting. We let $g_{i}: V \rightarrow \mathbb{R}^{d}$ be the corresponding $D$-embedding of the $i$ th layer, and let $\bar{g}_{i}$ be a $D$-Lipschitz extension of $g_{i}$ on $S$.

For applying Proposition4.1, we need to define the "puncture set" $P$ in such a way that $\bar{g}_{i}(S \backslash P)$ is guaranteed to be disjoint from $\bar{g}_{j}(S), j \neq i$. Since the punctures in the construction of $\mathbb{Y}$ have diameter 2 and each $\bar{g}_{i}$ is $D$-Lipschitz, it suffices to choose $P \subseteq S$ as the union balls of radius $3 D$ centered at all the loci (and intersected with $S$ ).


Figure 5: Embedding for a consistent instance ( $n=4, T=\{(3,1,2),(4,1,2),(4,1,3),(2,3,4)$, $(1,3,4)\})$.

If $R$ is sufficiently large in terms of $D$, then the complement $S \backslash P$ is connected; this is the moment in the proof where we need the assumption $d \geq 3$ (since if we puncture 1-dimensional sphere more than once, it falls apart!). For metric reasons, we then have $\Sigma_{i} \cap \Sigma_{j}^{*}=\emptyset$ for $i \neq j$, where $\Sigma_{i}:=\bar{g}_{i}(S), \Sigma_{i}^{*}:=\bar{g}_{i}(S \backslash P)$. It is straightforward to verify the assumptions of Proposition4.1 with $f_{i}=\bar{g}_{i}$ (we note that $V$ here and in the proposition are the same, but the $\varepsilon$ in the proposition, equal to $\frac{1}{16 D}$, need not be the same).

We thus conclude that the images $\bar{g}_{i}(S)$ are nested in the sense of Proposition 4.1, and we claim that the corresponding linear ordering $\preceq$ of $[n]$ is consistent for $T$. Indeed, for contradiction let us suppose that $(i, j, k) \in T$ but $j \preceq i \preceq k$, say. Then $\Sigma_{k}^{*}$ is contained in the unbounded component of $\Sigma_{i}$, while $\Sigma_{j}^{*}$ lies in a hole of $\Sigma_{i}$.

If $v$ is the locus of $(i, j, k)$, then none of the layers $i, j, k$ has a puncture near $v$, and thus for metric reasons, $\Sigma_{j}, \Sigma_{i}$, and $\Sigma_{k}$ do not intersect in the $3 D$-neighborhood of $g_{i}(v)$. It follows that $g_{j}(v)$ lies in the same connected component of $\mathbb{R}^{d} \backslash \Sigma_{i}$ as $\Sigma_{j}^{*}$, i.e., in a hole, while $g_{k}(v)$ is in the unbounded component.

We now consider the image under $g$ of the discrete path $\pi_{v, j, k}$ in $\mathbb{Y}$, and we extend this image by segments to a continuous path $\gamma$ connecting $g_{j}(v)$ and $g_{k}(v)$. Then $\gamma$ intersects $\Sigma_{i}$ at a point $x$. On $\gamma$ we can find a point $g(p)$ at distance at most $D \varepsilon$ from $x$, where $p \in \pi_{v, j, k}$, and in $\Sigma_{i}$ there is a point of the form $g(i, u), u \in V$, also at distance at most $D \varepsilon$ from $x$. But $\rho_{Y}((i, u), p) \geq 1$, and so $2 D \varepsilon \geq 1$ since $g$ was assumed noncontracting - a contradiction for $\varepsilon$ sufficiently small.

## 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

- Each $G_{i}$, as well as each $G_{i} \cap G_{j}$, is a connected subgraph of $G$.
- No vertex of $G$ belongs to more than 3 of the $G_{i}$.
- For every unordered triple $\{i, j, k\}$, there is a vertex $a_{i j k} \in V\left(G_{i}\right) \cap V\left(G_{j}\right) \cap V\left(G_{k}\right)$.

Proof. The vertex set of $G$ can be taken as $\binom{[n]}{2} \cup\binom{[n]}{3}$ (all 2-element and 3-element subsets of $[n])$, and the edges are of the form $\{\{i, j\},\{i, j, k\}\}, i, j, k$ all distinct. We let $G_{i}$ be the subgraph induced by the set of all $S \in V(G)$ with $i \in S$. Verifying the required properties is immediate.

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 $\varepsilon \ll 1 \ll R_{\text {edge }} \ll R_{\text {vert }}$ (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 $R_{\text {vert }}$, the maximum edge length is bounded by $R_{\text {vert }}$ times a polynomial in $n$, the minimum distance of every two vertex-disjoint edges is much larger than $R_{\text {edge }}$, 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 $R_{\text {vert }}$ and each edge $e$ by a cylinder $C_{e}$ of radius $R_{\text {edge }}$. We choose an $\varepsilon$-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 $i \in[n]$

$$
V_{i}:=\left(\bigcup_{e \in E\left(G_{i}\right)} V_{e}\right) \cup\left(\bigcup_{a \in V\left(G_{i}\right)} V_{a}\right),
$$

where the $G_{i}$ are the subgraphs as in Lemma 6.1. The metric space $\mathbb{Y}_{0}=\left(Y_{0}, \rho_{Y_{0}}\right)$ is given by

$$
\begin{aligned}
Y_{0} & =\left\{(i, v): i \in[n], v \in V_{i}\right\} \\
\rho_{Y_{0}}\left((i, v),\left(i^{\prime}, v^{\prime}\right)\right) & =\left\|v-v^{\prime}\right\|+\delta_{i i^{\prime}} .
\end{aligned}
$$

The $i$ th layer of $Y_{0}$ is $\{i\} \times V_{i}$.
Now for every triple $(i, j, k) \in T$, we choose a point $v \in V_{a_{i j k}}$, 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 $\varepsilon$. (Since the vertices $a_{i j k}$ 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 $\mathbb{Y}$.


Figure 6: An embedding for a consistent instance.

Consistent instances. 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. 6 for a (misleadingly planar) sketch for $n=4$, with $T$ the same as in Fig. 5 (for $n=4$, the graph $G$ can be taken very simple, as a $K_{4}$, with each $G_{i}$ a triangle).
$D$-embeddability implies consistency. We consider a noncontracting $D$-embedding $g: Y \rightarrow$ $\mathbb{R}^{d}$. Using Proposition 4.1 as in Section 5, we get that for each vertex $a$ of $G$ the embedding defines a linear ordering of at most 3 layers present at $a$, and similarly for each edge of $G$.

Next, we check that these orderings are locally consistent between a vertex $a$ and the adjacent edges.

Lemma 6.2 Let $\preceq_{a}$ be the linear ordering of the layers present at a vertex $a$, and let $\preceq_{e}$ be the linear ordering of the layers at an edge e incident to $a$. If layers $i$ and $j$ are present both at $a$ and at $e$, then $i \preceq_{a} j$ iff $i \preceq_{e} j$.

Proof. For notational convenience we will consider $i=1$ and $j=2$.
We consider three punctured pseudospheres: $S_{a}:=\partial B_{a}, S_{e}:=\partial C_{e}, S_{a e}:=\partial\left(B_{a} \cup C_{e}\right)$. The puncture sets $P_{a}, P_{e}, P_{a e}$ are defined as expected: $P_{a} \subset S_{a}$ is the union of suitable neighborhoods of the cylinders $C_{e^{\prime}}$ for all edges $e^{\prime}$ incident to $a, P_{e} \subset S_{e}$ consists of suitable neighborhoods of the top and bottom sides of the cylinder $C_{e}$, and $P_{a e} \subset S_{a e}$ is $P_{a} \cup P_{e}$ minus the parts of $P_{e}$ and $P_{a}$ surrounding the place where $C_{e}$ touches $B_{a}$; see Fig. 7.

We consider the restriction $g_{i}$ of the embedding $g: Y \rightarrow \mathbb{R}^{d}$ to the $i$ th layer, $i=1,2$. We fix a $D$-Lipschitz map $\bar{g}_{i, a}: S_{a} \rightarrow \mathbb{R}^{d}$ that coincides with $g_{i}$ on $V_{a}$, and similarly for $\bar{g}_{i, e}$ and $\bar{g}_{i, a e}$;


Figure 7: The ball $B_{a}$ and the cylinders adjacent to it (left); the punctured pseudospheres $S_{a}, S_{e}, S_{a e}$ (right); puncture sets indicated by a thick line.


Figure 8: Two consecutive punctured pseudospheres $S_{k}$ and $S_{k+1}$.
moreover, we make sure that $\bar{g}_{i, a}$ and $\bar{g}_{i, e}$ coincide with $\bar{g}_{i, a e}$ on the shared parts of the domains (to this end, we can first extend $g_{i}$ to all of $S_{a} \cup S_{e}$ and then take $\bar{g}_{i, a}, \bar{g}_{i, e}$, and $\bar{g}_{i, a e}$ to be restrictions of this common extension).

The relation $\preceq_{a}$ is determined by the nesting of the images $\Sigma_{i, a}:=\bar{g}_{i, a}\left(S_{a}\right)$, and similarly for $\preceq_{e}$. Instead of relating $\preceq_{a}$ and $\preceq_{e}$ directly, we relate both of them to $\preceq_{a e}$, which is the relation defined by the nesting of $\Sigma_{1, a e}$ and $\Sigma_{2, a e}$.

Let us first show that $\preceq_{a}$ and $\preceq_{a e}$ are the same. We will use the last part of Proposition 4.1, which tells us that the nesting relation doesn't depend on the behavior of the considered maps on the puncture set, and also the freedom to choose the puncture set.

We define a sequence of punctured pseudospheres $S_{0}=S_{a}, S_{1}, \ldots, S_{t}=S_{a e}$ that interpolate between $S_{a}$ and $S_{a e}$ : We start with the sphere $S_{a}$, and we gradually "grow" an attached cylinder from it, as in Fig. 8. The puncture sets $P_{k}$ are as indicated in the picture.

We let $\bar{g}_{i, k}: S_{k} \rightarrow \mathbb{R}^{d}$ be a suitable $D$-Lipschitz map that coincides with $\bar{g}_{i, a e}$ on the common part of their domain (where $\bar{g}_{i, 0}=\bar{g}_{i, a}$ and $\bar{g}_{i, t}=\bar{g}_{i, a e}$ ). Let $\Sigma_{i, k}:=\bar{g}_{i, k}\left(S_{k}\right)$.

For each $k=0,1, \ldots, t$, we thus get a "nesting" relation $\preceq_{k}$. If $\preceq_{a}$ and $\preceq_{a e}$ were not the same, we would get that for some $k, \preceq_{k}$ is different from $\preceq_{k+1}$.

To see that this is impossible, we choose a 1-Lipschitz map $h: S_{k+1} \rightarrow S_{k}$ that is the identity map on $S_{k+1} \cap S_{k}$ (we just contract a piece of the lateral surface of the cylinder $C_{e}$ ). Then we
define $g_{i}^{\prime}: S_{k+1} \rightarrow \mathbb{R}^{d}$ as the composition $\bar{g}_{i, k} h$.
On the one hand, if we use Proposition 4.1 with $S_{k+1}$ as the pseudosphere $S, P_{k+1}$ as the puncture set $P, \bar{g}_{i, k+1}$ as $f_{i}, i=1,2$, and $g_{i}^{\prime}$ as $\tilde{f}_{i}, i=1,2$, we get that the nesting relation $\preceq^{\prime}$ determined by $g_{1}^{\prime}$ and $g_{2}^{\prime}$ coincides with $\preceq_{k+1}$, since $\bar{g}_{i, k+1}$ agrees with $g_{i}^{\prime}$ outside the puncture set.

On the other hand, the nesting relation depends solely on the images of the considered maps. In our case we have $g_{i}^{\prime}\left(S_{k+1}\right)=\Sigma_{i, k}=\bar{g}_{i, k}\left(S_{k}\right), i=1,2$. For the images of the "non-puncture sets" we have the inclusion $\Sigma_{i, k}^{*}=\bar{g}_{i, k}\left(S_{k} \backslash P_{k}\right) \subseteq g_{i}^{\prime}\left(S_{k+1} \backslash P_{k+1}\right)$, and the latter set is path-connected. Therefore, considering the way the nesting relation is defined in Proposition 4.1, we get that $\preceq^{\prime}$ also coincides with $\preceq_{k}$. In conclusion, $\preceq_{a}$ is the same as $\preceq_{a e}$.

It remains to show that $\preceq_{e}$ and $\preceq_{a e}$ are also the same. We proceed similarly, but this time a single step is actually sufficient. We consider yet another punctured pseudosphere $S_{e a}$, which is $S_{a e}$ translated so that the origin is in the middle of the cylinder $C_{e}$ (as in $S_{e}$ ). The puncture set $P_{e a}$ of $S_{e a}$ includes the appropriately translated copy of $P_{a e}$ plus all the spherical part of $S_{e a}$ (we take advantage of the fact that the spherical part is angularly small in $S_{e a}$, provided that $C_{e}$ is sufficiently long compared to the radius of $S_{a}$ ).

The map $\bar{g}_{i, e a}$ is defined on $S_{e a}$ in the same way as $\bar{g}_{i, a e}$ on $S_{a e}$ (the only difference is the translation of the domain). The nesting relation $\preceq_{e a}$ obtained from these maps is the same as $\preceq_{a e}$ (since the images are the same, up to an enlargement of the puncture set).

The argument showing that $\preceq_{e a}$ coincides with $\preceq_{e}$ is then almost the same as the one above for $\preceq_{k}$ and $\preceq_{k+1}$, using a suitable 1-Lipschitz map $h^{\prime}: S_{e a} \rightarrow S_{e}$ that collapses the spherical part of $S_{e a}$. This concludes the proof of the lemma.

By Lemma 6.2, and since each $G_{i} \cap G_{j}$ is connected and nonempty, the local linear orderings at the vertices consistently define a relation $\preceq$ on $[n]$ such that for each $i \neq j$, exactly one of $i \preceq j$ and $j \preceq i$ holds. Since every three indices meet at a vertex of $G$, $\preceq$ is also transitive and thus a linear ordering. Finally, as in Section 5, the paths in $\mathbb{Y}$ make sure that $\preceq$ obeys all constraints in $T$.

It is not hard to see that $R_{\text {edge }}, R_{\text {vert }}, 1 / \varepsilon$ can be bounded by a fixed polynomial in $n$ and $D$, of degree independent of $d$, and thus $|Y|=O\left((n D)^{B d}\right)$ for some universal constant $B$. This concludes the proof of Theorem 1.2.

## 7 No Menger-type lemma for low-distortion planar embeddings

Here we prove Theorem 1.3. We set $w:=c_{1} \sqrt{\varepsilon} / k$ for a sufficiently small constant $c_{1}>0$. To construct $\mathbb{X}=\left(X, \rho_{X}\right)$, we consider a drawing of the graph $K_{3,3}$ (any other fixed nonplanar graph would do) as in the left part of Fig. 9 (here $w$ determines the width of the oval-shaped strip).

We place the $n$ points of $X$ on the edges as indicated, with a regular spacing of const $/ n$. The distances under $\rho_{X}$ are given by the Euclidean distances of the points in the considered drawing, except for the points $p_{1}, \ldots, p_{m}, q_{1}, \ldots, q_{m}$ that fall in the dotted rectangle ( $m$ is a fixed fraction of $n$ ). The distance of each $p_{i}$ or $q_{i}$ to a point outside the dotted rectangle is still the Euclidean distance in the drawing, but $\rho_{X}\left(p_{i}, q_{j}\right)$ is given by the Euclidean distances in the planar point configuration shown in Fig. 9 right, where the $p_{i}$ are placed on the top side of a rectangle of height $w$ and the $q_{i}$ on the bottom side.

Arguing as in [20], we can easily check that any embedding of $\mathbb{X}$ in the plane requires distortion $\Omega(n w)=\Omega(\sqrt{\varepsilon} n / k)$. Indeed, given a $D$-Lipschitz noncontracting mapping $f: X \rightarrow \mathbb{R}^{2}$, we extend it to a drawing of $K_{3,3}$ by connecting the images of the points of $X$ by straight segments. In every


Figure 9: The construction for the proof of Theorem 1.3.


Figure 10: Embedding a $k$-point subspace in Theorem 1.3.
drawing of $K_{3,3}$, there are two vertex-disjoint edges $e, e^{\prime}$ that cross. At distance at most $O(D / n)$ from the crossing we find images of two points $x, x^{\prime} \in X, x$ on $e$ and $x^{\prime}$ on $e^{\prime}$. However, we have $\rho_{X}\left(x, x^{\prime}\right) \geq w$, and thus $D=\Omega(n w)$ as claimed.

Next, let $\mathbb{Y}$ be a $k$-point subspace of $\mathbb{X}$. Then there is an $i_{0}$ such that none of $p_{i_{0}}, q_{i_{0}}, p_{i_{0}+1}$, $q_{i_{0}+1}, \ldots, p_{i_{0}+t-1}, q_{i_{0}+t-1}$ falls in $\mathbb{Y}$, where $t=\Omega(n / k)$. We embed $\mathbb{Y}$ as in Fig. 10. A simple calculation shows that the distortion of distances among the $p_{i}$ and $q_{i}$ is $1+O\left((w n / t)^{2}\right) \leq 1+\varepsilon$, while the distortion of distances involving some $p_{i}$ or $q_{i}$ and some other point is $O(1+w) \leq 1+\varepsilon$.

Remarks. A similar example, with correspondingly weaker parameters, can be constructed for embeddings in $\mathbb{R}^{d}$ for every fixed $d \geq 2$, using the Van Kampen-Flores simplicial complexes nonembeddable in $\mathbb{R}^{d}$ (as in [20]) in place of $K_{3,3}$. The details are somewhat technical and there is no substantial new idea involved, so we prefer to omit this part.

Another natural question is, whether an "approximate Menger lemma" for embedding in the plane might hold for some interesting subclass of all metric spaces. Since the construction in the proof of Theorem 1.3 is based on a nonplanar graph, a natural class of interest are planar-graph metrics ${ }^{9}$

For planar-graph metrics we have the following weaker analogy of Theorem 1.3: For all $n$,

[^6]

Figure 11: A construction for planar-graph metrics.
there are planar-graph metrics on $O(n)$ vertices whose embedding in $\mathbb{R}^{2}$ requires distortion $\Omega(\sqrt{n})$, while all subspaces on at most $\sqrt{n}$ points embed with distortion $O(1)$ (we don't get $1+\varepsilon$ as in Theorem 1.3, since planar-graph metrics usually don't embed almost isometrically in the Euclidean plane).

The construction is based on an idea from Bateni et al. [3] (in a simplified form); see Fig. [1], The metric space $\mathbb{X}$ is the vertex set of the depicted graph with the shortest-path metric. The edges drawn with a single line have length 1 and the edges drawn with a double line have length $n^{-1 / 2}$.

The reason for bad embeddability in the plane is as follows. The graph is essentially a subdivision of $K_{4}$ with a long path attached to each vertex. A low-distortion embedding of $\mathbb{X}$ in the plane yields a drawing of $K_{4}$ where vertex-disjoint edges don't cross, and in any such drawing of $K_{4}$ there is a vertex that is not incident to the outer face. Then the path attached to such a vertex doesn't have enough room in the inner face. The details are similar to an argument in [3] and we omit them. It is also easy to see that if we consider a subspace of $\sqrt{n}$ points, then there is a gap in the "ladder" in the bottom part of the graph of length $\Omega(\sqrt{n})$, and using such a gap, the path attached to the central vertex can cross to the outer face, which yields an $O(1)$-embedding.

## 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 lowdimensional 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 48 th 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. 17 th 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.
[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.


[^0]:    *An extended abstract of this paper will appear in Proc. 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS), Philadelphia, 2008.

[^1]:    ${ }^{1} \mathrm{~A}$ set $V$ in a metric space $\left(X, \rho_{X}\right)$ is called $\varepsilon$-dense if for each $x \in X$ there is $v \in V$ with $\rho_{X}(x, v) \leq \varepsilon$.

[^2]:    ${ }^{2}$ 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.

[^3]:    ${ }^{3}$ We recall that two continuous maps $f, g: X \rightarrow Y$ of topological spaces are homotopic, in symbols $f \sim g$, if there exists a continuous map $F: X \times[0,1] \rightarrow Y$ 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 $\left(M, \rho_{M}\right)$ is an $\varepsilon$-net if every two points of $N$ have distance greater than $\varepsilon$ and $N$ is inclusion-maximal with respect to this property; that is, every point of $M$ is at most $\varepsilon$ far from some point of $N$.
    ${ }^{5}$ 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 $\Sigma$ is the union of finitely many simplices, and then we could work with the perhaps more familiar singular or simplicial cohomology.

[^4]:    ${ }^{6}$ See Munkres [22] §72-74; the naturality is explicitly stated there only for polyhedral sets, since it is needed for the proof of the general case, but it follows for the general case as well by the limiting process used in the proof.

[^5]:    ${ }^{7}$ 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.
    ${ }^{8}$ 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.

[^6]:    ${ }^{9}$ A metric space $\left(X, \rho_{X}\right)$ s a planar-graph metric if there is a planar graph $G=(V, E)$ with positive real weights on edges such that $X \subseteq V$ and $\rho_{X}$ is given by the shortest-path metric of $G$.

