

Online Embedding of Metrics

Ilan Newman

Department of Computer Science, University of Haifa, Israel

<http://cs.haifa.ac.il/~ilan/>

ilan@cs.haifa.ac.il

Yuri Rabinovich

Department of Computer Science, University of Haifa, Israel

<http://cs.haifa.ac.il/~yuri/>

yuri@cs.haifa.ac.il

Abstract

We study deterministic online embeddings of metric spaces into normed spaces of various dimensions and into trees. We establish some upper and lower bounds on the distortion of such embedding, and pose some challenging open questions.

2012 ACM Subject Classification Networks → Network algorithms

Keywords and phrases Metric spaces, online embedding

Digital Object Identifier 10.4230/LIPIcs.SWAT.2020.32

Funding *Ilan Newman*: This Research was supported by The Israel Science Foundation, grant number 497/17.

1 Introduction

The modern theory of low-distortion embeddings of finite metrics spaces into various host spaces began to take shape with the appearance of the classical results of Johnson and Lindenstrauss [7]¹ and Bourgain [4]², in the last decades of the 20'th century. It was soon observed that this theory provides powerful tools for numerous theoretical and practical algorithmic problems. Nowadays, it is a mathematically deep and widely applicable developed theory, whose importance to algorithmic design is well recognized.

In this paper we study a relatively neglected aspect of metric embeddings, the *online* embeddings. In this setting, the vertices of the input finite metric space (V, d) are exposed one by one, together with their distances to the previously exposed vertices. Each newly exposed vertex v is mapped to the host space $(\mathcal{H}, d_{\mathcal{H}})$ before the next vertex is exposed, and without altering the embedding of previously exposed vertices. The quality of the resulting embedding $\phi : V \rightarrow \mathcal{H}$ is measured by its *expansion* and *contraction*:

$$\text{expansion : } \max_{v,u \in V} \frac{d_{\mathcal{H}}(\phi(v), \phi(u))}{d(v, u)} \qquad \text{contraction : } \max_{v,u \in V} \frac{d(v, u)}{d_{\mathcal{H}}(\phi(v), \phi(u))}$$

The product of the two is called the (multiplicative) *distortion* of ϕ . The distortion $\text{dist}(d \hookrightarrow d_{\mathcal{H}})$ of embedding (V, d) into $(\mathcal{H}, d_{\mathcal{H}})$ is the minimum possible distortion of any such mapping ϕ . Since usually (and in this paper in particular) the host space is scalable, $\text{dist}(d \hookrightarrow d_{\mathcal{H}})$ can be alternatively defined in the offline setting as the minimum possible expansion over non-contracting mappings, or the minimum possible contraction over non-expanding mappings.

¹ Any n -point Euclidean metric can be efficiently embedded into $\ell_2^{\frac{\log n}{\epsilon^2}}$ with $(1 + \epsilon)$ -distortion.

² Any n -point metric can be efficiently embedded into the Euclidean space of dimension $O(\log n)$.



In the online setting the above three notions may not, and probably do not, coincide. This is so since the scaling of ϕ cannot be performed in the end, after having finished constructing the entire mapping. For the same reason, in the online setting, the *maximum* between the distortion and the contraction may be a more suitable measure of similarity than the multiplicative distortion. Also, the knowledge of $n = |V|$ in advance matters, and may potentially be of help.

In addition to deterministic online embedding algorithms, it is natural to consider *probabilistic* online embeddings against non-adaptive adversary. In this setting, instead of considering the distortion between d and $d_{\mathcal{H}}$ for a fixed embedding, we consider the expected distortion between d and a random embedding that is selected from some pre-designed distribution.

In this paper we focus only on *deterministic embeddings* into the standard normed spaces $\ell_2, \ell_1, \ell_\infty$ of various dimensions, and into trees. Our results clarify what can be achieved in dimension 1, and in dimension exponential in n . What happens in between is a challenging open problem. We also present a lower bound on online embedding of a size- n metrics into ℓ_2 of unbounded dimension.

It is our hope that the findings of the present paper may provide a good starting point for further studies of deterministic online embeddings.

1.1 Previous Work

To the best of our knowledge, the first result about online embedding appeared implicitly in a paper of the authors [8]. The authors show a $(\sqrt{\log n})$ lower bound on the distortion of an *offline* embedding of a shortest path metric of a certain family of series-parallel graphs $\{D_n\}$ into ℓ_2 . Without ever mentioning the term “online”, the proof, in fact, establishes a lower bound of \sqrt{n} on the distortion of an online embedding of the shortest path metric of a certain family of graphs G_{2n} on $2n$ vertices that are subgraphs of D_n . Although [8] received due attention, and its online implications were noticed e.g. by the authors of [6]³, the explicit statement (which was and still remains state of the art in its context) has never appeared in print, and went largely unnoticed. Here we amend this situation (see Section 2).

Another related result appeared in [1] (Th. 3.1) in a rather unrelated context. It claims the following. Let (V, d) an arbitrary metric space with $|V| = n$. Assume that V is exposed in a random uniform order. Then the greedy online algorithm that attaches each new point v to the closest one among the points exposed so far, say u , by an edge of length $d(v, u)$, produces a random dominating tree T so that $E[d_T]$ expands d by $O(n^2)$.

If the order is fixed, a similar but simpler analysis implies that d_T expands d by at most $O(2^n)$. Since this turns out to be rather tight (up to the basis of the exponent) for a deterministic embedding into a tree, we shall discuss it in more details in Section 3. Another somewhat related notion, that we do not discuss here, is that of terminal embedding and using extension techniques [5].

The first (and, to our knowledge, the only) published paper explicitly dedicated to online embeddings is [6]. Observing that a large part of the offline embedding procedure from [2] can be implemented online, the authors in [6] establish quite strong results for probabilistic online embeddings. Most of these results depend on the so called *aspect ratio* Δ of the input metric d – that is, the ratio between the largest and the smallest distance in it. The main results of [6] are as follows (it is assumed that $|V| = n$):

³ It served as partial motivation for their paper. [Private communication]

1. A metric space (V, d) can be probabilistically online embedded into $\ell_p^{\log n \cdot (\log \Delta)^{1/p}}$ with distortion $O(\log n \cdot \log \Delta)$ for any $p \in [1, \infty]$.
For $p = \infty$, (V, d) can also be embedded in $\ell_\infty^{\log^{O(1)} n}$ with distortion $O(\log n \cdot \sqrt{\log \Delta})$.
On the negative side, (V, d) cannot be online embedded into ℓ_2^D with distortion better than $\Omega(n^{1/D-1})$ even when d is $(1 + \epsilon)$ -close to a submetric of ℓ_2^D .
2. A metric space (V, d) can be probabilistically online embedded into a distribution of a non-contracting ultrametrics (and subsequently tree-metrics) with distortion $O(\log n \cdot \log \Delta)$.
On the negative side, (V, d) cannot be probabilistically online embedded into a distribution of a non-contracting ultrametrics with distortion better than $\min\{n, \log \Delta\}$.

A very recent result [3] also discuss probabilistic online embeddings into trees, in the context of terminal-embedding. In particular they also obtain lower bounds on probabilistic embeddings into trees that are parameterized by the aspect ratio.

1.2 Our Results

We are interested in *deterministic* online embeddings into normed spaces, and in particular in the interplay between the distortion and the dimension of the host space. Unlike [6], we seek bounds independent of the aspect ratio.

1. **Embedding Into ℓ_2 :** There exists a family of metrics $\{d_{2n}\}$ such that each d_{2n} (a metric on $2n$ points) requires distortion \sqrt{n} in any deterministic online embedding into ℓ_2 of *any dimension*. The metrics $\{d_{2n}\}$ are the shortest-path metrics of a family $\{G_{2n}\}$ of weighted series-parallel graphs. These metrics are quite simple; e.g., they embed into the line with a constant universally bounded distortion.

By John's Theorem from the theory of finite-dimensional normed spaces, this implies an $\sqrt{n/D}$ lower bound on online embedding of d_{2n} into any normed space of dimension D . Comparing to the corresponding lower bounds of [6] (and ignoring the restrictions on d), we conclude that their result is stronger for $D = 2, 3$, and incompatible or weaker for other dimensions.

Our only positive results for online embedding into ℓ_2 follow from the embedding into the line.

2. **Embedding Into the line, and into trees.** As mentioned above, a simple greedy online embedding algorithms results in a dominating tree whose metric distorts the input metric d_n , on n points, by at most $O(2^n)$. Using a more complicated argument, we show that d can be online embedded into the *line* with distortion $O(n \cdot 6^n)$.

We also establish a lower bound of $\Omega(2^{n/2})$ for online embedding metrics on n points into trees. The “hard” metrics used in the proof are in fact submetrics of a (continuous) cycle, and they embed (offline) in the line with a constant universally bounded distortion.

3. **Distortion and dimension.** What is the smallest dimension D such that d can be embedded into ℓ_∞^D with distortion at most $1 + \epsilon$?⁴ Our first, rather surprising result, is that even a metric d on 4 points requires $D = \Omega(\log \frac{1}{\epsilon})$ in this setting.

On the positive side, we (efficiently) prove that $D = (\frac{4n}{\epsilon})^n$ suffices.

4. **Isometric online embeddings.** We show that size- n tree metrics d (i.e., arbitrary submetrics of the shortest-path metrics of weighted trees) isometrically online embed into ℓ_1^{n-1} . This implies that such d isometrically online embeds into $\ell_\infty^{2^{n-2}}$ for $n > 1$. One conclusion is that if d probabilistically online embeds with expansion a into a distribution of tree metrics supported on at most k trees, then d embeds with the same expansion into $\ell_1^{k(n-1)}$ and $\ell_\infty^{2^{k(n-1)-1}}$.

⁴ It is well known that any metric d of size n can be isometrically embedded into ℓ_∞^{n-1} . Thus, unlike any other ℓ_p , ℓ_∞ is universal in the sense that any metrics is isometric to its submetric.

Open questions:

- Can every n -points metric be online embedded into ℓ_2 (of some dimension) with $\text{poly}(n)$ small distortion?
- Can every n -points metric be online embedded into ℓ_∞^D , where D is at most polynomial in n , and with $\text{poly}(n)$ distortion. In particular this is open for $D = 2$, with a polynomially small distortion?
- At what rate does the quality of the best online embedding (e.g., into ℓ_∞^D) improve when D grows?

2 A lower bound for embeddings into ℓ_2

As mentioned in the previous section, the following theorem is implied by the proof of the main result of [8]:

► **Theorem 1.** *There is a family of metrics $\{d_{2n}\}$ on $2n$ points for any natural n , that requires expansion $\geq \sqrt{n}$ in any non-contracting online embedding into ℓ_2 of any dimension (including infinite dimension).*

Given an online non-contracting embedding algorithm A , the “hard” $\{d_{2n}\}$ is constructed as follows. It will be the shortest-path metric of the following weighted graph G_{2n} . G_2 is simply unit-weighted K_2 . The graph G_{2n+2} is obtained by choosing an edge $e = (v, u)$ of weight 2^{2-n} in G_{2n} , and replacing it by a 4-cycle $v-x-u-y-v$ with edges of weight 2^{3-n} . It remains to specify the edge e . It is proven in [8], inductively, that one of weight 2^{2-n} edges in G_{2n} is expanded by A , by at least \sqrt{n} . Further, this implies that of the four new edges (v, x) , (x, u) , (u, y) and (y, u) , at least one edge must be expanded by A by at least \sqrt{n} . This is the new edge to be chosen by the adversary.

As mentioned above, the metrics d_{2n} are very simple. E.g., it is an easy matter to verify that each d_{2n} (offline) embeds into the line with distortion ≤ 3 , and isometrically embeds into ℓ_1 .

Currently, we do not know how tight is the above bound, and whether is it at all possible to obtain a polynomially small in n (online) distortion for online embedding into ℓ_2 . We do know that it is possible for tree metrics (in view of Theorem 15), and that in general it is at most exponential (by Theorem 10).

3 Online embedding into trees

In tree embeddings we refer to online embeddings that constructs a tree whose vertices may contain *Steiner points*. That is, the constructed tree, besides the points corresponding to the input metric, may contain additional points. At each step, once a new vertex is exposed, the embedding algorithm picks an existing edge of the tree, subdivide it (without changing its total weight) by creating a new Steiner point, and attaches to it the new vertex by a new edge of a corresponding weight. The new edge is always a leaf, except when the weight is 0.

► **Theorem 2.** *Any metric on n points can be deterministically online embedded into a tree with distortion $\leq 2^{n-1} - 1$, even without using Steiner points.*

Proof. Just connect the new point v to the previously exposed point u that is the closest to v in the metric d , by an edge of weight $d(v, u)$.

The analysis is essentially the same as in [1]. Let \tilde{d} denote the tree metric approximating d . Clearly, \tilde{d} is not-contracting. Let α_k denote the its expansion after k steps. Then, $\alpha_2 = 1$, and $\alpha_{k+1} \leq 2\alpha_k + 1$. Indeed, let x be the new point, and assume it was connected to y . Then, for any previously exposed vertex a ,

$$\begin{aligned} \tilde{d}(a, x) &= \tilde{d}(a, y) + d(x, y) \leq \alpha_k \cdot d(a, y) + d(x, y) \leq \alpha_k \cdot (d(a, x) + d(x, y)) + d(x, y) \\ &\leq (2\alpha_k + 1) \cdot d(a, x) \end{aligned}$$

where the penultimate inequality is by the triangle inequality, and the last one follows from the choice of y .

The recursive formula $\alpha_{k+1} \leq 2\alpha_k + 1$ implies that $\alpha_n \leq 2^{n-1} - 1$. \blacktriangleleft

In view of the above theorem, this is rather tight:

► Theorem 3. *There is a class of metrics on n points for which any online embedding algorithm for that class into a tree metric, results in a distortion of at least $2^{(n-4)/2}$. Furthermore, every metric in the class is (offline) embeddable into a line with constant distortion.*

Proof. The metric that will be exposed is a finite submetric of the continuous unit cycle C . Let d_C be the shortest path metric induced on C . We will show that for every $k \geq 1$, the tree that is constructed on the first $4 + 2k$ points distorts d_C on the induced $4 + 2k$ points by no less than 2^k .

Working with the infinite metric space C (instead of a finite metric space), simplifies notions. One may consider the case in which n , the number of points in the metric space that is going to be exposed is given to the algorithm at the beginning. Even then, the following proof works. Moreover, we may restrict ourselves to the finite submetric space of C induced by 2^n points that are uniformly placed on C .

We start with the following simple facts. For two points x, y $P_T(x, y)$ will always refer to the path between x and y in the tree T that would be relevant to the context.

▷ **Claim 4.** Let u_1, u_2, u_3, u_4 be four vertices in a tree $T = (Y, E)$. Let $P(u_i, u_j)$, $1 \leq i < j \leq 4$ be the path in the tree from u_i to u_j . Then either $P(u_1, u_2) \cap P(u_3, u_4) \neq \emptyset$ or $P(u_2, u_3) \cap P(u_1, u_4) \neq \emptyset$.

The lower bound on the distortion will follow from the following claim.

▷ **Claim 5.** Let p, q, r, s be points in a metric space such that $d(p, q), d(r, s) \leq \alpha$ while $d(p, r), d(p, s), d(q, r), d(q, s) \geq \beta$. Assume also that p, q, r, s are embedded into a weighted tree, T , such that $P_T(p, q) \cap P_T(r, s) \neq \emptyset$. Then the tree distance d_T distort d by at least β/α .

Proof. Assume that the expansion is $\gamma \geq 1$. Then in the tree, $d_T(p, q), d_T(r, s) \leq \gamma \cdot \alpha$. In particular it follows that,

$$\sum_{e \in P_1} w(e) + \sum_{e \in P_2} w(e) = d_T(p, q) + d_T(r, s) \leq 2\gamma \cdot \alpha \quad (1)$$

However, as the paths $P_1 = P_T(p, q)$ and $P_2 = P_T(r, s)$ intersect, it follows that their union include the paths $\mathcal{P} = \{P_T(p, s), P_T(s, q), P_T(q, r), P_T(r, p)\}$. More over, every edge in $P_1 \cup P_2$ appears in exactly two of the paths from \mathcal{P} . Hence we conclude that,

$$\sum_{e \in P_T(p, s)} w(e) + \sum_{e \in P_T(s, q)} w(e) + \sum_{e \in P_T(q, r)} w(e) + \sum_{e \in P_T(r, p)} w(e) \leq 2 \sum_{e \in P_1} w(e) + 2 \sum_{e \in P_2} w(e) \quad (2)$$

Where the last inequality may be strict as the edges in the intersection of P_1 and P_2 contribute four times to the right hand side.

Combining Equations (1) and (2), we conclude that for at least one path $P \in \mathcal{P}$, the length of P is at least $\gamma \cdot \alpha$. Assume w.l.o.g that $P = P(p, s)$, then the contraction is at least $\nu = \frac{d(p,s)}{\gamma \cdot \alpha} \geq \frac{\beta}{\gamma \cdot \alpha}$ which implies that the distortion is at least $\nu \cdot \gamma \geq \beta/\alpha$. \triangleleft

We return to the proof of the theorem. Let $d = d_C$ the metric induced by C . For two subsets $S, T \subseteq C$ let $d(S, T) = \min_{s \in S, t \in T} d(s, t)$ (in our applications this minimum will always exist). We fix one point in the cycle $v(0) \in C$ as a reference point, this then defines every point by its distance along the cycle going clockwise. Thus we denote by $v(\alpha)$, $0 \leq \alpha < 1$ the point of length α from $v(0)$ when going along the cycle. For two points $x = v(\alpha), y = v(\beta)$ let $C[x : y]$ be the segment of the cycle on the shortest path between x and y . The mid point of $C[x : y]$ is the point on the geodesic path between x, y which is of equal distance from x and y (this is not well defined only if $d(x, y) = 1/2$, but we will never use this definition in this case).

Fix an online embedding algorithm for n points from C into a tree T , and denote the resulted metric is d_T . The adversary first exposes $u_1 = v(0), u_2 = v(\frac{1}{4}), u_3 = v(\frac{1}{2}), u_4 = v(\frac{3}{4})$. Let T_1 be the tree constructed by the algorithm just after this point. Using fact 4 (with that order on the points) we may assume w.l.o.g that $P(u_1, u_2) \cap P(u_3, u_4) \neq \emptyset$. Note that $d(C[u_1 : u_2], C[u_3 : u_4]) \geq 1/4$.

The adversary will work in phases, each time exposing 2 points. The initial phase (numbered as $k = 0$ exposing 4 points) results in T_0 above. Let $T = T_k$ be the tree that is constructed by the algorithm at steps $k = 1, \dots$ after exposing $4 + 2k$ points. The adversary will always hold two pairs of points that are already exposed $x, y \in C[u_1 : u_2], x', y' \in C[u_3 : u_4]$, maintaining the invariant that: all points in $C[x : y] \setminus \{x, y\}$ and in $C[x' : y'] \setminus \{x', y'\}$ are not exposed, and $P_T(x, y) \cap P_T(x', y') \neq \emptyset$. It will also be the case that $d_C(C[x : y], C[x' : y']) \geq 1/4$, while $d_C(x, y) = d_C(x', y') = 2^{-k-2}$.

For $k = 0$ the points $x = u_1, y = u_2, x' = u_3, y' = u_4$ already comply with the invariants above. Assume that after phase k we already have exposed $4 + 2k$ and the adversary holds x, y, x', y' as required. Then at phase $k+1$ the adversary exposes two new vertices: z that is the mid point in $C[x : y]$ and z_1 that is the mid point of $C[x', y']$. Since $P_T(x, y) \cap P_T(x', y') \neq \emptyset$ then at least one of $P_T(x, z), P_T(z, y)$ intersects $P_T(x', y')$. We replace y with z if $P_T(x, z)$ intersects $P_T(x', y')$, otherwise, we replace x with z . Similarly, we replace either x' or y' with z_1 , so that the resulting two paths still intersect. It is easy to see that the distances are as claimed.

Finally, by Claim 5, applied on x, y, x', y' (in this order) at the end of any phase k , we conclude that the tree distance d_T distort d_C on the four points by at least 2^k .

We end this proof by noting that the actual metric that is exposed is (offline) embeddable into a tree and even into a line with a constant distortion. This can be done by e.g., 'cutting' C at the point $v(7/8)$ and embedding each point x at $v(x)$ in the resulting interval. \blacktriangleleft

4 Embedding into the line

Theorem 3 implies that online embedding of general metrics into the line results in a distortion that in the worst case is at least exponential in the number of points. This is true even for online embedding of tree metrics into the line (using a similar argument as in the proof of Theorem 3). Here we show that any metric (V, d) on n points can be online embedded into the line with distortion that is most exponential (in the number of points exposed so far). We don't assume here that n , the number of points or any upper bound on this number is given in advance.

► **Theorem 6.** *Let (V, d) be a metric, then V can be online embedded into the line (without a priori knowing $n = |V|$), with distortion bounded by $O(n6^n)$.*

Proof. For every point x that is already embedded let $\phi(x)$ be its embedding.

Assume at stage i that x_i is exposed and let z be the closet point to x_i from the previously exposed points, with $d(x_i, z) = d$. Let I be the left most interval of length $3^{-i}d$ that is right of z and is empty of any previously exposed point. We then place x_i in the mid point of I . Note that since there are only $i - 1$ previously exposed points (including z), then there must be such empty interval at distance at most $(i - 2) \cdot 3^{-i}d$.

In the following we call z the *father* of x_i in this embedding and denote it as $father(x_i)$. We are going to bound separately the expansion and the contraction.

Let $\gamma(k)$ denote the bound on the expansion after the k th point is embedded. Bounding $\gamma(k)$ is by a similar argument to the tree embedding. Let x_k be the last point that is embedded, let $z = father(x_k)$ and let y be any previously exposed point. The triangle inequality asserts that,

$$d(y, z) \leq d(y, x_k) + d(x_k, z) \leq 2d(x_k, y) \quad (3)$$

Thus, by the definition of the embedding of x_k and the induction hypothesis,

$$\begin{aligned} |\phi(x_k) - \phi(y)| &\leq |\phi(x_k) - \phi(z)| + |\phi(z) - \phi(y)| \leq \\ &+(k - 1.5) \cdot 3^{-k}d(x_k, z) + \gamma(k - 1)d(z, y) \leq k3^{-k} \cdot d(x_k, y) + \gamma(k - 1)d(z, y) \end{aligned} \quad (4)$$

where the last inequality is by the fact that $d(x_k, z) \leq d(x_k, y)$.

Using equation (3) we get,

$$|\phi(x_k) - \phi(y)| \leq (k3^{-k} + 2\gamma(k - 1)) \cdot d(x_k, y). \quad (5)$$

We get the following recursion on $\gamma(k)$: $\gamma(k) \leq k3^{-k} + 2\gamma(k - 1)$ which implies that $\gamma(k) \leq 3 \cdot 2^k$.

To bound the contraction let a, b be any two exposed points. By the embedding algorithm there are two sequences $z = y_1, y_2, \dots, y_k = a$ and $z = w_1, w_2, \dots, w_\ell = b$ where $y_i = father(y_{i+1})$, $w_i = father(w_{i+1})$ and $\{y_i\}_1^k, \{w_i\}_1^\ell$ are disjoint. A marginal case is when $z = a$ and one of the sequences is empty. The argument for the marginal case will be presented at the end of the proof (after the proof of Claim 8).

Let $\delta_i = d(y_{i-1}, y_i), i = 2, \dots, k$ and $\nu_i = d(w_{i-1}, w_i), i = 2, \dots, \ell$. For any point x let $order(x) = \ell$ if $x = x_\ell$ namely, x is the ℓ exposed point (do not confuse the $order(x)$ with its location in the sequences of y_i 's or w_i 's).

Let $D = \max\{d(x, father(x))\}$ where x ranges over all points except z in the two sequences above, and assume w.l.o.g. that the last exposed point among the two sequences, x_j , for which $d(x_j, father(x_j)) = D$ is y_i (namely, that the maximum is achieved in the sequence that corresponds to a). Let $s = order(y_i)$.

By our algorithm $y_i = x_s$ is embedded in the middle of an empty interval I of size $3^{-s}D$.

We use the following claims.

For $j = 1, \dots, k - i$ let $r_j = order(y_{i+j}) - s$.

► **Claim 7.** For any $j \geq 1$, y_{i+j} is embedded inside I and $0 \leq \phi(y_{i+s}) - \phi(y_i) \leq \frac{|I|}{2} \cdot (1 - 2^{-r_j})$.

Proof. We first describe the situation for the case $j = 1$. The case for larger j is similar. Let $r_1 = order(y_{i+1}) = r$. Namely, $r - 1$ points $x_{s+1}, \dots, x_{s+r-1}$ are exposed after y_i and before y_{i+1} .

Recall that at time i when y_i is embedded, I is empty, and y_i is placed in the middle of I splitting I into two empty intervals I_L, I_R of size $|I|/2$ each. According to the algorithm y_{i+1} needs to be embedded in the middle of an interval of size $\alpha = 3^{-(r+s)} \cdot \delta_{i+1} \leq 3^{-r} \cdot 3^{-s} \cdot D \leq 3^{-r} \cdot |I|$ that is empty at time $s+r$. If $r=1$, namely if y_{i+1} is exposed right after $y_i = x_s$, then obviously there is a α -size empty interval right of $\phi(y_i)$, as I_R is empty at this point and $|I|/2 > \alpha$. Generally, for $r > 1$, some of the $r-1$ points that are exposed between y_i and y_{i+1} may occupy parts of I forcing y_{i+1} to be embedded further to the right.

Each time a point $x \neq y_{i+1}$ is placed in I_R it must be in the middle of an empty interval splitting the right empty interval of I_R into two empty subintervals, hence leaving an empty interval of at least half the size at the right of I_R . Hence after placing at most $r-1$ such points, there will still be an empty interval of size $|I_R|/2^{r-1} > \alpha$. Hence there is a suitable empty interval for y_{i+1} in I_R and it follows that $\phi(y_{i+1}) - \phi(y_i) \leq |I|/2 - 2^{-r+1}|I_R| + \alpha/2 \leq \frac{|I|}{2} \cdot (1 - 2^{-r})$.

In the general case for y_{i+j} the argument is identical except that $r_j - 1$ points might have been embedded into I_R before y_{i+j} . \triangleleft

\triangleright **Claim 8.** $|\phi(a) - \phi(b)| \geq 2^{-(r_k+1)}|I|$.

Proof. Claim 7 asserts that $a = y_k$ is embedded inside I to the right of y_i and $\phi(a) - \phi(y_i) \leq \frac{|I|}{2} \cdot (1 - 2^{-r_k})$.

We now consider the place where b is embedded. Let t be the largest so that w_t is exposed before y_i . Again, since I is empty when y_i is exposed, w_t must be embedded to the right or to the left of I . If w_t is embedded to the right of I then b , that is embedded right of w_t , is right of I and hence $\phi(b) - \phi(a) \geq \frac{|I|}{2} \cdot 2^{-r_k}$ implying the claim.

On the other hand, if w_t is embedded to the left of I , then by a similar calculation that is done in Claim 7, all points w_{t+j} are embedded at most at distance $\frac{|I|}{2}$ from w_t . (Since they all are exposed after y_i and in particular have $\nu_{t+j} < D$ and $\text{order}(w_{t+j}) > s$). Since $\phi(y_i) - \phi(w_t) \geq |I|/2$ by the assumption that w_t is left of I , we conclude that $\phi(a) - \phi(b) \geq \phi(a) - \phi(y_i) + \phi(y_i) - \phi(w_t) \geq \frac{|I|}{2} \cdot 2^{-r_k}$ in this case too.

This completes the proof of Claim 8. \triangleleft

Now with this lower bound on $|\phi(a) - \phi(b)|$, to bound the contraction it is enough to upper bound $d(a, b)$. Indeed,

$$d(a, b) \leq \sum_1^k d(y_i, y_{i-1}) + \sum_1^\ell d(w_i, w_{i-1}) \leq n \cdot D \implies$$

$$\frac{d(a, b)}{|\phi(a) - \phi(b)|} \leq \frac{nD}{2^{-(r_k+1)}|I|} \leq n2^{r_k+1} \cdot 3^s \leq n3^n$$

To complete the proof, consider the case where one of the sequences is empty. Namely w.l.o.g $z = b$. If b is exposed before y_i then we are at the same situation as in Claim 8, implying the same lower bound on $\phi(y_i) - \phi(b)$. If $z = y_i$ then Claim 7 asserts that $a = y_k$ is embedded in I , and $\phi(a) - \phi(b) \geq \sum_{j=2}^k (\phi(y_j) - \phi(y_{j-1})) \geq \sum_2^k \delta_j \cdot 3^{-r_j}/2$, where the inequality is by the fact that for every j , y_j is embedded to the right of y_{j-1} at distance at least $3^{-r_j} \delta_j/2$. But this last expression is at least $\frac{3^{-n}}{2} \sum_2^k \delta_j \geq 3^{-n} d(a, b)/2$ proving that in this case the contraction is bounded by $2 \cdot 3^n$ as well.

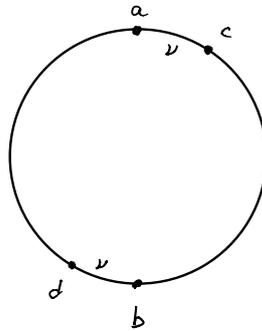
This completes the proof of the Theorem. \blacktriangleleft

5 Online embedding into ℓ_∞ with $(1 + \epsilon)$ distortion

It is well known that any metric on n points can be (offline) embedded isometrically online into ℓ_∞^{n-1} . It therefore comes as a surprise that even the metrics on 4 points cannot be online isometrically embedding into ℓ_∞ of any finite dimension. This will be proven using a special class of 4-points metrics that are submetrics of a continuous cycle.

► **Theorem 9.** *There exists μ on four points for which any online embedding into ℓ_∞^D incurs a distortion $(1 + \Omega(1/4^{2D}))$. Consequently, to ensure a distortion $(1 + \epsilon)$, one needs $\Omega(\log(1/\epsilon))$ dimensions.*

Proof. The metrics under discussion look as depicted in Figure 1. They are all defined by 4 points on the cycle whose circumference is of size 2. All these metrics contain two antipodal points a, b as in Figure 1, that are exposed first and with $d(a, b) = 1$. After a, b are embedded the next two points c, d are exposed. c, d are also antipodal and are defined by the distance $d(a, c) = \nu$.



■ **Figure 1** The metric μ_4 . a, b are exposed first and then ν is set (defining c, d).

We view ϕ as D online non-contacting embeddings into the line $\Phi = \{\psi_1, \dots, \psi_d\}$, where $\psi_i : \{a, b, c, d\} \rightarrow \mathbb{R}, i = 1, \dots, D$. The adversary reveals first the antipodal points a, b . It will then choose ν appropriately, and reveal the corresponding antipodal points c, d .

Let $\delta = 4^{-(D+1)}$. Assume that a, b are exposed and w.l.o.g., $0 = \psi(a) \leq \psi(b)$ for every $\psi \in \Phi$. Moreover, we may assume that every ψ_i is not expanding by more than $1 + \delta$, as otherwise we are done. Hence by multiplying by $\frac{1}{1+\delta}$ we may assume that every $\psi \in \Phi$ is non-expanding.

We partition the interval $[0, 1]$ into $d + 1$ sets $B_0 = (1 - 4^{-D}, 1]$, $B_i = (1 - 4^{i-D}, 1 - 4^{i-D-1}]$, $i = 1, \dots, D - 1$, and finally $B_D = [0, 3/4]$.

After exposing a, b , $\psi(b)$ is determined for every $\psi \in \Phi$. This partitions Φ into $D + 1$ classes $\tilde{B}_0, \dots, \tilde{B}_d$ by letting $\psi \in \tilde{B}_i$ if $\psi(b) \in B_i$. Hence for some $i \in [D + 1]$, $\tilde{B}_i = \emptyset$. Fix such an i , and set $\nu = 4^{i-D}/3$, which define c and d .

Consider first the case $j > i$, then for $\psi \in \tilde{B}_j$, $\psi(b) < 1 - 4^{i-D} \leq 1 - 3\nu$.

Since ψ is non-expanding, $\mu(a, c) = \nu$ implies that $\psi(c) \in I_c = [-\nu, \nu]$. Similarly, $\psi(d) \in I_d = [1 - \psi(b) - \nu, 1 - \psi(b) + \nu]$. But $\max\{|y - x|, y \in I_d, x \in I_c\} \leq 1 - \psi(b) + 2\nu$. Hence the contraction of $\mu(c, d)$ in this case is at least $\frac{1}{1 - \psi(b) + 2\nu} \geq \frac{1}{1 - 4^{i-D-1}}$.

On the other hand, for $\psi \in \tilde{B}_j$ and $j < i$, $\psi(b) \geq 1 - 4^{i-1-D}$. But then since ψ is non-expanding, it follows that $\psi(d) \leq 1 - \nu$ (on account of $\mu(a, d)$), and $\psi(c) \geq \psi(b) - (1 - \nu) = \nu - (1 - \psi(b))$. It follows that the $\psi(d) - \psi(c) \leq 1 - 2\nu + (1 - \psi(b)) \leq 1 - \frac{2 \cdot 4^{i-D}}{3} + 4^{i-1-D} \leq 1 - 4^{i-D-1}$. Hence each such ψ contracts $\mu(c, d)$ by at least $\frac{1}{1 - 4^{i-D-1}}$.

32:10 Online Embedding of Metrics

We conclude that for the above setting of ν , every ψ contracts $\mu(c, d)$ by at least $\frac{1}{1-4^{-(D+1)}}$. Recall that we have started the proof by multiplying Φ by $\frac{1}{1+\delta}$. Hence the distortion of Φ is at least $\frac{1}{(1+\delta)(1-4^{-(D+1)})} \geq \frac{1}{1-4^{-2(D+1)}}$. \blacktriangleleft

Let us note, without providing more details here, that for metrics μ as above the about lower bound is tight and cannot be strengthened. We conjecture that every metric on 4 points can be online embedded in ℓ_∞^D with distortion $1 + \exp(-D)$.

Next we show a result complementary to Theorem 9. That is – any metric can be embedded into ℓ_∞ with distortion arbitrary close to 1, using large enough dimension.

► **Theorem 10.** *Let (V, d) be any metric space, $\epsilon > 0$ an arbitrary small constant, and let $n = |V|$ (or an upper bound on $|V|$) be known in advance. Then using $(\frac{2n}{\ln(1+\epsilon)})^{2n}$ coordinates one can embed d online in a 1-Lipschitz embedding with contraction bounded by $(1 + \epsilon)$.*

Proof. The proof idea is to approximate the universal embedding of d into ℓ_∞^∞ .

Let $V' \subseteq V$ be points from V and $z \in V'$ one fixed point. Let ϕ_z^U (U here stands for *universal*) be the following embedding of V' into the line. $\phi_z^U(z) = 0$ and for every $x \in V'$, $\phi_z^U(x) = d(z, x)$.

The following claim is standard and immediate from the definition.

► **Claim 11.** ϕ_z^U is a 1-Lipschitz embedding of d and for every $x \in V'$ it does not distort $d(z, x)$.

It follows then that if for every $z \in V$ there is a coordinate on which ϕ_z^U is realized, then the embedding is an isometry of d in ℓ_∞ . Hence, ℓ_∞ of dimension $n - 1$ is universal for any metric space on n points.

Here we will online approximate each of these coordinates by preparing in advance a coordinate (in fact a collection of coordinates) for each possible new z . Suppose that the points a_1, \dots, a_k are exposed (not necessarily in that order), and that for a new point z , there is a line (coordinate) in which a_1, \dots, a_k are embedded in increasing order that is consistent with the distance order to z . We show below that under some restrictions on the embedding of these first k points, there is an augmentation of this embedding to any possible consistent z .

Since there are only finite number of ordering of the first k points with respect to their distance from z , we will prepare in advance a line (in fact, a set of lines) for every possible ordering. This will allow us to embed every possible new coming z .

We start with the following Claim asserting that a consistent ordering on the line can be augmented to a new point, under some suitable restriction.

► **Claim 12.** Let $k < n$ and $\{a_1, \dots, a_k\} \subset V$ a set of arbitrary points. Let $\delta > 0$ be a small constant and $1 \leq \ell_i \leq \frac{n}{\delta}$, $i = 2, \dots, k$, a sequence of integers. Let $\tilde{\phi}$ a fixed 1-Lipschitz embedding of a_1, \dots, a_k on the line with $\tilde{\phi}(a_{i+1}) = \tilde{\phi}(a_i) + \frac{\ell_{i+1}-1}{n} \cdot \delta d(a_i, a_{i+1})$, $i = 1, \dots, k-1$. Then for any $z \in V \setminus \{a_1, \dots, a_k\}$ such that

1. $d(z, a_i) \leq d(z, a_{i+1})$, $i = 1, \dots, k-1$.
2. For every $i = 1, \dots, k-1$,

$$\frac{(\ell_{i+1} - 1)\delta}{n} \cdot d(a_{i+1}, a_i) < d(z, a_{i+1}) - d(z, a_i) \leq \frac{\ell_{i+1}\delta}{n} \cdot d(a_{i+1}, a_i)$$

The augmentation of $\tilde{\phi}$ with $\tilde{\phi}(z) = \tilde{\phi}(a_1) - d(z, a_1)$ is 1-Lipschitz, contracting the distances $d(z, a_i)$, $i = 1, \dots, k$ by at most $e^{2\delta}$.

We call such $\tilde{\phi}$ “additive shifted approximation” of ϕ_z^U .

Proof. We will show that for any possible z for which the premises of the Claim hold, the augmented embedding $\tilde{\phi}$ above is an approximation of ϕ_z^U .

Fix a_1, \dots, a_k and let $\tilde{\phi}$ be the augmented embedding for an arbitrary z for which the assumptions of the claim hold. Let $c_i = e^{2\delta i/n}$. We prove by induction on $i = 1, \dots, k$ that

$$c_i d(z, a_i) \leq \tilde{d}(z, a_i) = \tilde{\phi}(a_i) - \tilde{\phi}(z) \leq d(z, a_i). \quad (6)$$

Which will assert that the embedding is 1-Lipschitz and with contraction of $d(z, a_k)$ which is at most c_k .

Indeed for $i = 1$, $\tilde{d}(z, a_1) = d(z, a_1)$.

Assume that Equation (6) is already proved for i and let us prove it for $i + 1$.

By definition, $\tilde{d}(z, a_{i+1}) = \tilde{d}(z, a_i) + \frac{(\ell_{i+1}-1)\delta}{n} \cdot d(a_i, a_{i+1}) \leq d(z, a_i) + d(z, a_{i+1}) - d(z, a_i) \leq d(z, a_{i+1})$, where the inequality follows from the induction hypothesis that $\tilde{d}(z, a_i) \leq d(z, a_i)$ and the 2nd condition in the claim. This establish the fact that the mapping is non-expanding.

The contraction c_{i+1} is bounded by (again using condition 2 of the claim)

$$\frac{d(z, a_{i+1})}{\tilde{d}(z, a_{i+1})} \leq \frac{d(z, a_i) + \ell_{i+1} \cdot \frac{\delta}{n} \cdot d(a_{i+1}, a_i)}{\tilde{d}(z, a_i) + (\ell_{i+1} - 1) \cdot \frac{\delta}{n} \cdot d(a_{i+1}, a_i)} \quad (7)$$

Next we note that by assumption $d(z, a_{i+1}) - d(z, a_i) \leq \ell_{i+1} \cdot \frac{\delta}{n} \cdot d(a_i, a_{i+1})$, while by the triangle inequality $d(a_i, a_{i+1}) \leq d(z, a_{i+1}) + d(z, a_i)$. Combining these two conditions implies that,

$$d(z, a_i) \geq \frac{1}{2} \cdot (1 - \ell_{i+1} \cdot \frac{\delta}{n}) d(a_i, a_{i+1}) \quad (8)$$

Plugging this in Equation (7), using that by induction $\tilde{d}(z, a_i) \geq d(z, a_i)/c_i$, we get,

$$\begin{aligned} c_{i+1} &\leq c_i \cdot \frac{d(z, a_i) + \ell_{i+1} \cdot \frac{\delta}{n} \cdot d(a_i, a_{i+1})}{d(z, a_i) + c_i \cdot (\ell_{i+1} - 1) \frac{\delta}{n} \cdot d(a_i, a_{i+1})} \\ &\leq c_i \cdot \frac{d(z, a_i) + \ell_{i+1} \cdot \frac{\delta}{n} \cdot d(a_i, a_{i+1})}{d(z, a_i) + (\ell_{i+1} - 1) \frac{\delta}{n} \cdot d(a_i, a_{i+1})} \leq \\ &c_i \cdot \left(1 + \frac{\frac{\delta}{n} \cdot d(a_i, a_{i+1})}{d(z, a_i) + (\ell_{i+1} - 1) \frac{\delta}{n} \cdot d(a_i, a_{i+1})}\right) \end{aligned}$$

Using again equation (8) for $d(z, a_i)$ in the denominator we get,

$$c_{i+1} \leq c_i \cdot \left(1 + \frac{\frac{2\delta}{n}}{1 + (\ell_{i+1} - 2) \frac{\delta}{n}}\right)$$

The last expression is the largest when $\ell_{i+1} = 1$ for which we get $c_{i+1} \leq c_i \cdot (1 + 2\frac{\delta}{n})$ and the claim follows. \triangleleft

To complete the proof, we will show that when exposing the $k + 1$ point z , having already embedded the first k points in a 1-Lipschitz embedding, there is a coordinate for which the points are placed along the line according to non-decreasing order of the distances from z , and with pairwise distances $\tilde{\phi}(a_{i+1}) - \tilde{\phi}(a_i)$ that are approximated as in the second item of the Claim.

Indeed fix $\epsilon > 0$ and let $\delta < \frac{1}{2} \ln(1+\epsilon)$, namely, such that $e^{2\delta} < 1+\epsilon$. For every $i = 2, \dots, n$ (note that n needs to be known in advance), we will make sure that after exposing the k th point, for any permutation $\pi \in \mathcal{S}_k$, and any setting of $(s_2, \dots, s_k) \in \{0, 1, \dots, \lfloor \frac{n}{\delta} \rfloor\}^k$ we have

32:12 Online Embedding of Metrics

at least one coordinate, namely a line and a 1-Lipschitz embedding $\tilde{\psi}$ of a_i , $i = 1, \dots, k$ such that the points appear in order as specified by the permutation π , and in which $\psi(a_{\pi(i+1)}) - \psi(a_{\pi(i)}) = s_{i+1} \cdot \frac{\delta}{n} d(a_{\pi(i+1)}, a_{\pi(i)})$. Note, that we assume inductively that for all coordinates (that is, lines) these embedding are required to be 1-Lipschitz.

Having such situation, once z is exposed, assume that a_1, \dots, a_k is a re-enumeration of the points by their distance to z , and let $s_i, i = 2, \dots, k$ be such that $\frac{(\ell_{i+1}-1)\delta}{n} \cdot d(a_{i+1}, a_i) < d(z, a_{i+1}) - d(z, a_i) \leq \frac{\ell_{i+1}\delta}{n} \cdot d(a_{i+1}, a_i)$. By the triangle inequality, such sequence $s_i, i = 2, \dots, k$ exists, and hence by assumption there is a line and an embedding of a_1, \dots, a_k for which the conditions of Claim 12 hold with respect to z . Then Claim 12 asserts that z can be placed in that line and the corresponding augmented embedding that is a “online” shifted additive approximation of ϕ_z^U is 1-Lipschitz, and with bounded contraction as needed.

For every other line we only need to place z so that it will remain 1-Lipschitz. Indeed since the embedding of a_1, \dots, a_k is 1-Lipschitz on the line, it is folklore that z can be placed too, so to result in a 1-Lipschitz embedding (e.g., by using the Helly property for the line).

One last thing to observe, is that in order to take care for future points, we also need to make sure that any relevant order of $\nu \in \mathcal{S}_{k+1}$, namely including z and any set of relevant integers s_2, \dots, s_{k+1} is also realized. To do this, for any possible such set of integers, and any permutation $\pi \in \mathcal{S}_k$, we prepare enough identical copies of the same embedding of $\{a_1, \dots, a_{k+1}\} \setminus \{z\}$ so to be able to place z in any of the corresponding $k+1$ intervals, and in any of the possible $\frac{n}{\delta}$ placers in the interval, so to cover all possible sequences s_2, \dots, s_{k+1} . Thus to estimate the required dimension, let $f(k)$ denote the number of lines needed for step k , in which we assume that every order and every sequence of numbers s_1, \dots, s_k is realized. By the previous discussion we need $f(k+1) = f(k) \cdot (k+1) \cdot \frac{n}{\delta}$ lines for step $k+1$. Since $f(1) = 1$, the recurrence implies that $f(n) = n! \cdot (\frac{n}{\delta})^{n-1}$. We conclude that the dimension needed for the online embedding above is at most $f(n) < (\frac{n}{\delta})^{2n}$ in order to embed any n -point metric. ◀

6 Isometric online embeddings

We conclude with a number of remarks on isometric embeddings. First, observe that the tree metrics d_T (like the Euclidean metrics) are essentially rigid, i.e., there exists an essentially unique (minimal) weighted tree T^* with Steiner points realizing d_T as its submetric. Moreover, this T^* can be constructed in an online manner, regardless of the order of exposure.

► **Theorem 13.** *Every tree metric d_T can be isometrically embedded into a metric of a weighted tree T^* (using Steiner points). The knowledge of n is not required.*

Skipping the details, the embedding at each step, given a new point x , introduces a new Steiner point y_x into the tree constructed so far, and attaches x to y_x by a new edge of weight w_x . Interestingly, the use of Steiner points is essential:

► **Lemma 14.** *There is a family of tree metrics that suffer an exponential distortion in every online embedding into a tree that does not use Steiner points, even when n is known in advance.*

The proof of Lemma 14 is based on the same ideas as in the proof of Theorem 3. We omit further details from this draft.

Next, we claim that tree metrics isometrically embed online into ℓ_1 .

► **Theorem 15.** *Every tree metric on n points is isometrically online embeddable into ℓ_1^{n-1} .*

We do not present a proof of this theorem here. The general idea is to follow the isometric embedding into a weighted tree T^* as outlined above. The invariant property of the embedding is that the adjacent points in the tree will differ by a simple coordinate. Thus, adding a new Steiner point will require no increase in dimension. Attaching a new point x to the corresponding Steiner point y_x involves taking the vector representing y_x , and adding to it a new coordinate with value w_x . The vectors constructed so far are assumed to have value 0 on this coordinate.

Our last remark is that if a metric d is online-embeddable into ℓ_1 of an a priori known dimension $D(n)$, then d can also be online embedded into ℓ_∞ of dimension $2^{D(n)-1}$. This is so, since the embeddings: $x = (x_1, \dots, x_D) \rightarrow (\langle x, \epsilon^1 \rangle, \dots, \langle x, \epsilon^{2^D} \rangle)$, where ϵ^i range over all possible choices of D -dimensional ± 1 vectors, is an isometry from ℓ_1^D into $\ell_\infty^{2^D}$, and, moreover, it is online constructible. 2^D can be improved to 2^{D-1} by fixing the first sign to be 1. Thus, e.g.,

► **Theorem 16.** *Every tree metric on n points is isometrically online embeddable into $\ell_\infty^{2^{n-2}}$.*

References

- 1 Sanjeev Arora, László Lovász, Ilan Newman, Yuval Rabani, Yuri Rabinovich, and Santosh S. Vempala. Local versus global properties of metric spaces. *SIAM J. Comput.*, 41(1):250–271, 2012. doi:10.1137/090780304.
- 2 Yair Bartal. On approximating arbitrary metrics by tree metrics. In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 161–168. ACM, 1998. doi:10.1145/276698.276725.
- 3 Yair Bartal, Nova Fandina, and Seeun William Umboh. Online probabilistic metric embedding: A general framework for bypassing inherent bounds. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1538–1557, 2020. doi:10.1137/1.9781611975994.95.
- 4 Jean Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel J. Math.*, 52:46–52, 1985. doi:10.1007/BF02776078.
- 5 Michael Elkin, Arnold Filtser, and Ofer Neiman. Terminal embeddings. *Theor. Comput. Sci.*, 697:1–36, 2017. doi:10.1016/j.tcs.2017.06.021.
- 6 Piotr Indyk, Avner Magen, Anastasios Sidiropoulos, and Anastasios Zouzias. Online embeddings. In Maria J. Serna, Ronen Shaltiel, Klaus Jansen, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010, Barcelona, Spain, September 1-3, 2010. Proceedings*, volume 6302 of *Lecture Notes in Computer Science*, pages 246–259. Springer, 2010. doi:10.1007/978-3-642-15369-3_19.
- 7 William B. Johnson and Joram Lindenstrauss. Extensions of 1-lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26:198–206, 1984.
- 8 Ilan Newman and Yuri Rabinovich. A lower bound on the distortion of embedding planar metrics into euclidean space. *Discrete Comput. Geom.*, 29:77–81, 2002. doi:10.1007/s00454-002-2813-5.