Discrete & Computational Geometry
https://doi.org/10.1007/s00454-023-00565-2

Labelings vs. Embeddings: On Distributed and Prioritized Representations of Distances

Arnold Filtser[1] · Lee-Ad Gottlieb[2] · Robert Krauthgamer[3]

Received: 15 December 2021 / Revised: 1 December 2022 / Accepted: 25 February 2023

Abstract

We investigate for which metric spaces the performance of distance labeling and of ℓ_∞-embeddings differ, and how significant can this difference be. Recall that a distance labeling is a distributed representation of distances in a metric space (X, d), where each point $x \in X$ is assigned a succinct label, such that the distance between any two points $x, y \in X$ can be approximated given only their labels. A highly structured special case is an embedding into ℓ_∞, where each point $x \in X$ is assigned a vector $f(x)$ such that $\|f(x) - f(y)\|_\infty$ is approximately $d(x, y)$. The performance of a distance labeling or an ℓ_∞-embedding is measured via its distortion and its label-size/dimension. We also study the analogous question for the prioritized versions of these two measures. Here, a priority order $\pi = (x_1, \dots, x_n)$ of the point set X is given, and higher-priority points should have shorter labels. Formally, a distance labeling has prioritized label-size $\alpha(\cdot)$ if every x_j has label size at most $\alpha(j)$. Similarly, an

Editor in Charge: Kenneth Clarkson

A preliminary version of this paper appeared in the proceedings of SODA 2020 [13]. This version contains the proofs of Theorems 5.1 and 5.2, omitted from the preliminary version. Arnold Filtser: This research was supported by the Israel science foundation (Grant No. 1042/22). Lee-Ad Gottlieb: Work partially supported by ISF grant 1602/19. Robert Krauthgamer: Work partially supported by ONR Award N00014-18-1-2364, the Israel Science Foundation Grant # 1086/18, and a Minerva Foundation grant.

Arnold Filtser
arnold.filtser@biu.ac.il

Lee-Ad Gottlieb
leead@ariel.ac.il

Robert Krauthgamer
robert.krauthgamer@weizmann.ac.il

1 Computer Science Department, Bar Ilan University, 5290002 Ramat-Gan, Israel

2 School of Computer Science, Ariel University, 40700 Ariel, Israel

3 Faculty of Mathematics and Computer Science, The Weizmann Institute of Science, 76100 Rehovot, Israel

Published online: 15 September 2023

embedding $f: X \to \ell_\infty$ has prioritized dimension $\alpha(\cdot)$ if $f(x_j)$ is non-zero only in the first $\alpha(j)$ coordinates. In addition, we compare these prioritized measures to their classical (worst-case) versions. We answer these questions in several scenarios, uncovering a surprisingly diverse range of behaviors. First, in some cases labelings and embeddings have very similar worst-case performance, but in other cases there is a huge disparity. However in the prioritized setting, we most often find a strict separation between the performance of labelings and embeddings. And finally, when comparing the classical and prioritized settings, we find that the worst-case bound for label size often "translates" to a prioritized one, but also find a surprising exception to this rule.

Keywords Metric embedding · Distance labeling · ℓ_∞

Mathematics Subject Classification 30L05 · 46B85 · 05C12 · 05C78 · 68R12

1 Introduction

It is often useful to succinctly represent the pairwise distances in a metric space (X, d) in a distributed manner. A common model, called *distance labeling*, assigns to each point $x \in X$ a label $l(x)$, such that some algorithm $\mathcal{A}$ (oblivious to (X, d)) can compute the distance between any two points $x, y \in X$ given only their labels $l(x), l(y)$, i.e., $\mathcal{A}(l(x), l(y)) = d(x, y)$. The goal is to construct a labeling whose label-size, defined as $\max_{x \in X} |l(x)|$, is small. For general n-point metric spaces, Gavoille et al. [14] constructed a labeling scheme with label size of $O(n)$ words, and also proved this bound to be tight.[1]

To obtain smaller label size, one often considers algorithms that approximate the distances. A distance labeling is said to have *distortion* $t \geqslant 1$ if

$$\forall\, x, y \in X, \qquad d(x, y) \leqslant \mathcal{A}(l(x), l(y)) \leqslant t \cdot d(x, y).$$

While the lower bound of [14] holds even for distortion $t < 3$, Thorup and Zwick [26] constructed a labeling scheme with distortion $2t - 1$ and label size $O(n^{1/t} \log n)$ for every integer $t \geqslant 2$. These bounds are almost tight (assuming the Erdős girth conjecture), and demonstrate that for distortion $O(\log n)$, label size $O(\log n)$ is possible.[2]

From an algorithmic viewpoint, there is a significant advantage to labels possessing additional structure, for example labels that are vectors in a normed space. This structure can lead to improved algorithms, for example nearest neighbor search [5, 15]. A natural candidate for vector labels is the ℓ_∞ space, since every finite metric space embeds into it isometrically (i.e., with no distortion). As such isometric embeddings

[1] We measure size in words to avoid issues of bit representation. In the common scenario where distances are polynomially-bounded integers, every word has $O(\log n)$ bits, where $n = |X|$. The bounds in [14] are given in bits and are for unweighted graphs. Nevertheless, once we consider weighted graphs, $\Theta(n)$ words are sufficient and necessary for exact distance labeling, see Theorem 2.1.

[2] A much earlier technique to construct labeling scheme with distortion $O(\log n)$ is Bourgain's [6] embedding into $O(\log^2 n)$-dimensional ℓ_2, providing $O(\log^2 n)$ label size.

require $\Omega(n)$ dimensions [19], one may consider instead embeddings with small distortion. Formally, an embedding $f : X \to \ell_\infty$ is said to have *distortion* $t \geqslant 1$ if

$$\forall x, y \in X, \qquad d(x, y) \leqslant \|f(x) - f(y)\|_\infty \leqslant t \cdot d(x, y).$$

Matoušek [20] showed that for every integer $t \geqslant 2$, every metric space embeds with distortion $2t - 1$ into ℓ_∞ of dimension $O(n^{1/t}\, t \log n)$ (which again is almost tight assuming the Erdős girth conjecture). For distortion $O(\log n)$, Abraham et al. [1] later improved the dimension to $O(\log n)$.

In this paper, we take the perspective that ℓ_∞-embeddings are a particular form of distance labelings, and study the trade-offs these two models offer between distortion and dimension/label-size. While the inherent structure of ℓ_∞-embeddings makes them preferable, one may suspect that their additional structure precludes the tight trade-off achieved using generic labelings. Yet we have seen that for general metric spaces, the performance of ℓ_∞-embeddings is essentially equivalent to that of generic labelings. This observation motivates us to consider more restricted input metrics, such as ℓ_p spaces, planar graph metrics, and trees. The central question we address is the following.

Question 1.1 In what settings are generic distance labelings more succinct than ℓ_∞-embeddings, and how significant is the gap between them?

Priorities. Elkin et al. [9] introduced the problems of *prioritized distortion* and *prioritized dimension*; they posited that some points have higher importance or priority, and it is desirable that these points achieve improved performance. Formally, given a priority ordering $\pi = \{x_1, \ldots, x_n\}$ on the point set X, we say that embedding $f : X \to \ell_\infty$ possesses *prioritized contractive distortion*[3] $\alpha : \mathbb{N} \to \mathbb{N}$ (w.r.t. π) if

$$\forall\, j < i, \qquad \frac{d(x_j, x_i)}{\alpha(j)} \leqslant \|f(x_j) - f(x_i)\|_\infty \leqslant d(x_j, x_i). \tag{1.1}$$

Prioritized distortion is defined similarly for distance labeling. Furthermore, we say that a labeling scheme has prioritized label-size $\beta : \mathbb{N} \to \mathbb{N}$, if every x_j has label length $|l(x_j)| \leqslant \beta(j)$. We say that embedding $f : X \to \ell_\infty$ has prioritized dimension β if every $f(x_j)$ is non-zero only in the first $\beta(j)$ coordinates (i.e., $f_i(x_j) = 0$ whenever $i > \beta(j)$). Here too ℓ_∞-embeddings are a more structured case of labelings, and we again ask what are the possible trade-offs and how these two compare. It is worth noting that the priority functions α, β are defined on all of $\mathbb{N}$ and apply when embedding every finite metric space; in particular, they are not allowed to depend on $n = |X|$. Analogously to Question 1.1, we may also ask here about prioritized label size and dimension:

[3] In the original definition of prioritized distortion in [9], the requirement of equation (1.1) is replaced by the requirement $d(x_j, x_i) \leqslant \|f(x_j) - f(x_i)\|_\infty \leqslant \alpha(j) \cdot d(x_j, x_i)$. We add the word *contractive* to emphasize this difference. Prioritized contractive distortion is somewhat weaker in that it does not imply scaling distortion (see Sect. 1.2).

Question 1.2 In what settings are distance labelings with prioritized label size more succinct than ℓ_∞-embeddings with prioritized dimension, and how significant is the gap between them?

In many embedding results, the (worst-case) distortion is a function of the size of the metric space $n = |X|$. Elkin et al. [9] demonstrated a general phenomenon: Often a worst-case distortion $\alpha(n)$ can be replaced with a prioritized distortion $\widetilde{O}(\alpha(j))$ using the same α.[4] For example, every finite metric space embeds into a distribution over trees with prioritized expected distortion $O(\log j)$, which extends the $O(\log n)$ distortion known from [12]. Recently, Bartal et al. [4] showed that every finite metric space embeds into ℓ_2 with prioritized distortion $O(\log j)$, which extends the $O(\log n)$ distortion known from [6]. In fact, we are not aware of any setting where it is impossible to generalize a worst-case distortion guarantee to a prioritized guarantee. The final question we raise is the following.

Question 1.3 Does this analogy between worst-case and prioritized distortion extend also to dimension and to label-size, or perhaps their worst-case and prioritized versions exhibit a disparity?

1.1 Results: Old and New

Our main results and most relevant previous bounds are discussed below and summarized in Table 1. Additional related work is described in Sect. 1.2.

General Metrics. As discussed above, embeddings and labeling schemes for general metrics have essentially the same label size/dimension for all distortion parameters. For prioritized labelings and embeddings, the comparison is more complex. For exact labeling scheme, one can obtain label size $O(j)$ by simply storing in the label of the point x_j its distances to $x_1, \ldots, x_{j-1}$ (recall that we count words). This is essentially optimal, even if we allow distortion up to 3, see Theorem 2.1. In contrast, for embeddings into ℓ_∞, we show in Theorem 2.2 that prioritized dimension is impossible for distortion less than $3/2$. Specifically, we provide an example where the images of x_1 and x_2 must differ in at least $\Omega(n)$ coordinates for arbitrarily large n. This proves a strong separation between embeddings and labelings, and also demonstrates an embedding result that has no prioritized counterpart.

For prioritized distortion $O(\log j)$, Elkin et al. [9] constructed a labeling with prioritized label size of $O(\log j)$. We construct in Theorem 2.3 ℓ_∞-embeddings with different tradeoffs between the prioritized distortion α and dimension β. Two representative examples are prioritized distortion $\alpha(j) = O(\log j)$ with prioritized dimension $\beta(j) = O(j)$, and $\alpha(j) = O(\log\log j)$ with $\beta(j) = O(j^2)$. This is significantly better than for the $O(1)$-distortion case, yet considerably weaker than results on labeling.

Additional interesting results in this context were given in [9], showing that every metric space embeds into every ℓ_p, $p \in [1, \infty]$, with prioritized distortion $O(\log^{4+\epsilon} j)$ and prioritized dimension $O(\log^4 j)$ (for every constant $\epsilon > 0$). Furthermore, independently and concurrently to our work, Elkin and Neiman [11] obtained two additional

[4] We use $\widetilde{O}$ notation to suppress constants and logarithmic factors, that is $\widetilde{O}(\alpha(j)) = \alpha(j) \cdot \text{polylog}(\alpha(j))$.

Table 1 Summary of our findings

	Worst-case label-size/dimension					
	Distortion		Distance labeling		Embedding into ℓ_∞	
1	General metric	< 3	$\Theta(n)^*$	[14]	$\Theta(n)$	[21]
2	General metric	$O(\log n)$	$O(\log n)$	[26]	$\Theta(\log n)$	[1]
3	ℓ_p for $p \in [1, 2]$	$1+\epsilon$	$O(\epsilon^{-2}\log n)$	Theorem 3.1	$\Theta(n)^{**}$	Theorem 3.2
4	Tree	1	$O(\log n)$	[25]	$\Theta(\log n)$	[19]
5	Planar	1	$\Theta(\sqrt{n})$	[14]	$\Theta(n)$	[19]
6	Treewidth k	1	$O(k\log n)$	[14]	$\Theta(n)^\dagger$	[19]
	Prioritized label-size/dimension					
	Distortion		Distance labeling		Embedding into ℓ_∞	
7	General metric	$< 3/2$	$\Theta(j)^*$	Theorem 2.1	$\Theta(n)^\ddagger$	Theorem 2.2
8	General metric	$O(\log j)$	$O(\log j)$	[9]	$O(j)$	Corollary 2.4
9	ℓ_p for $p \in [1, 2]$	$1+\epsilon$	$O(\epsilon^{-2}\log j)$	Theorem 3.1	$j^{\Omega(1/\epsilon)}$	Theorem 3.3
10	Tree	1	$O(\log j)$	[9]	$\Theta(\log j)$	Theorem 4.2
11	Planar	1	$\Theta(j)$	Theorem 5.2	$\Theta(n)^\ddagger$	Theorem 5.1
12	Treewidth k	1	$O(k\log j)$	[9]	$\Theta(n)^{\dagger,\ddagger}$	Theorem 5.1

Question 1.1 is answered by comparing the last two columns of lines 1–6; in the very general and very restricted families (lines 1, 2, 4), labelings and embeddings perform similarly, while other families (lines 3, 5, 6) exhibit a strict separation. Question 1.2 is answered by comparing the last two columns of lines 7–12; we see a strict separation between them in all families other than trees (line 10). Question 1.3 is answered by comparing each line $i = 1, \ldots, 6$ with line $i+6$; for distance labeling, worst-case bound $\beta(n)$ translates to prioritized $O(\beta(j))$ except for planar graphs (lines 5, 11), while for embeddings, dimension translates to its prioritized version only for trees (lines 4, 10)

*The upper bound is for distortion 1 (i.e., isometric embedding)

**Holds for $1+\epsilon < \sqrt{2}$ and $p \in [1, \infty]$

†Holds for $k \geqslant 2$

‡This excludes priority dimension for any function $\alpha : \mathbb{N} \to \mathbb{N}$ that is independent of $n = |X|$

embeddings into ℓ_∞, for any integer parameter $k \geqslant 1$, there are embeddings with: (1) prioritized distortion $2\lceil k \log j / \log n \rceil - 1$ and dimension $O(kn^{1/k}\log n)$ (not prioritized); and (2) prioritized distortion $2k \log\log j + 1$ and prioritized dimension $O(kj^{2/k}\log n)$ (note that the dimension bounds here also depend on $n = |X|$ and hence are not truly prioritized). See Table 2 for a comparison of these results with ours.

ℓ_p Spaces. The seminal Johnson–Lindenstrauss Lemma [16] states that every n-point subset of ℓ_2 embeds with distortion $1+\epsilon$ into $\ell_2^{O(\epsilon^{-2}\log n)}$ (where as usual ℓ_p^d denotes the d-dimensional ℓ_p space), and this readily implies a labeling with distortion $1+\epsilon$ and label size $O(\epsilon^{-2}\log n)$. Since every ℓ_p, $p \in [1, 2]$, embeds isometrically into squared-L_2 (equivalently, its snowflake embeds into L_2), this implies a labeling with the same performance for ℓ_p as well, see Theorem 3.1. Furthermore, we show in Theorem 3.1 (using [22]) that this labeling can be prioritized to achieve distortion $1+\epsilon$ with label size $O(\epsilon^{-2}\log j)$.

Table 2 ℓ_∞-embeddings and distance labelings of general metrics with different trade-offs between prioritized distortion and dimension/label size

Distance labelings for general metrics				
Prioritized distortion		Prioritized label size	Notes	Refs.
1	$2\cdot\lceil k\log j/\log n\rceil-1$	$O(n^{1/k}\log j)$	$\forall k\in\mathbb{N}$	[9]
2	$2k-1$	$O(j^{1/k}\log j)$	$\forall k\in\mathbb{N}$	[9]
Embeddings of general metrics				
Prioritized distortion		Prioritized dimension	Notes	Refs.
3	$O(\log^{4+\epsilon} j)$	$O(\log^4 j)$	$\forall$ constant ϵ	[9]
4	$2\cdot\lceil k\log j/\log n\rceil-1$	$O(kn^{1/k}\log n)$	$\forall k\in\mathbb{N}$	[11]
	$2\cdot\lceil\log j\rceil-1$	$O(\log^2 n)$		[11]
5	$2k\log\log j+1$	$O(k(j^{2/k}+\log k)\log n)$	$\forall k\in\mathbb{N}$	[11]
6	$2\cdot\lceil k\log j/\log n\rceil$	$n^{1/k}j$	$\forall k\in\mathbb{N}$	Corollary 2.4
	$2\cdot\lceil\log j\rceil$	$2j$		Corollary 2.4
7	$2\cdot\lceil\log\log j\rceil$	j^2		Corollary 2.4

The labeling results are superior to their embedding counterparts. Line 6 is obtained by plugging in $t=\log n/k$ in Corollary 2.4. Comparing the result in line 4 to ours in line 6, in the most interesting regime of distortion $2\log j$, we achieve a truly prioritized result (with dimension independent of n), while [11] avoids linear dependencies in the dimension. Our result in line 7 is strictly superior to that of line 5, which is not truly prioritized. However, [11] provides a much wider spectrum of possible trade-offs

For ℓ_∞-embeddings, the performance is significantly worse. We show in Theorem 3.2 that for certain n-point subsets of ℓ_p, for any $p\in[1,\infty]$, embedding into ℓ_∞ with distortion less than $\sqrt{2}$ requires $\Omega(n)$ coordinates (recall that $O(n)$ coordinates are sufficient to isometrically embed every n-point metric into ℓ_∞). For prioritized embeddings into ℓ_∞ with distortion $1+\epsilon$, we prove a lower bound of $j^{\Omega(1/\epsilon)}$ on the prioritized dimension, see Theorem 3.3.

Tree Metrics. Trees are a success story, where both labelings and embeddings have the same performance. Here we study metric spaces that induced by the shortest path metric of weighted trees. In their seminal paper on metric embeddings, Linial et al. [19] proved that every n-node tree embeds isometrically into $\ell_\infty^{O(\log n)}$. In the context of routing, Thorup and Zwick [25] constructed an exact distance labeling with label size $O(\log n)$ (where routing decisions can be done in constant time), and Elkin et al. [9] modified this to achieve prioritized label size $O(\log j)$. Our contribution is a prioritized version of [19], i.e., an isometric embedding of a tree metric into ℓ_∞ with prioritized dimension $O(\log j)$, see Theorem 4.2. We note that an equivalent result was proved independently and concurrently by Elkin and Neiman [11].

Planar Graphs and Restricted Topologies. Here we study metric spaces that induced by the shortest path metric of weighted graphs with restricted topologies. We first consider exact distance labeling and isometric embeddings. Gavoille et al. [14] showed that planar graphs admit exact labeling with label size $O(\sqrt{n})$, and proved a matching lower bound.[5] They further showed that treewidth-k graphs admit exact labeling with label size $O(k \log n)$. Linial et al. [19] proved that an isometric embedding of the n-cycle graph into ℓ_∞, and in fact into any normed space, requires $\Omega(n)$ coordinates.[6] Notice that the cycle graph is both planar and has treewidth 2; hence, there is a strict separation between distance labeling and ℓ_∞-embedding.

For exact prioritized distance labeling, we prove that planar graphs require prioritized label size $\Omega(j)$ (based on [14]), see Theorem 5.2. This bound is tight, as prioritized label size $O(j)$ is possible already for general graphs (Theorem 2.1). We conclude that priorities make exact distance labelings much harder for planar graphs.[7] This lower bound for exact prioritized labeling holds for unweighted graphs as well, hence this type of labeling is now well understood. For treewidth-k graphs, Elkin et al. [9] constructed exact labeling with prioritized label size $O(k \log j)$. For isometric embeddings into ℓ_∞, we show in Theorem 5.1 that no prioritized dimension is possible for the cycle graph, which provides a lower bound for both planar and treewidth-2 graphs. This implies a dramatic separation for these families.

Additional results on labelings with $1+\epsilon$ distortion, and embeddings with constant distortion are described in Sect. 1.2.

1.2 Related Work

For distortion $1+\epsilon$ in planar graphs, Thorup [24] and Klein [17] constructed distance labels of size $O(\log n/\epsilon)$. Abraham and Gavoille [3] generalized this result to K_r-minor-free graphs, achieving label size $O(g(r) \log n/\epsilon)$.[8] No low-dimension embedding into ℓ_∞ with distortion $1+\epsilon$ is known for planar graphs or even treewidth-2 graphs. If one allows larger distortion, Krauthgamer et al. [18] proved that planar graphs embed with distortion $O(1)$ into $\ell_\infty^{O(\log n)}$, or more generally that K_r-minor-free graphs embed with distortion $O(r^2)$ into $\ell_\infty^{O(3^r \log n)}$. Abraham et al. [2] showed that K_r-minor-free graphs embed with distortion $O(1)$ into $\ell_\infty^{O(g(r) \log^2 n)}$. Turning to priorities, Elkin et al. [9] constructed prioritized versions of distance labeling with distortion $1+\epsilon$. Specifically, for planar and K_r-minor-free graphs they achieve label sizes of $O(\log j/\epsilon)$ and $O(g(r) \log j/\epsilon)$, respectively. No prioritized embeddings are known, nor lower bounds thereof.

Elkin et al. [8] studied the problem of *terminal* distortion, where there is specified a subset $K \subset X$ of terminal points, and the goal is to embed the entire space (X, d) while preserving pairwise distances among $K \times X$. For additional embeddings with terminal distortion see [4, 10]. Embeddings with terminal distortion can be used used to construct embeddings with prioritized distortion. We utilize this approach in Theorems 3.1 and 4.2.

[5] This lower bound, as well as all other lower bounds from [14], count bits rather than words.

[6] Their proof is much more general than what is required for ℓ_∞. For a simpler proof for the special case studied here, see Theorem 5.1.

[7] Interestingly, for unweighted planar graphs, Gavoille et al. [14] prove only a lower bound of $\Omega(n^{1/3})$ on the label size, and closing the gap to the upper bound $O(\sqrt{n})$ remains an important open question.

[8] The function $g(r)$ depends only on r and is taken from the structure theorem of Robertson and Seymour [23].

Abraham et al. [1] studied *scaling* distortion, which provides improved distortion for $1-\epsilon$ fractions of the pairs, simultaneously for all $\epsilon \in (0,1)$, as a function of ϵ. A stronger version called *coarse scaling* distortion has improved distortion guarantees for the farthest pairs. Bartal et al. [4] showed that scaling distortion and prioritized distortion (in the sense of [9]) are essentially equivalent, but this is not known to hold for the prioritized contractive distortion we use in the current paper (see footnote 3).

1.3 Preliminaries

The ℓ_p-norm of a vector $x=(x_1,\ldots,x_d)\in\mathbb{R}^d$ is $\|x\|_p := \big(\sum_{i=1}^d |x_i|^p\big)^{1/p}$, where $\|x\|_\infty := \max_i |x_i|$. An embedding f between two metric spaces (X,d_X) and (Y,d_Y) has distortion ct if for every $x,y\in X$,

$$\frac{d_X(x,y)}{c} \leqslant d_Y(f(x),f(y)) \leqslant t\cdot d_X(x,y).$$

t (resp. c) is the *expansion* (resp. *contraction*) of f. If the expansion is 1, we say that f is *Lipschitz*, while if $c=1$ we say that the embedding is *non-contractive*. Embedding with distortion 1 (where $c=t=1$) is called *isometric*.

Embedding $f\colon X\to\ell_\infty^d$ can be viewed as a collection of embeddings $\{f_i\}_{i=1}^d$ into the line $\mathbb{R}$. By scaling we can assume that the embedding is non-contractive. That is, if f has distortion t then for every $x,y\in X$ and i, $|f_i(x)-f_i(y)|\leqslant t\cdot d_X(x,y)$ and there is some index $i_{x,y}$ such that $d_X(x,y)\geqslant |f_{i_{x,y}}(x)-f_{i_{x,y}}(y)|$. We say that the pair x,y is *satisfied* by the coordinate $i_{x,y}$.

We consider connected undirected graphs $G=(V,E)$ with edge weights $w\colon E\to\mathbb{R}_{>0}$. Let d_G denote the shortest path metric in G. For a vertex $x\in V$ and a set $A\subseteq V$, let $d_G(x,A) := \min_{a\in A} d(x,a)$, where $d_G(x,\emptyset) := \infty$. We often abuse notation and write the graph G instead of its vertex set V.

We always measure the size of a label by the number of words needed to store it (where each word contains $O(\log n)$ bits). For ease of presentation, we will ignore issues of representation and bit counting. In particular, we will assume that every pairwise distance can be represented in a single word. We note however that the lower bounds of [14] are given in bits, and therefore our Theorem 5.2 is as well.

All logarithms are in base 2. Given a set A, $\binom{A}{2}=\{\{x,y\}\mid x,y\in A,\ x\neq y\}$ denotes all the subsets of size 2. The notation $x=(1\pm\epsilon)\cdot y$ means $(1-\epsilon)y\leqslant x\leqslant (1+\epsilon)y$.

2 General Graphs

In this section we discuss our result on succinct representations of general metric spaces. We start with the regime of small distortion. Recall that there exist both exact distance labelings with $O(n)$ label size [14] as well as isometric embeddings into ℓ_∞^n [21], and both results are essentially tight (even if one allows distortion <3). In the following theorem we provide lower and upper bounds for exact distance labelings with prioritized label size.

Theorem 2.1 *Given an n-point metric space (X, d) and priority ordering $X = \{x_1, \ldots, x_n\}$, there is an exact labeling scheme with prioritized label size j. This is asymptotically tight, that is every exact labeling scheme must have prioritized label size $\Omega(j)$. Furthermore, for $t < 3$, every labeling scheme with distortion t must have prioritized label size $\widetilde{\Omega}(j)$.*

Proof We begin by constructing the labeling scheme. The label of x_j simply consists of the index j and $d(x_1, x_j), d(x_2, x_j), \ldots, d(x_{j-1}, x_j)$. The size bound and algorithm for answering queries are straightforward. If one allows distortion $t < 3$, [14] proved that every labeling scheme with distortion t must have label size of $\Omega(n)$ bits, or $\widetilde{\Omega}(n)$ words. As some vertex must have a label of size $\widetilde{\Omega}(n)$, the prioritized lower bound $\widetilde{\Omega}(j)$ follows.

Finally, we prove the $\Omega(j)$ lower bound for exact distance labeling. We begin by arguing that some label must be of length $\Omega(n)$ (in words), and then the $\Omega(j)$ lower bound for prioritized label size follows. The proof follows the steps of [14]. Consider a complete graph with $\binom{n}{2}$ edges all having integer weights in $\{n+1, n+2, \ldots, 2n\}$. Note that there are $n^{\binom{n}{2}}$ such graphs, where each choice of weights defines a different shortest path metric. Given an exact labeling scheme, the labels $l(x_1), \ldots, l(x_n)$ precisely encode the graph. Following arguments from [14], the sum of lengths of the labels must be at least logarithmic in the number of different graphs. Thus

$$\max_i |l(x_i)| \geqslant \frac{\log n^{\binom{n}{2}}}{n} = \Omega(n \log n).$$

We conclude that some label length must be of $\Omega(n \log n)$ bits, or $\Omega(n)$ words. □

While under the standard worst-case model distance labelings and embeddings into ℓ_∞ behave identically, we show that the prioritized versions are very different. In the following theorem we show that no prioritized dimension is possible, even if one allows distortion $< 3/2$ (note that for much larger distortions, prioritized dimension is possible. See [9] and Corollary 2.4).

Theorem 2.2 *There is no function $\alpha : \mathbb{N} \to \mathbb{N}$ such that every metric space can be embedded into ℓ_∞ with prioritized dimension α and distortion $t < 3/2$ (for any fixed t).*

Proof Consider the family $\mathcal{G}$ of unweighted bipartite graphs $G = (V = L \cup R, E)$ where $|L| = |R| = n$, for large enough n. We first argue that there is a graph $G \in \mathcal{G}$ with the following properties:

(i) For every $u, v \in R$ or $u, v \in L$, we have $d_G(u, v) = 2$.
(ii) Every embedding $f : G \to \ell_\infty$ with distortion $2t$ requires $\Omega(n)$ coordinates.

The existence of G follows by a counting argument similar to [21]. Note that $|\mathcal{G}| = 2^{n^2}$. Denote by $\mathcal{G}' \subseteq \mathcal{G}$ the graphs in $\mathcal{G}$ fulfilling property (i). Our first step is to lower bound $|\mathcal{G}'|$. Sample uniformly a graph $G \in \mathcal{G}$. For $u, v \in R$ (resp. $u, v \in L$) let $I_{u,v}$ be an indicator for the event $d_G(u, v) \neq 2$. $I_{u,v}$ occurs if and only if u and v do not have a common neighbor in L (resp. R). Then $\Pr[I_{u,v}] = (3/4)^n$. By a union bound, the probability that property (i) does not hold is at most $2 \cdot \binom{n}{2} \cdot (3/4)^n$. We conclude that

$$|\mathcal{G}'| \geqslant 2^{n^2} \cdot \left(1 - 2\binom{n}{2}\left(\frac{3}{4}\right)^n\right) \geqslant \frac{2^{n^2}}{2}.$$

Matoušek [21, Proposition 3.3.1] implicitly proved that for any subset $\mathcal{G}'$ of $\mathcal{G}$, if all of $\mathcal{G}'$ embeds into ℓ_∞^d with distortion $2t < 3$, then

$$c^{dn} \geqslant |\mathcal{G}'|,$$

where $c > 1$ is a constant depending on $3 - 2t$. Thus $d = \Omega(n)$. We conclude that there is a graph $G \in \mathcal{G}$ fulfilling both properties (i) and (ii).

Consider such a graph $G = (V = L \cup R, E)$. Note that property (i) implies that there are no isolated vertices, and moreover for every $u \in R$, $v \in L$, $d_G(u, v) \in \{1, 3\}$. Let G' be the graph G along with two new vertices l, r where l (resp. r) has edges to all vertices in R (resp. L). Note that for every $u, v \in V$, $d_G(u, v) = d_{G'}(u, v)$. Set $L' = L \cup \{l\}$ and $R' = R \cup \{r\}$.

Claim Every embedding $f : G' \to \ell_\infty$ with distortion $t < 3/2$ has $\Omega(n)$ coordinates i for which $f_i(l) \neq f_i(r)$.

Proof We assume that the embedding f has expansion at most t, and for every pair of vertices there is a coordinate where the pair is satisfied (i.e., not contracted). Set $\mathcal{A}_i = \left\{\{u, v\} \in \binom{L' \cup R'}{2} \mid d_{G'}(u, v) = i\right\}$ to be all the vertex pairs at distance exactly i. Note that $\binom{L' \cup R'}{2} = \mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{A}_3$. To satisfy all the pairs in $\binom{L' \cup R'}{2}$, $\Omega(n)$ coordinates are required (this is property (ii)). We will show that we can satisfy all the pairs in $\mathcal{A}_1 \cup \mathcal{A}_2$ using $O(\log n)$ coordinates only. Thus satisfying all the pairs in $\mathcal{A}_3$ requires $\Omega(n)$ coordinates.

The clique K_n can be embedded isometrically into $\ell_\infty^{\lceil \log n \rceil}$ [19]. Such an embedding can be constructed by simply mapping K_n to different combinations of $\{0, 1\}^{\lceil \log n \rceil}$. As 1 is the minimal distance, we can just embed all the $2n + 2$ vertices as a clique using $O(\log n)$ coordinates. By doing so, all the pairs in $\mathcal{A}_1$ will be satisfied. $\mathcal{A}_2$ equals $\binom{L'}{2} \cup \binom{R'}{2}$. Note that the metric induced on $\binom{L'}{2}$ is just a clique with edges of length 2. Thus we can embed all of L' to the vectors $\{\pm 1\}^{O(\log n)}$. Additionally send all of R' to $\mathbf{0}$. Note that by doing so we satisfied all the pairs in $\binom{L'}{2}$ while incurring no expansion. Similarly we can satisfy all the pairs in $\binom{R'}{2}$ using $O(\log n)$ additional coordinates.

Next consider an arbitrary embedding $f : G' \to \ell_\infty$ with distortion $t < 3/2$. We argue that in a coordinate $f_i : G' \to \mathbb{R}$ where $f_i(l) = f_i(r)$, no pair of $\mathcal{A}_3$ is satisfied. Indeed, every vertex $v \in L' \cup R'$ is at distance 1 from either l or r. As we have expansion at most t, in a coordinate i where $f_i(l) = f_i(r)$ the maximal distance between a pair of vertices v, u is $2t$. In particular, for every $\{v, u\} \in \mathcal{A}_3$, $|f_i(x) - f_i(y)| \leqslant 2t < 3$. Thus no pair $\{v, u\} \in \mathcal{A}_3$ is satisfied. As there must be $\Omega(n)$ coordinates where some pair from $\mathcal{A}_3$ is satisfied, necessarily there are $\Omega(n)$ coordinates where $f_i(l) \neq f_i(r)$. $\square$

We conclude that there are $\Omega(n)$ coordinates where at least one of l, r is not mapped to 0. Set π to be any priority ordering wherein l and r have priorities 1 and 2 respectively. For every priority function $\alpha\colon \mathbb{N} \to \mathbb{N}$, by taking $n \gg \alpha(2), \alpha(1)$, there is no embedding with prioritized dimension α with respect to π. The theorem follows. □

Considering that for distortion less than 3/2 no prioritized dimension is possible, it is natural to ask for what distortions are prioritized embeddings possible. Some previous results of this nature are described in the introduction [9, 11]. As exact distance labeling is possible using $O(j)$ labels, it is also natural to ask what distortion can be obtained with prioritized dimension $O(j)$. The following is a meta theorem constructing various trade-offs. We present some specific implications in Corollary 2.4. A comparison between our results and other results appears in Table 2.

Consider a monotone function $\beta\colon \mathbb{N} \to \mathbb{N}$. For $j \in \mathbb{N}$, let $\chi_\beta(j)$ be the minimal i such that $\beta(\chi_\beta(j)) \geqslant j$.

Theorem 2.3 *Given a metric space (X, d) with priority ordering $X = \{x_1, \dots, x_n\}$ and a function $\beta\colon \mathbb{N} \to \mathbb{N}$, there is an embedding $f\colon X \to \ell_\infty$ with prioritized dimension $\beta(\chi_\beta(j))$ and contractive prioritized distortion $2 \cdot \chi_\beta(j)$.*

Before presenting the proof of Theorem 2.3, we provide some of the intuition behind it. Recall that the Fréchet embedding [21] (also called Kuratowski embedding) is an embedding into ℓ_∞^n, where the j'th coordinate for a point x is simply x's distance to x_j. While this is an isometric embedding, every point is non-zero in $n-1$ coordinates. In order to obtain prioritized dimension, we will set the j coordinate of a point x to be its distance to the set that contains x_j together with all points $x_{j'}$ for sufficiently small j' (where the value of j' is determined by the function β). This "padding" will ensure prioritized dimension, but also induce larger distortion as a function of β.

Proof of Theorem 2.3 We suggest that while inspecting the proof, it may be helpful for the reader to focus on the setting $\beta(i) = 2^i$, wherein $\chi_\beta(j) = \lceil \log j \rceil$. Set $S_0 = \emptyset$ and $S_i = \{x_j \mid j \leqslant \beta(i)\}$. We define embedding f by setting its j'th coordinate to be

$$f_j(x) := d\big(x, S_{\chi_\beta(j)-1} \cup \{x_j\}\big).$$

Note that for every j' such that $\chi_\beta(j') > \chi_\beta(j)$, $f_{j'}(x_j) = 0$. Note also that there may be many points $x_{j'}$ with $j' < j$ and yet $f_j(x_{j'}) \neq 0$. Thus x_j is non-zero only in the first $\beta(\chi_\beta(j))$ coordinates as required.

Next we argue the prioritized distortion. It is clear that f is Lipschitz. Consider a pair of vertices x_j, y. Set $\Delta = d(x_j, y)$, and $\alpha_i = d(\{x_j, y\}, S_i)$. Then $\infty = \alpha_0 > \alpha_1 \geqslant \alpha_2 \geqslant \dots \geqslant \alpha_{\chi_\beta(j)} = 0$. We argue that there must be some index i such that $\alpha_{i+1} \leqslant \min\{\alpha_i, \Delta/2\} - \Delta/(2\chi_\beta(j))$. Suppose for contradiction otherwise (i.e., no such index exist). We argue by induction on $q \in [1, \chi_\beta(j)]$ that $\alpha_{\chi_\beta(j)-q} < q\Delta/(2\chi_\beta(j))$. For the base case note that $0 = \alpha_{\chi_\beta(j)} > \min\{\alpha_{\chi_\beta(j)-1}, \Delta/2\} - \Delta/(2\chi_\beta(j))$, implying $\alpha_{\chi_\beta(j)-1} < \Delta/(2\chi_\beta(j))$. For general q, using the induction hypothesis $q\Delta/(2\chi_\beta(j)) > \alpha_{\chi_\beta(j)-q} > \min\{\alpha_{\chi_\beta(j)-q-1}, \Delta/2\} - \Delta/(2\chi_\beta(j))$, implying $\min\{\alpha_{\chi_\beta(j)-q-1}, \Delta/2\} < (q+1)\Delta/(2\chi_\beta(j))$ and therefore $\alpha_{\chi_\beta(j)-q-1} < (q+1)\Delta/(2\chi_\beta(j))$. Overall we conclude that $\alpha_0 = \alpha_{\chi_\beta(j)-\chi_\beta(j)} < \chi_\beta(j) \cdot \Delta/(2\chi_\beta(j)) = \Delta/2$, a contradiction as $\alpha_0 = \infty$.

Choose $z \in S_{i+1}$ such that $d(\{x_j, y\}, z) = \alpha_{i+1}$, and suppose that $z = x_q$. Assume without loss of generality that $d(x_j, z) = d(\{x_j, y\}, z) = \alpha_{i+1}$, and so $d(y, z) \geqslant d(x_j, y) - d(x_j, z) \geqslant \Delta - \Delta/2 + \Delta/(2\chi_\beta(j)) > \Delta/2$. It holds that $d(y, S_i \cup \{z\}) = \min\{d(y, S_i), d(y, z)\} \geqslant \min\{\alpha_i, \Delta/2\}$. Thus

$$\|f(y) - f(x_j)\|_\infty \geqslant |f_q(y)_\infty - f_q(x_j)| = \left|d(y, S_i \cup \{z\}) - d(x_j, S_i \cup \{z\})\right| \geqslant \left|\min\left\{\alpha_i, \frac{\Delta}{2}\right\} - \alpha_{i+1}\right| \geqslant \frac{\Delta}{2\chi_\beta(j)}.$$

Prioritized distortion $2 \cdot \chi_\beta(j)$ follows. □

Corollary 2.4 *Given a metric space (X, d) with priority ordering $X = \{x_1, \ldots, x_n\}$,*

- *For every $t \in \mathbb{N}$, there is an embedding $f\colon X \to \ell_\infty$ with prioritized distortion $2 \cdot \lceil \log j/t \rceil$ and prioritized dimension $2^t \cdot j$.*
- *There is an embedding $f\colon X \to \ell_\infty$ with prioritized distortion $2 \cdot \lceil \log\log j \rceil$ and prioritized dimension j^2.*

Proof The first case follow by choosing the function $\beta(i) = 2^{ti}$. Here $\chi_\beta(j) = \lceil \log_{2^t} j \rceil = \lceil \log j/t \rceil$, and thus the prioritized distortion is $2 \cdot \lceil \log j/t \rceil$ while the prioritized dimension is $\beta(\chi_\beta(j)) = 2^{t \cdot \lceil \log j/t \rceil} < 2^{t + \log j} = 2^t j$. For the second case choose $\beta(i) = 2^{2^i}$. Here $\chi_\beta(j) = \lceil \log\log j \rceil$, and thus the prioritized distortion is $2 \cdot \lceil \log\log j \rceil$ and the prioritized dimension $\beta(\chi_\beta(j)) = 2^{2^{\lceil \log\log j \rceil}} < 2^{2 \cdot 2^{\log\log j}} = j^2$. □

Note that the first case implies prioritized distortion $2 \cdot \lceil \log j \rceil$ and prioritized dimension $2j$.

3 ℓ_p Spaces

In this section we consider representations of ℓ_p spaces. As these spaces are somewhat restricted, we focus on the $1 + \epsilon$ distortion regime. We begin with the upper bound for distance labeling.

Theorem 3.1 *For every $\epsilon > 0$, $p \in [1, 2]$, and n points in ℓ_p, there is a $(1+\epsilon)$-labeling scheme with label size $O(\epsilon^{-2} \log n)$. Furthermore, given a priority ordering π, there is a $(1 + \epsilon)$-labeling scheme with prioritized label size $O(\epsilon^{-2} \log j)$.*

Proof We begin by constructing a labeling scheme for a set X on n points in ℓ_2. Then we will generalize the result to ℓ_p for $p \in [1, 2]$.

As a consequence of the Johnson–Lindenstrauss Lemma [16], there is an embedding $f\colon X \to \ell_2^{O(\epsilon^{-2} \log n)}$ with $1 + \epsilon$ distortion. By simply storing $f(x)$ as the label of $x \in X$, we obtain a $1 + \epsilon$ labeling scheme with $O(\epsilon^{-2} \log n)$ label size. Next, we consider a set X with priority ordering $\pi = \{x_1, x_2, \ldots, x_n\}$. Narayanan and Nelson [22] constructed a terminal version of the JL transform: Specifically, given a set K of k points in ℓ_2 there is an embedding f of the entire ℓ_2 space into $\ell_2^{O(\epsilon^{-2}\log k)}$ such that for every $x \in K$ and $y \in \ell_2$, $\|f(x)-f(y)\|_2 = (1\pm\epsilon)\|x-y\|_2$. For $i = 0,1,\dots,\lceil\log\log n\rceil$, set $S_i = \{x_j \mid j \leqslant 2^{2^i}\}$. Let $f_i\colon X \to \ell_2^{O(\log|S_i|)}$ be a terminal JL transform w.r.t. S_i. The label of x_j will consist of $f_0(x_j), f_1(x_j), \dots, f_{\lceil\log\log j\rceil}(x_j)$. Given a query on $x_j, x_{j'}$, where $j < j'$, our answer will be $\|f_{\lceil\log\log j\rceil}(x_j) - f_{\lceil\log\log j\rceil}(x_{j'})\|_2$. The distortion follows as $x_j \in S_{\lceil\log\log j\rceil}$ (hence [22] guarantees $\|f_{\lceil\log\log j\rceil}(x_j) - f_{\lceil\log\log j\rceil}(x_{j'})\|_2 = (1\pm\epsilon)\|x_j - x_{j'}\|_2$). The length of the label of x_j is bounded by

$$\sum_{i=0}^{\lceil\log\log j\rceil} O(\epsilon^{-2}\log|S_i|) = O(\epsilon^{-2})\sum_{i=0}^{\lceil\log\log j\rceil} 2^i$$
$$= O(\epsilon^{-2})\cdot 2^{\lceil\log\log j\rceil+1} = O(\epsilon^{-2}\log j),$$

words, as required.

To generalize the labeling schemes to ℓ_p for $p \in [1,2]$, we note that every $p \in [1,2]$, ℓ_p embeds isometrically into squared-L_2, or equivalently, the snowflake of ℓ_p embeds into L_2 (see e.g. [7]). Specifically, for a set $X \subseteq \ell_p$, there is a function $f_X\colon X \to \ell_2$, such that for every $x, y \in X$, $\|x-y\|_p = \|f(x)-f(y)\|_2^2$. Then a labeling scheme for ℓ_2 implies the same performance for ℓ_p as well, the only change being that the computed distances must be squared. □

Next we turn our attention to lower bounds. Every n-point set in ℓ_2 embeds isometrically into any other ℓ_p space, for $p \in [1,\infty]$ (see e.g. [21]). This implies that any lower bound that we prove for ℓ_2 will holds as well for any other ℓ_p space (as the hard example will reside in ℓ_p as well).

Theorem 3.2 *For every $p \in [1,\infty]$ and $n \in \mathbb{N}$, there is a set A of $2n$ points in ℓ_p, such that every embedding of A into ℓ_∞ with distortion smaller than $2^{\max\{1/2,1-1/p\}}$ has dimension at least n.*

Proof Set $A = \{e_1, -e_1, e_2, -e_2, \dots, e_n, -e_n\}$, the standard basis and its *antipodal* points (here $\{e_i, -e_i\}$ is an *antipodal pair*). Fix p, and we will prove that every embedding of $A \subseteq \ell_p$ with distortion smaller than $2^{1-1/p}$ into ℓ_∞ requires at least n coordinates. As mentioned above, the lower bound for $p = 2$ holds for all ℓ_p as well; thus the theorem will follow.

We argue that each coordinate can satisfy at most a single antipodal pair. As there are n such pairs, the lower bound follows. Consider a single coordinate $f\colon A \to \mathbb{R}$. Assume by way of contradiction that there are $e_i, -e_i, e_j, -e_j \in A$, $i \neq j$, such that $2 \leqslant |f(e_i) - f(-e_i)|, |f(e_j) - f(-e_j)|$. As $f(e_i), f(-e_i), f(e_j), f(-e_j) \in \mathbb{R}$, by case analysis there must be a pair consisting of one point from $\{f(e_i), f(-e_i)\}$, and one point from $\{f(e_j), f(-e_j)\}$ at distance at least $\min\{|f(e_i) - f(-e_i)|, |f(e_j) - f(-e_j)|\} \geqslant 2$. But the actual distance between this pair is only $2^{1/p}$. Thus f has distortion $2/2^{1/p} = 2^{1-1/p}$, a contradiction. □

Note that Theorem 3.2 implies a lower bound of $\Omega(j)$ on the prioritized dimension of an embedding from ℓ_p into ℓ_∞, with distortion smaller than $\sqrt{2}$. Next, for distortion $1+\epsilon$ we prove a stronger lower bound with exponential dependency on ϵ.

Theorem 3.3 *For every $\epsilon \in (0, 1)$ and $p \in [1, \infty]$ there is a set of points in ℓ_p and a priority ordering, such that every embedding of them into ℓ_∞ with distortion $1 + \epsilon$ has prioritized dimension at least $j^{1/6\epsilon}$.*

Proof As above, we may assume that $p = 2$. Furthermore, we will assume that $\epsilon < 1/6$, as otherwise a better lower bound follows from Theorem 3.2. Let n be large enough, and $H_n = \{\pm 1\}^n \subseteq \ell_2^n$ be the Hamming cube. We start by creating a symmetric subset $A \subset H_n$ (i.e., $A = -A$), such that all the points in A differ in more than $\epsilon' n$ coordinates, for $\epsilon' = 3\epsilon$. The set A is created in a greedy manner. Initially set $S = H_n$ and $A = \emptyset$. First pick an arbitrary pair $x, -x \in S$ from S and add them to A. Delete from S all the points that differ in fewer than $\epsilon' n$ coordinates from either x or $-x$. Note that when $y \in S$ is deleted, so is its *antipodal* point $-y$. Thus, both S, A are maintained to be symmetric. We continue with this process until S is empty. The number of points that differ by at most $\epsilon' n$ coordinates from any point $v \in H$ is

$$\sum_{i=0}^{\epsilon' n} \binom{n}{i} \leqslant \binom{n}{\epsilon' n}\left(1 + \frac{\epsilon' n}{n - 2\epsilon' n + 1}\right) < 2\binom{n}{\epsilon' n}.$$

Therefore for each added vertex we deleted fewer than $2\binom{n}{\epsilon' n}$ points. We conclude that the size of A is at least

$$|A| \geqslant \frac{2^n}{2\binom{n}{\epsilon' n}} \geqslant \frac{1}{2} \cdot \frac{2^n}{\left(\frac{en}{\epsilon' n}\right)^{\epsilon' n}} = \frac{2^{(1-\epsilon' \log(e/\epsilon'))n}}{2} > 2 \cdot 2^{n/2}. \tag{3.1}$$

We argue that an embedding f of A into $\mathbb{R}$ can satisfy at most a single antipodal pair $x, -x$. Indeed, assume by way of contradiction that there is $f : A \to \mathbb{R}$ and $x, y \in A$ such that $\sqrt{4n} \leqslant |f(x) - f(-x)|, |f(y) - f(-y)| \leqslant (1 + \epsilon)\sqrt{4n}$. Similarly to the proof of Theorem 3.2, by case analysis, there must be a pair $z \in \{x, -x\}$ and $w \in \{y, -y\}$ such that $|f(z) - f(w)| \geqslant \min\{|f(x) - f(-x)|, |f(y) - f(-y)|\} \geqslant \sqrt{4n}$. As both $x, -x$ differ from both $y, -y$ in more than $\epsilon' n$ coordinates, z coincides with w in at least $\epsilon' n$ coordinates. In particular $\|z - w\|_2 \leqslant \sqrt{4n(1 - \epsilon')}$. Thus f has distortion at least

$$\frac{|f(z) - f(w)|}{\|z - w\|_2} \geqslant \frac{\sqrt{4n}}{\sqrt{4n(1 - \epsilon')}} > 1 + \epsilon,$$

a contradiction.

Next, let $Y = \{\pm 1\}^{\epsilon' n}\{0\}^{(1-\epsilon')n}$ be the set of all points that attain values $\{\pm 1\}$ in the first $\epsilon' n$ coordinates, with all other coordinates 0. Consider a coordinate $f : X \to \mathbb{R}$ that sends all of Y to $\mathbf{0}$. We argue that f will not satisfy any antipodal pair in A. Indeed, consider an antipodal pair $x, -x$. Let $y \in Y$ be the point agreeing with x on the first $\epsilon' n$ coordinates and 0 everywhere else. It holds that

$$\begin{aligned}|f(x)-f(-x)| &\leqslant |f(x)-f(y)|+|f(y)-f(-y)|+|f(-y)-f(-x)|\\ &\leqslant (1+\epsilon)(\|x-y\|_2+0+\|(-x)-(-y)\|_2)\\ &= (1+\epsilon)\cdot 2\cdot\sqrt{(1-\epsilon')n} < \sqrt{4n}.\end{aligned}$$

As each coordinate can satisfy at most a single antipodal pair from A, we conclude that every $1+\epsilon$ embedding of X into ℓ_∞ must be non-zero on Y in at least $|A|/2$ coordinates.

We can now conclude the proof: Assume by way of contradiction that for any set in ℓ_2 there is a $1+\epsilon$ embedding into ℓ_∞ with prioritized dimension $j^{1/6\epsilon}$. Set priority π for $X = A\cup Y$ with the points in Y occupying the first $|Y|$ places. By our assumption, there is a $1+\epsilon$ embedding where the points of Y are non-zero only in the first

$$|Y|^{1/6\epsilon} = (2^{\epsilon' n})^{1/2\epsilon'} = 2^{n/2} \overset{(3.1)}{<} \frac{|A|}{2}$$

coordinates. Thus the embedding cannot satisfy all the pairs in A, a contradiction. □

4 Trees

In this section, we present an embedding of trees into ℓ_∞ with prioritized dimension $O(\log j)$. We begin by sketching the classic isometric embedding of trees into $\ell_\infty^{O(\log n)}$ due to [19]. This embedding utilizes a balanced decomposition (a technique also used in distance labelings for trees): First, identify a separator vertex s among the vertex set V, such that we can decompose T into two trees T_1, T_2, each containing at most $2n/3+1$ vertices, where $T_1\cap T_2 = \{s\}$. Now create a new coordinate wherein each vertex $v\in T_1$ assumes value $d(v,s)$, while each vertex $x\in T_2$ assumes value $-d(x,s)$. This coordinate satisfies all pairwise distances $T_1\times T_2$. Recursively (and separately) embed T_1 and T_2 into ℓ_∞, recalling that each has its own copy of s. The two embeddings are then merged by translating T_2 so that its copy of s is mapped to the same vector assumed by the copy of s in T_1.

Given a priority ordering on the vertices $v_1, v_2, \ldots, v_n$, our goal is to create an isometric embedding into ℓ_∞ with prioritized dimension $O(\log j)$. A natural first step would be to devise a terminal embedding: Given terminal set K, embed T into $\ell_\infty^{O(\log |K|)}$ while preserving all pairwise distances $K\times V$. A terminal embedding can be constructed following the lines of the classic embedding by modifying the separator decision rule, and ensuring that after $O(\log |K|)$ recursive steps each terminal is found in a different subtree. However, a terminal embedding of this type is too weak to yield a prioritized embedding, since the mapping of all terminals into $\mathbf{0}$ (subsequent to their first $O(\log k)$ non-zero coordinates) interferes with the distances between non-terminal pairs.

To circumvent this problem, we shall "fold" the terminals one above the other, until ultimately all terminals will fall on a single representative vertex (see Lemma 4.1). During such a folding, some of the non-terminal vertices will fold upon each other as well, but our terminal embedding will be sufficiently robust to ensure that their distances are retained. We will then use this result on terminal embeddings of trees into ℓ_∞ (Lemma 4.1) to derive the stronger result, priority embeddings of trees into ℓ_∞ (Theorem 4.2).

4.1 Terminal Lemma

Lemma 4.1 *Given a weighted tree $T = (V, E, w)$ and a set K of k terminals, there exist a pair of embeddings $f : T \to \ell_\infty^{O(\log k)}$ and $g : T \to \mathcal{T}$ (into another weighted tree $\mathcal{T}$) such that the following properties hold:*

(i) Lipschitz: *For every $x, y \in V$, $\|f(x) - f(y)\|_\infty \leqslant d_T(x, y)$ and $d_{\mathcal{T}}(g(x), g(y)) \leqslant d_T(x, y)$.*
(ii) Preservation: *For every $x, y \in V$, either $\|f(x) - f(y)\|_\infty = d_T(x, y)$ or $d_{\mathcal{T}}(g(x), g(y)) = d_T(x, y)$, or both.*
(iii) Terminal Collapse: *g maps all of K into a single vertex, i.e., $|g(K)| = 1$.*

Proof We may assume that all terminals of K are leafs, as otherwise we can simply add a dummy vertex in place of each terminal, and connect the terminal to the dummy vertex with an edge of weight 0. The proof is by induction on k.

Base cases. For the case $k = 1$ we can just return the tree as is, along with the null embedding into ℓ_∞. Next we prove the case of $k = 2$. Denote the two terminals by t_1, t_2, and let P be the unique path in T connecting t_1, t_2. Let $c \in V$ be the midpoint of t_1 and t_2, such that $d_T(t_1, c) = d_T(t_2, c)$. (If c does not exist in V, then add c to V, and split the corresponding middle edge into two new edges joined at c.) Now "fold" P around c. That is, create a new tree $\mathcal{T}$, where path P is replaced by a new path that ends at c, and every $x \in P$ is found on the new path at distance exactly $d_T(x, c)$ from c. Any pair of points in P equidistant from c are merged—and in particular t_1 and t_2 are now the same point, which is the other endpoint of the new path. All the other edges and vertices remain the same. As a result, we obtain an embedding $g : d_T \to \mathcal{T}$ (see Fig. 1 for an illustration). It is clear that g is Lipschitz, and moreover $|g(\{t_1, t_2\})| = 1$.

Having specified the function g, we now describe the function f: separate T into two trees T_1, T_2 where $T_1 \cap T_2 = \{c\}$. Set the function $f : V \to \mathbb{R}$ as follows:

$$f(v) = \begin{cases} d_T(v, c) & v \in T_1, \\ -d_T(v, c) & v \in T_2 \setminus \{c\}. \end{cases} \tag{4.1}$$

See Fig. 1 for an illustration of function f. We argue that f is Lipschitz: Consider a pair of vertices u, v. If $u, v \in T_i$ (for some i), then by the triangle inequality $|f(u) - f(v)| = |d_T(u, c) - d_T(v, c)| \leqslant d_T(u, v)$. Otherwise, assume without loss of generality that $u \in T_1$ while $v \in T_2$. The shortest path from u to v must pass through c, thus

$$|f(u) - f(v)| = |d_T(u, c) + d_T(v, c)| = d_T(u, v). \tag{4.2}$$

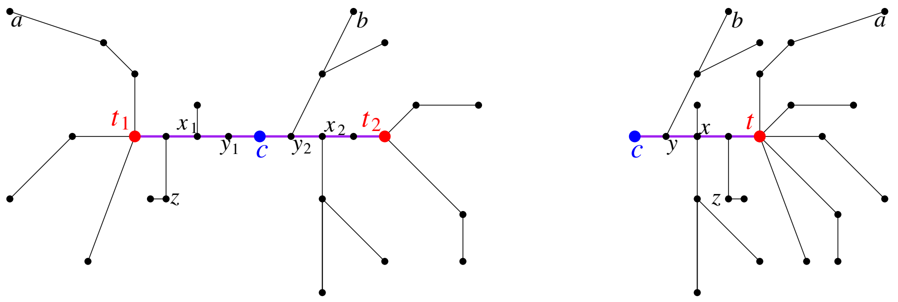

Fig. 1 On the left is illustrated the tree T with two terminals t_1, t_2. The path P between the terminals is colored in purple. The (possibly imaginary) vertex c lies at the midpoint of t_1 and t_2. On the right is illustrated the tree $\mathcal{T}$ which is obtained by "folding" the path P around c. In this example, all the edges in T are of unit weight, except for the edge $\{y_1, y_2\}$ that has weight 2. The values of the function $f\colon T \to \mathbb{R}$ (see (4.1)) are: $f(t_1) = 4$, $f(t_2) = -4$, $f(a) = 7$, $f(b) = -3$, $f(x_1) = 2$, $f(x_2) = -2$, $f(z) = 4$

It remains only to prove the second property (preservation). Consider a pair of vertices u, v. If $u \in T_1$ and $v \in T_2$, then by (4.2), $|f(u) - f(v)| = d_T(u, v)$. Otherwise, if $u, v \in T_i$, the shortest path between u and v in T is isomorphic to the shortest path in $\mathcal{T}$, and so $d_{\mathcal{T}}(u, v) = d_T(u, v)$ as required.

Induction step. For $k > 2$ terminals, we will assume by induction that for every tree with $k' < k$ terminals there are embeddings f, g as required above, such that f uses at most $a \log k'$ coordinates, for $a = 2/\log(3/2)$. Consider a tree T, and a terminal set K of size k. Let $s \in V$ be a separator vertex, such that T can be separated into two trees T_1, T_2 where $T_1 \cap T_2 = \{s\}$, and each T_i contains at most $2k/3$ terminals. As all the terminals are leafs, $s \notin K$. Create a single new coordinate $h^s\colon V \to \mathbb{R}$ defined as follows:

$$h^s(x) = \begin{cases} d_T(x, s) & x \in T_1, \\ -d_T(x, s) & x \in T_2. \end{cases}$$

It is clear that h^s is Lipschitz, and that for every $x \in T_1$, $y \in T_2$, $|h^s(x) - h^s(y)| = d_T(x, y)$. For $i \in \{1, 2\}$, invoke the induction hypothesis on T_i with terminal set $K_i = T_i \cap K$, creating embedding pair $f_i\colon T \to \ell_\infty^{a \log |K_i|}$ and $g_i\colon T \to \mathcal{T}_i$ which together satisfy requirements (i)–(iii). By padding with 0-valued coordinates, we can assume that both f_1 and f_2 use exactly $a \cdot \log(2k/3)$ coordinates. Moreover, by translation we can assume that $f_1(s) = f_2(s) = \mathbf{0}$ (note that there are no prioritized/terminal dimension guarantees here). Set f_{12} to be the combined function of f_1, f_2:

$$f_{12}(x) = \begin{cases} f_1(x) & x \in T_1, \\ f_2(x) & x \in T_2. \end{cases}$$

We argue that the function f_{12} is Lipschitz: For $x, y \in T_i$, $\|f_{12}(x) - f_{12}(y)\|_\infty = \|f_i(x) - f_i(y)\|_\infty \leqslant d_{T_i}(x, y) = d_T(x, y)$. On the other hand for $x \in T_1$ any $y \in T_2$, using the triangle inequality

$$
\begin{aligned}
\|f_{12}(x)-f_{12}(y)\|_\infty &\leqslant \|f_{12}(x)-f_{12}(s)\|_\infty+\|f_{12}(s)-f_{12}(y)\|_\infty\\
&\leqslant d_{T_1}(x,s)+d_{T_2}(s,y)=d_T(x,s)+d_T(s,y)=d_T(x,y).
\end{aligned}
$$

Set f_{12s} to be the concatenation of f_{12} with h^s, and it is clear that f_{12s} is Lipschitz as well. This completes the description of the embedding into ℓ_∞.

For the embedding into the tree, let $\mathcal{T}_{12}$ be composed of the trees $\mathcal{T}_1$ and $\mathcal{T}_2$ glued together in $g_1(s)$, $g_2(s)$. Similarly define $g_{12}\colon T\to\mathcal{T}_{12}$ as follows:

$$
g_{12}(x)=\begin{cases} g_1(x) & x\in T_1,\\ g_2(x) & x\in T_2.\end{cases}
$$

Using the triangle inequality in the same manner as for f_{12}, it is clear that g_{12} is Lipschitz. We argue that requirement (ii) holds w.r.t. f_{12s}, g_{12}. Indeed, for u, v in T_i,

$$
\begin{aligned}
&\max\bigl\{\|f_{12s}(x)-f_{12s}(y)\|_\infty, d_{\mathcal{T}_{12}}(g_{12}(x),g_{12}(y))\bigr\}\\
&\qquad\geqslant \max\bigl\{\|f_i(x)-f_i(y)\|_\infty, d_{\mathcal{T}_i}(g_i(x),g_i(y))\bigr\}=d_{T_i}(x,y)=d_T(x,y).
\end{aligned}
$$

On the other hand, for $v\in T_1$, $u\in T_2$,

$$
\max\bigl\{\|f_{12s}(v)-f_{12s}(u)\|_\infty, d_{\mathcal{T}_{12}}(g_{12}(v),g_{12}(u))\bigr\}\geqslant |h^s(v)-h^s(u)|=d_T(v,u).
$$

However, requirement (iii) does not immediately hold, as $\mathcal{T}_{12}$ contains two terminals $g_1(K_1)$, $g_2(K_2)$. Invoke the lemma for the case of $k=2$ to create two embeddings $\hat f\colon\mathcal{T}_{12}\to\mathbb{R}$, $\hat g\colon\mathcal{T}_{12}\to\mathcal{T}$ that fulfill requirements (i)–(iii). Set $f=f_{12s}\oplus\hat f(g_{12})$ to be the concatenation of $f_{12\,s}$ with $\hat f(g_{12})$ and $g=\hat g(g_{12})$ to be the composition of $\hat g$ with g_{12} ending in the tree $\mathcal{T}$. It is clear that both f, g are Lipschitz, as the Lipschitz property is preserved under concatenation and composition, thus establishing (i). Moreover, g maps all terminals to a single vertex. Requirement (ii) also holds:

$$
\begin{aligned}
d_T(u,v)&=\max\bigl\{\|f_{12s}(v)-f_{12s}(u)\|_\infty, d_{\mathcal{T}_{12}}(g_{12}(v),g_{12}(u))\bigr\}\\
&=\max\bigl\{\|f_{12s}(v)-f_{12s}(u)\|_\infty,\\
&\qquad\qquad |\hat f(g_{12}(v))-\hat f(g_{12}(u))|, d_{\mathcal{T}}(\hat g(g_{12}(v)),\hat g(g_{12}(v)))\bigr\}\\
&=\max\bigl\{\|f(v)-f(u)\|_\infty, d_{\mathcal{T}}(g(v),g(v))\bigr\}.
\end{aligned}
$$

Finally, and recalling that $a=2/\log(3/2)$, the number of coordinates is bounded by

$$
a\log\frac{2k}{3}+1+1=a\log k+a\log\frac{2}{3}+2=a\log k.
$$

The lemma now follows. □

4.2 Prioritized Embedding of Trees into ℓ_∞

Theorem 4.2 *Given a weighted tree $T=(V,K,w)$ and a priority ordering π over V, there is an isometric embedding f into ℓ_∞ with prioritized dimension $O(\log j)$.*

Proof Let $\pi=\{x_1,x_2,\dots,x_n\}$ be a priority order. Set $S_i=\{x_i : i\leqslant 2^{2^i}\}$ for $1\leqslant i\leqslant\lceil\log\log n\rceil$. Using Lemma 4.1, w.r.t. terminal set S_1 construct embeddings $f_1\colon T\to\ell_\infty^{O(\log|S_1|)}$ and $g_1\colon T\to T_1$. It holds that $g_1(S_1)$ is a single vertex in T_1, and for every $u,v\in V$, $d_T(u,v)=\max\{\|f_1(u)-f_1(v)\|_\infty, d_{T_1}(g_1(u),g_1(v))\}$. Next, using Lemma 4.1 again, w.r.t. terminal set $g_1(S_2)$, construct embeddings $f_2\colon g_1(T)\to\ell_\infty^{O(\log|S_2|)}$ and $g_2\colon g_1(T)\to T_2$. By translation, we can assume that $f_2(g_1(S_1))=\mathbf{0}$. Furthermore, $g_2(g_1(S_2))$ is a single vertex in T_2. It also holds that

$$d_T(u,v)=\max\bigl\{\|f_1(u)-f_1(v)\|_\infty, \|f_2(g_1(u))-f_2(g_1(v))\|_\infty, d_{T_2}(g_2(g_1(u)),g_2(g_1(v)))\bigr\}.$$

More generally, in the i-th step, we invoke Lemma 4.1 on T_{i-1} (w.r.t. terminal set $g_{i-1}(g_{i-2}(\cdots(g_1(S_i)))))$ to construct tree T_i and embeddings f_i, g_i. By induction, we constructed trees $T_1,\dots,T_i$ and embeddings $f_1\colon T\to\ell_\infty^{O(\log|S_1|)},\dots, f_i\colon T_{i-1}\to\ell_\infty^{O(\log|S_i|)}$, $g_1\colon T\to T_1,\dots,g_i\colon T_{i-1}\to T_i$ such that for all $q\in[1,i]$, $g_q(g_{q-1}(\dots(g_1(S_q))))$ is a single vertex in T_q and $f_q(g_{q-1}(\cdots(g_1(S_{q-1}))))=\{\mathbf{0}\}$. Furthermore,

$$\begin{aligned}d_T(u,v)=\max\bigl\{&\|f_1(u)-f_1(v)\|_\infty,\dots,\\ &\|f_i(g_{i-1}(\cdots(g_1(u))))-f_i(g_{i-1}(\cdots(g_1(u))))\|_\infty,\\ &d_{T_i}\bigl(g_i(g_{i-1}(\cdots(g_1(u)))),g_i(g_{i-1}(\cdots(g_1(u))))\bigr)\bigr\}.\end{aligned}\tag{4.3}$$

Denote $\alpha=\lceil\log\log n\rceil$. After α steps we get functions and trees as above. Set

$$f=f_1\oplus(f_2\circ g_1)\oplus(f_3\circ g_2\circ g_1)\oplus\dots\oplus(f_\alpha\circ g_{\alpha-1}\circ\dots\circ g_1)\colon T\to\ell_\infty.$$

We argue that f is an isomorphic embedding with prioritized dimension $O(\log j)$ as promised. Note that all vertices of V belong to S_α and hence mapped by $g_\alpha(g_{\alpha-1}(\cdots(g_1)))$ to the same vertex. Thus for every $u,v\in V$,

$$d_{T_\alpha}\bigl(g_\alpha(g_{\alpha-1}(\cdots(g_1(u)))),g_\alpha(g_{\alpha-1}(\cdots(g_1(v))))\bigr)=0.$$

By (4.3) we get

$$\begin{aligned}d_T(u,v)&=\max\bigl\{\|f_1(u)-f_1(v)\|_\infty,\dots,\\ &\qquad\|f_\alpha(g_{\alpha-1}(\cdots(g_1(u))))-f_\alpha(g_{\alpha-1}(\cdots(g_1(u))))\|_\infty\bigr\}\\ &=\|f(u)-f(v)\|_\infty.\end{aligned}$$

Finally we argue that f has prioritized dimension $O(\log j)$. Consider $x_j\in S_{\lceil\log\log j\rceil}$. For every $i>\lceil\log\log j\rceil$ it holds that $f_i(g_{i-1}(g_{i-2}(\cdots(g_1(x_j)))))=\mathbf{0}$ (as $x_j\in$

S_{i-1}). Therefore x_j might be non-zero only in the first

$$\sum_{i=1}^{\lceil \log\log j \rceil} O(\log |S_i|) = O\left(\sum_{i=1}^{\lceil \log\log j \rceil} 2^i\right) = O\left(2^{\lceil \log\log j \rceil + 1}\right) = O(\log j)$$

coordinates. □

5 Planar Graphs

The theorem below demonstrates that any isometric embedding of the cycle graph C_{2n} into ℓ_∞ requires dimension n. Furthermore, no prioritized dimension is possible for isometric embeddings of the cycle graph. The cycle graph is an interesting example as it is both planar and has treewidth 2. The non-prioritized lower bound is a special case of a theorem proved in [19], which applies to general norms. Nonetheless, the proof provided here is much simpler.

Theorem 5.1 *For every $n \in N$, every isometric embedding of C_{2n} (the unweighted cycle graph) into ℓ_∞ requires at least n coordinates. Furthermore, there is no function $\alpha : \mathbb{N} \to \mathbb{N}$ for which the family of cycle graphs $\{C_n\}_{n\in\mathbb{N}}$ can be embedded into ℓ_∞ with prioritized dimension α.*

Proof Denote the vertices of C_{2n} by $V = \{v_0, v_1, \ldots, v_{2n-1}\}$. The maximum distance is n, and it is realized on all the antipodal pairs $\{v_0, v_n\}, \{v_1, v_{n+1}\}, \ldots, \{v_{n-1}, v_{2n-1}\}$. We argue that in a single embedding into the line $\mathbb{R}$, at most one antipodal pair can be satisfied, that is realize distance n. Indeed, suppose by way of contradiction that there is a non-expansive function $f : C_{2n} \to \mathbb{R}$ such that $|f(v_j) - f(v_{n+j})| = |f(v_i) - f(v_{n+i})| = n$ for $i \neq j$, then necessarily $\max\{|f(v_i) - f(v_j)|, |f(v_i) - f(v_{n+j})|, |f(v_{n+i}) - f(v_j)|, |f(v_{n+i}) - f(v_{n+j})|\} \geqslant n$, a contradiction. As there are n antipodal pairs, every isometric embedding requires at least n coordinates.

For the second part, for sufficiently large n set a priority ordering π of C_{2n} where v_n, v_{n+1} have priorities 1 and 2 respectively. Consider a single Lipschitz coordinate $f : C_{2n} \to \mathbb{R}$ sending both v_n, v_{n+1} to 0. By the triangle inequality, for every antipodal pair $\{v_i, v_{n+i}\}$, it holds that

$$\begin{aligned} |f(v_i) - f(v_{n+i})| &\leqslant |f(v_i) - f(v_n)| + |f(v_n) - f(v_{n+1})| + |f(v_{n+1}) - f(v_{n+i})| \\ &\leqslant (n-i) + 0 + (i-1) = n - 1 < n. \end{aligned}$$

Thus no antipodal pair could be satisfied. We conclude that $f(v_i) \neq f(v_{n+i})$ in at least n coordinates. In particular for $\alpha(2) < n$, priority distortion α is impossible. □

Theorem 5.2 *Every isometric prioritized labeling scheme for planar graphs must have prioritized label size of at least $\Omega(j)$ (in bits). This lower bound holds even for unweighted planar graphs.*

Proof Recall that [14] proved an $\Omega(n^{1/3})$ lower bound on the label size for exact distance labeling for unweighted planar graphs. We will use the same example graph G from [14]. We refer to [14] for the description of G; here it suffices to describe its relevant properties. Given a parameter n, $G = (V, E)$ is an unweighted planar graph with $O(n^3)$ vertices, among which $O(n)$ lie on the outer face, denoted $\widetilde{V} \subset V$. Set $E = E_1 \cup E_2$, where $|E_2| = \Omega(n^2)$. For every subset $A \subset E_2$, denote by $G_A = (V, E_1 \cup A)$ the graph G wherein the edge-set $E_2 \setminus A$ has been removed (equivalently, where only the edge-set $E_1 \cup A$ is retained). [14] showed that given all pairwise distances between the outerface vertices $\{d_{G_A}(v, u) \mid v, u \in \widetilde{V}\}$, one can recover the set A. Note that $\log 2^{|E_2|} = \Omega(n^2)$ bits are required to encode the set A.

Suppose by way of contradiction that there is an exact prioritized labeling scheme with $o(j)$ labels size (in bits). Given a graph G_A, we define a priority ordering where the vertices of $\widetilde{V}$ occupy the first $|\widetilde{V}|$ places. Given all the labels of $\widetilde{V}$, we can encode the set A by simply concatenating all the labels. Therefore the sum of the lengths of the labels of $\widetilde{V}$ must be $\Omega(n^2)$. However, by our assumption, the sum of their lengths is only $\sum_{j=1}^{|\widetilde{V}|} o(j) = o(n^2)$, a contradiction (for sufficiently large n). □

6 Conclusions and Open Questions

We uncover a wide spectrum of settings and bounds that answer our questions. For Question 1.1, in the simplest case of trees, labeling and embeddings have similar behavior, and both admit prioritization with similar bounds. For the least restricted case of general graphs/metrics, we find similarly that labelings and embeddings exhibit similar behavior across various distortion parameters. However between these two extremes, for ℓ_p spaces, planar graphs and treewidth-k graphs, we see significant separations between labelings and embeddings.

For Question 1.2, we show that labelings admit far superior prioritized versions than their embedding counterparts in all settings other than trees, and most notably for general graphs and for planar/bounded-treewidth graphs, where no prioritized dimension is possible. In ℓ_p spaces, while we did not rule out the possibility of prioritized dimension, we demonstrate a surprising exponential gap between labelings and embeddings (also in the dependence on ϵ).

For Question 1.3 we saw that labeling schemes have prioritized versions, in all cases other than planar graphs where instead of the desired $O(\sqrt{j})$ label size we show that $\Theta(j)$ is surprisingly necessary. For embeddings into ℓ_∞ we showed that for larger distortion some prioritized dimension is possible, even though it is much worse than its labeling counterpart.

Our results leave a few open questions that may be of independent interest:

- How many coordinates are required in order to embed planar graphs – or even treewidth-2 graphs – into ℓ_∞ with distortion $1 + \epsilon$?
- What is the required label size for $1 + \epsilon$ distance labeling for ℓ_p spaces, for $p > 2$?
- Is it possible to embed ℓ_p spaces, $p \in [1, \infty]$, into ℓ_∞ with distortion $1 + \epsilon$ and some prioritized dimension? Theorem 3.3 provided a $j^{\Omega(1/\epsilon)}$ lower bound, but did not rule out this possibility. The same question applies when considering constant distortion.

– All results on embedding of general graphs into ℓ_∞ with both prioritized distortion and dimension (our Theorem 2.2, [9, Theorem 15], and [11, Theorems 2 and 3]) feature prioritized *contractive* distortion. What is possible w.r.t. classic prioritzed distortion (see footnote 3)?

References

1. Abraham, I., Bartal, Y., Neiman, O.: Advances in metric embedding theory. Adv. Math. **228**(6), 3026–3126 (2011)
2. Abraham, I., Filtser, A., Gupta, A., Neiman, O.: Metric embedding via shortest path decompositions. In: 50th Annual ACM SIGACT Symposium on Theory of Computing (Los Angeles 2018), pp. 952–963. ACM, New York (2018)
3. Abraham, I., Gavoille, C.: Object location using path separators. In: 25th Annual ACM Symposium on Principles of Distributed Computing (Denver 2006), pp. 188–197. ACM, New York (2006)
4. Bartal, Y., Filtser, A., Neiman, O.: On notions of distortion and an almost minimum spanning tree with constant average distortion. J. Comput. System Sci. **105**, 116–129 (2019)
5. Bartal, Y., Gottlieb, L.-A.: Approximate nearest neighbor search for ℓ_p-spaces ($2<p<\infty$) via embeddings. Theor. Comput. Sci. **757**, 27–35 (2019)
6. Bourgain, J.: On Lipschitz embedding of finite metric spaces in Hilbert space. Isr. J. Math. **52**(1–2), 46–52 (1985)
7. Deza, M.M., Laurent, M.: Geometry of Cuts and Metrics. Algorithms and Combinatorics, vol. 15. Springer, Berlin (1997)
8. Elkin, M., Filtser, A., Neiman, O.: Terminal embeddings. Theor. Comput. Sci. **697**, 1–36 (2017)
9. Elkin, M., Filtser, A., Neiman, O.: Prioritized metric structures and embedding. SIAM J. Comput. **47**(3), 829–858 (2018)
10. Elkin, M., Neiman, O.: Near isometric terminal embeddings for doubling metrics. In: 34th International Symposium on Computational Geometry (Budapest 2018). Leibniz International Proceedings in Informatics, vol. 99, # 36. Leibniz-Zent. Inform., Wadern (2018)
11. Elkin, M., Neiman, O.: Lossless prioritized embeddings. In: 31st Annual ACM-SIAM Symposium on Discrete Algorithms (Salt Lake City 2020), pp. 1049–1062. SIAM, Philadelphia (2020)
12. Fakcharoenphol, J., Rao, S., Talwar, K.: A tight bound on approximating arbitrary metrics by tree metrics. J. Comput. System Sci. **69**(3), 485–497 (2004)
13. Filtser, A., Gottlieb, L.-A., Krauthgamer, R.: Labelings vs. embeddings: On distributed representations of distances. In: 31st Annual ACM-SIAM Symposium on Discrete Algorithms (Salt Lake City 2020), pp. 1063–1075. SIAM, Philadelphia (2020)
14. Gavoille, C., Peleg, D., Pérennes, S., Raz, R.: Distance labeling in graphs. J. Algorithms **53**(1), 85–112 (2004)
15. Indyk, P.: On approximate nearest neighbors in non-Euclidean spaces. In: 39th Annual Symposium on Foundations of Computer Science (Palo Alto 1998), pp. 148–155. IEEE (1998)
16. Johnson, W., Lindenstrauss, J.: Extensions of Lipschitz mappings into a Hilbert space. In: Conference in Modern Analysis and Probability (New Haven 1982). Contemporary Mathematics, vol. 26, pp. 189–206. American Mathematical Society, Providence (1984)
17. Klein, P.: Preprocessing an undirected planar network to enable fast approximate distance queries. In: 13th Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco 2002), pp. 820–827. ACM, New York (2002)
18. Krauthgamer, R., Lee, J.R., Mendel, M., Naor, A.: Measured descent: a new embedding method for finite metrics. Geom. Funct. Anal. **15**(4), 839–858 (2005)
19. Linial, N., London, E., Rabinovich, Yu.: The geometry of graphs and some of its algorithmic applications. Combinatorica **15**(2), 215–245 (1995)
20. Matoušek, J.: On the distortion required for embedding finite metric spaces into normed spaces. Isr. J. Math. **93**(1), 333–344 (1996)
21. Matoušek, J.: Lecture notes on metric embeddings. Technical report, ETH Zürich (2013). https://kam.mff.cuni.cz/~matousek/ba-a4.pdf

22. Narayanan, S., Nelson, J.: Optimal terminal dimensionality reduction in Euclidean space. In: 51st Annual ACM SIGACT Symposium on Theory of Computing (Phoenix 2019), pp. 1064–1069. ACM, New York (2019)
23. Robertson, N., Seymour, P.: Graph minors. XVI. Excluding a non-planar graph. J. Combin. Theory Ser. B **89**(1), 43–76 (2003)
24. Thorup, M.: Compact oracles for reachability and approximate distances in planar digraphs. J. ACM **51**(6), 993–1024 (2004)
25. Thorup, M., Zwick, U.: Compact routing schemes. In: 13th Annual ACM Symposium on Parallel Algorithms and Architectures (Hersonissos 2001), pp. 1–10. ACM, New York (2001)
26. Thorup, M., Zwick, U.: Approximate distance oracles. J. ACM **52**(1), 1–24 (2005)