

Lossless Prioritized Embeddings

Michael Elkin^{*1} and Ofer Neiman^{†1}

¹Ben-Gurion University of the Negev. Email: {elkinm,neimano}@cs.bgu.ac.il

Abstract

Given metric spaces (X, d) and (Y, ρ) and an ordering x_1, x_2, \dots, x_n of (X, d) , an embedding $f : X \rightarrow Y$ is said to have a *prioritized distortion* $\alpha(\cdot)$, for a function $\alpha(\cdot)$, if for any pair x_j, x' of distinct points in X , the distortion provided by f for this pair is at most $\alpha(j)$. If Y is a normed space, the embedding is said to have *prioritized dimension* $\beta(\cdot)$, if $f(x_j)$ may have at most $\beta(j)$ nonzero coordinates.

The notion of prioritized embedding was introduced by Filtser and the current authors in [EFN18], where a rather general methodology for constructing such embeddings was developed. Though this methodology enabled [EFN18] to come up with many prioritized embeddings, it typically incurs some *loss* in the distortion. In other words, in the worst-case, prioritized embeddings obtained via this methodology incur distortion which is at least a constant factor off, compared to the distortion of the classical counterparts of these embeddings. This constant loss is problematic for isometric embeddings. It is also troublesome for Matousek's embedding of general metrics into ℓ_∞ , which for a parameter $k = 1, 2, \dots$, provides distortion $2k-1$ and dimension $O(k \log n \cdot n^{1/k})$.

In this paper we devise two *lossless* prioritized embeddings. The first one is an *isometric prioritized* embedding of tree metrics into ℓ_∞ with dimension $O(\log j)$, matching the worst-case guarantee of $O(\log n)$ of the classical embedding of Linal et al. [LLR95]. The second one is a prioritized Matousek's embedding of general metrics into ℓ_∞ , which for a parameter $k = 1, 2, \dots$, provides prioritized distortion $2\lceil k \frac{\log j}{\log n} \rceil - 1$ and dimension $O(k \log n \cdot n^{1/k})$, again matching the worst-case guarantee $2k - 1$ in the distortion of the classical Matousek's embedding.

We also provide a dimension-prioritized variant of Matousek's embedding. Finally, we devise prioritized embeddings of general metrics into (single) ultra-metric and of general graphs into (single) spanning tree with

asymptotically optimal distortion.

1 Introduction

An injective function $f : X \rightarrow Y$ between the metric spaces (X, d) and (Y, ρ) is called an *embedding*. The embedding is said to have *distortion* at most α , if for every pair $x, x' \in X$ of distinct points we have $1 \leq \frac{\rho(f(x), f(x'))}{d(x, x')} \leq \alpha$.¹

Study of low-distortion embeddings, initiated with seminal works of [JL84, Ass83, Bou85] in the eighties has flourished since then, and has found numerous algorithmic applications. See, e.g., [LLR95, Bar96, Bar98, FRT04, GKL03, LMN04, ABN11, EFN18, EFN17, BFN16, AFGN18], and the references therein. Closely related to it is the study of metric structures, such as graph spanners and low-stretch spanning trees, distance labeling schemes, distance oracles and routing schemes [PS89, AKPW95, EP04, EEST08, TZ05, Tho01, Pel99, CE06, AB17, EN19], and the references therein.

In all known metric embeddings either the distortion or the dimension of the target metric space (or, in most cases, both) depend at least logarithmically on the cardinality $n = |X|$ of the point set, and moreover, this dependence is typically unavoidable [LLR95, ABN11]. This dependence was always viewed as an inherent downside of the entire approach. To alleviate the issue, Filtser and the current authors [EFN18] introduced² the notions of *prioritized* embeddings and metric structures. For some function $\alpha(\cdot)$, given an ordering x_1, x_2, \dots, x_n

¹This definition applies actually to *non-contracting* embeddings. For a *non-expanding* of Lipschitz embedding, one requires $1/\alpha \leq \frac{\rho(f(x), f(x'))}{d(x, x')} \leq 1$. Yet more generally, there may also be some normalizing term η , so that $\eta \leq \frac{\rho(f(x), f(x'))}{d(x, x')} \leq \alpha \cdot \eta$. The quotient (possibly normalized) $\frac{\rho(f(x), f(x'))}{d(x, x')}$ is called the *distortion* of the pair x, x' under the embedding f .

²There were earlier works [KSW09, ABC⁺05, ABN11] that studied partial and scaling embeddings, that provide guarantees in terms of a parameter $\epsilon > 0$, for all pairs (x, x') such that x' is not among the $\epsilon \cdot n$ closest points to x . However, these embeddings do not allow one to select for which pairs of points a better distortion is provided.

^{*}This research was supported by the ISF grant No. (724/15).

[†]Supported in part by ISF grant No. (1817/17) and by BSF grant No. 2015813.

of the points in the source metric space X , an embedding $f : X \rightarrow Y$ is said to be *prioritized* with *distortion* $\alpha(\cdot)$, if it guarantees that each pair (x_j, x') of distinct points of X admits distortion at most $\alpha(j)$ (as opposed to $\alpha = \alpha(n)$ guaranteed by classical, non-prioritized embeddings). In the case that Y is a normed space, it is said to be *prioritized* with *dimension* $\beta(\cdot)$, if the number of nonzero coordinates in $f(x_j)$ is at most $\beta(j)$. If in addition these nonzero coordinates are the first $\beta(j)$ coordinates of $f(x_j)$, it is called a *cohesive* prioritized dimension.

Even though these notions might look too demanding, surprisingly [EFN18] showed that they are achievable for numerous settings. In particular, a prioritized analogue of Bourgain’s embedding [Bou85, ABN08] shown there states that any n -point metric space can be embedded in $\mathbb{R}^{O(\log n)}$ with prioritized distortion $O(\log j \cdot (\log \log j)^{1/2+\epsilon})$, for an arbitrarily small constant $\epsilon > 0$. This was later refined in [BFN16] to the optimal bound of $O(\log j)$. Moreover, it was shown in [EFN18] that one can achieve prioritized distortion *and* (cohesive) dimension of *poly*($\log j$) simultaneously in embedding general metrics into Euclidean ones.

Another notable embedding shown in [EFN18] is an embedding of general metrics into a (single) tree metric. For the classical scenario a tight bound of $\Theta(n)$ for such embeddings is long well-known [RR98]. In [EFN18] Filtser and the current authors devised for this scenario a prioritized embedding with distortion $O(\alpha(\cdot))$ for any function $\alpha(\cdot)$ that satisfies

$$(1.1) \quad \sum_{j=1}^{\infty} \frac{1}{\alpha(j)} < 1.$$

Moreover, it was shown there that the condition (1.1) is necessary, i.e., if $\alpha(\cdot)$ does not satisfy the condition, no such embedding with prioritized distortion at most $\alpha(\cdot)/8$ is possible. Observe that linear functions do not satisfy the condition (1.1), and thus a certain *loss* when extending classical embeddings to prioritized setting *must* in general be incurred. (Observe also that $\alpha(n) = c \cdot n \cdot \log n (\log \log n)^{1.01}$ does satisfy (1.1) for a constant c , and thus the loss in this case is not too large, i.e., it is at most polylogarithmic in the original classical distortion.) On the other hand, for the FRT embedding of general metrics into ultrametrics [Bar96, Bar98, FRT04], it was shown in [EFN18] that an embedding with prioritized distortion $O(\log j)$ exists, i.e., in this case the loss is at most some constant.

In fact, [EFN18] presented a rather general methodology for converting classical embeddings and metric structures into prioritized ones (with respect to distortion). Roughly speaking, the methodology consists of

two parts. In the first part one devises a *strong terminal* analogue of the classical embedding.³ In the second part, one partitions the point set X into a collection of disjoint subsets $X = \bigcup X_i$ according to the given priority ordering, where the set sizes $|X_1|, |X_2|, \dots$ grow quickly.⁴ One then builds strong terminal embeddings f_i with respect to each of the sets X_i as terminal sets, and carefully combines these embeddings into the ultimate prioritized embedding f .

Typically both these steps incur some loss in distortion. The loss incurred by the first part of this scheme is at most constant, under very general conditions described in [EFN17]. The loss incurred in the second part is not yet well-understood. Specifically, in the prioritized variant of Bourgain’s embedding [EFN18] it is $O((\log \log j)^{1/2+\epsilon})$, for an arbitrarily small $\epsilon > 0$,⁵ while in the prioritized variant of the FRT embedding [EFN18] the loss is at most a constant.

The constant loss in the first part of the scheme has to do with an outer extension, implicitly developed in [EFN17], and explicated in [MMMR18, NN19]. Bi-Lipschitz outer extensions have been a focus of recent research [MMMR18, NN19, EN18], where they were studied in the context of Johnson-Lindenstrauss dimension reduction [MMMR18, NN19], and in the context of doubling metrics [EN18]. In both these contexts it was shown that the loss can be made at most $1 + \epsilon$, for an arbitrarily small $\epsilon > 0$. However, in general, Mahabadi et al. [MMMR18] showed that it is unfortunately not the case, as there are scenarios when Bi-Lipschitz outer extensions incur loss at least c , for some constant $c > 1$.

Generally speaking, the loss is most troublesome in two scenarios. The first one is when the embedding is isometric or near-isometric (i.e., has distortion $1 + \epsilon$, for an arbitrarily small $\epsilon > 0$), as it is the case for the Johnson-Lindenstrauss dimension reduction [MMMR18, NN19] and for embedding of doubling metrics into ℓ_∞ [EN18]. The second scenario is when the specific leading constant factor in the distortion is important. Consider Matousek’s embedding [Mat96] of general metrics into ℓ_∞ . Given an integer parameter $k = 1, 2, \dots$, this embedding has distortion $2k - 1$ and dimension $O(kn^{1/k} \cdot \log n)$. Employing the generic Lipschitz outer extension of [EFN17] on this embedding results in a prioritized embedding with distortion

³We say that an embedding $f : X \rightarrow Y$ has *terminal* distortion $\alpha(\cdot)$ with respect to a subset $K \subseteq X$, $|K| = k$, of k terminals, if it guarantees distortion $\alpha(k)$ for all pairs $(x, x') \in K \times X$. The embedding has *strong terminal* distortion, if it has terminal distortion, and in addition it has distortion $\alpha(n)$ for all pairs of points.

⁴The specific growth rate depends on the embedding at hand.

⁵The improvement in [BFN16] was achieved by a different approach.

tion $4 \cdot (2^{\lceil \frac{k \log j}{\log n} \rceil} - 1)$ and dimension $O(k^2 \cdot n^{1/k} \cdot \log n)$. (See Section 3 for more details.) In other words, in this case there is a significant loss both in the distortion and in the dimension. (Note that for distortion $8k - 4$ the classical Matousek embedding provides dimension $O(k \cdot n^{1/4k} \log n)$, i.e., *polynomially smaller* than the dimension obtained via the generic methodology.)

In this paper we address both these scenarios. We devise an isometric embedding of tree metrics into ℓ_∞ with cohesive *prioritized dimension* $O(\log j)$. Note that in the context of *isometric* embeddings, no loss whatsoever can be tolerated. This is unlike the case of near-isometric embeddings, where one can afford some loss in the distortion. This is indeed the case in the terminal Johnson-Lindenstrauss dimension reduction of Mahabadi et al. [MMMR18] and Narayanan and Nelson [NN19], where the distortion grows from $1 + \epsilon$ to $1 + O(\sqrt{\epsilon})$ in [MMMR18] and to $1 + O(\epsilon)$ in [NN19]. The loss is similar (i.e., $1 + \epsilon$ becomes $1 + O(\epsilon)$) in the terminal embedding of [EN18] (by the current authors) of doubling metrics to ℓ_∞ . Note also that the aforementioned embeddings [MMMR18, NN19, EN18] are *not prioritized*⁶, while our embedding of tree metrics into ℓ_∞ is. This embedding of ours is both isometric and *lossless prioritized*, while previously there were no known strong terminal or prioritized isometric embeddings. Moreover, the prioritized dimension $O(\log j)$ of our embedding is optimal [Mat02].

Our second contribution addresses the second scenario discussed above. Specifically, we devise a *lossless prioritized* variant of Matousek’s embedding. We show that for any positive integer parameter $k = 1, 2, \dots$, there is an embedding of general metrics into ℓ_∞ with prioritized distortion $2^{\lceil \frac{k \log j}{\log n} \rceil} - 1$ and dimension $O(k \cdot n^{1/k} \cdot \log n)$. Observe that by substituting $j = n$ one obtains the parameters of classical Matousek’s embedding *without any loss* in distortion, and the dimension remains asymptotically the same.

We devise a dimension-prioritized variant of Matousek’s embedding. Specifically, this variant of our embedding, for any positive integer parameter $k = 1, 2, \dots$, provides distortion $2k \log \log j + 1$ and prioritized dimension $O(k \cdot j^{1/k} \cdot \log n)$, and also a variant with a slightly larger cohesive prioritized dimension. (This embedding does leave some room for improvement. One can hope to obtain distortion $2k - 1$ and dimension as above, but with $\log n$ replaced by $\log j$.)

⁶It is moreover open if these embeddings can be made prioritized. This question was explicitly studied in [MMMR18], where a prioritized version of JL dimension reduction was shown, albeit with a much larger (than $1 + \epsilon$) distortion of $O(\log \log j)$, and under a weaker notion of prioritized distortion than the one we study here.

Finally, our third contribution has to do with the aforementioned prioritized embedding of [EFN18] of general metrics into (single) tree metrics. As was discussed above, this embedding provides prioritized dimension $\alpha(\cdot)$ for any function α that satisfies the condition (1.1), and this condition is tight. In our opinion, this embedding is a *Drosophila* of prioritized embeddings, i.e., it is fundamental to the area, and deserves further exploration.

Most notable deficiency of this construction of [EFN18] is that if the original metric is induced by a graph, this is not necessarily the case for the tree that the construction provides. Rather this tree may contain edges that are not present in the original graph. The problem whether one can always achieve a single *spanning* tree with optimal prioritized stretch was left open by [EFN18].

In this paper we answer this question in the affirmative, and devise an embedding of general metrics into a (single) ultrametric (see Section 4.1 for the definition), and an embedding of general graphs into a (single) *spanning* tree. Both these embeddings have the same (tight up to constant factors) prioritized distortion $O(\alpha(\cdot))$, for any function $\alpha(\cdot)$ that satisfies the condition (1.1). (The tightness follows from [EFN18].)

1.1 Technical Overview

Embedding General Metrics into ℓ_∞ . The embedding of Matousek [Mat96] is a Frechet one, i.e., for every index i , the i th coordinate of an image $f(x)$ of an input point $x \in X$ is a distance from x to a certain subset $S_i \subseteq X$. The set S_i (henceforth, a *Frechet* subset) is constructed by sampling points from X independently at random with appropriate probabilities. In the analysis of distortion for a fixed pair of points $x, x' \in X$ such that $d(x, x') = d$, one sets $\delta = \frac{d}{2k-1}$. One then considers pairs of balls $(B_0 = B(x, 0 \cdot \delta), B'_1 = B(x', 1 \cdot \delta)), (B'_1, B_2 = B(x, 2 \cdot \delta)), \dots, (B_{k-1}, B'_k = B(x', k \cdot \delta))$.⁷ Intuitively, for some index $i \leq k - 1$, the ratio between the sizes of B'_{i+1} and B_i is at most $n^{1/k}$, and at this point, with high probability, one of the two balls hits a Frechet set S_j of points selected with sampling probability $n^{-(1-i/k)}$, while the other ball does not. Therefore, in the respective coordinate $f(x)$ and $f(x')$ (where $f(\cdot)$ is the ultimate Lipschitz embedding) will differ by at least δ , ensuring that the distortion is at most $2k - 1$.

In our prioritized variant of this embedding we carefully design sampling probabilities so that higher-priority points are more likely to be included in the sets S_j than lower-priority ones. In the analysis of distortion

⁷For a point $x \in X$, and a non-negative r , we denote by $B(x, r) = \{y \mid d(x, y) \leq r\}$ the *radius- r ball around x* .

for a pair (x, x') that contains a high-priority point x , we build a different, much shorter sequence of ball-pairs $(B_i, B'_{i+1}), (B'_{i+1}, B_{i+2}), \dots, (B_{k-1}, B'_k)$, where the index i depends on the priority of x . The balls now will have radii that are multiples of $\delta_i = \frac{d}{2^{(k-i)-1}}$, as opposed to $\delta = \frac{d}{2^{k-1}}$ used in the classical embedding of Matousek. In other words, while Matousek's analysis of distortion for a pair (x, x') with $d(x, x') = d$ considers balls of radii $0, \frac{d}{2^{k-1}}, \frac{2d}{2^{k-1}}, \dots, \frac{(k-1)d}{2^{k-1}}$, our analysis of the distortion considers a sequence of radii that depends not just on the distance between x and x' , but also on the priorities of these points. Since this sequence of radii is shorter than k for pairs containing at least one high-priority point, the resulting distortion is smaller than $2k - 1$. Moreover, for points with very high priorities (i.e., priorities at most $n^{1/k}$), the resulting distortion between them and *all other points* is exactly 1. For points with priorities between $n^{1/k}$ and $n^{2/k}$, all pairs involving them admit distortion at most 3, etc..

Our approach also leads to *dimension-prioritized* variant of Matousek's embedding. As high-priority points are more likely to be included in the Frechet sets, they end up having more zeroes in the respective coordinates. However, the argument becomes insufficient when the sampling probabilities are very small, e.g., $1/n$. So we need to have only very few Frechet sets with such a small sampling probability (rather than $\approx n^{1/k}$ such sets in the standard embedding). In order to allow the distortion analysis to go through, the idea is to refine the sampling probabilities, and introduce $\log \log n$ scales that have smaller "gaps" between consecutive levels as the probabilities get closer to $1/n$. This increased number of scales translates to the (prioritized) factor of $\log \log j$ in the distortion.

Embedding of trees to ℓ_∞ with cohesive prioritized dimension. The isometric embedding of a tree into ℓ_∞ [LLR95] is based on centroids: every tree T has a vertex s , such that there are two trees T_1, T_2 who cover T and share only s , each of size at most $2|T|/3$. One can define an embedding to \mathbb{R} that preserves exactly distances between each $u \in T_1$ to every $v \in T_2$, and then embed the two trees recursively.

How can we guarantee that high priority points will receive few nonzero coordinates in such an embedding? We can try to follow the [EFN18] methodology described above: partition the vertex sets into few terminal sets K_1, K_2, \dots , and apply a strong terminal embedding on each K_i . However, it is not clear how to combine the different embeddings without incurring large dimension $\Theta(\log n)$ for all points. One possible attempt is to contract all the vertices in K_i to a single vertex z , and only then apply the embedding for K_{i+1} , ensuring that z will be mapped to 0 in that embedding.

The issue is that contraction may completely change the distances.

To handle this problem, we introduce a novel graph operation, that we call *folding*. Rather than contracting K_i , we iteratively "fold" the vertices in K_i one to the other, keeping half of the shortest path between them intact. While this operation can also change distances, we prove a structural lemma characterizing which pairs may be affected, and ensure that such pairs are handled by the previous level embedding.

Prioritized embedding of arbitrary metrics into a single spanning tree. The method to build a prioritized embedding into a single metric tree in [EFN18] was to re-weigh the edges (x, y) of the input metric X according to the priorities of their endpoints. Then one computed an MST of the resulting graph, and this MST was returned as the tree with the desired prioritized distortion. It is not clear if this approach can be used to construct a *spanning* tree with a similar property.

On the other hand, the approach that we take here is first to recursively construct an ultrametric with near-optimal prioritized distortion. This is done using hierarchical ball partitions; In each iteration we carve a suitable ball, setting the distance between separated pairs to be the diameter, and continue recursively on the two parts. The ball's radius is chosen carefully as to not separate points if their distortion will be too large. The construction of a spanning tree then imitates the construction of the ultrametric by using the mechanism of petal decompositions [AN19] instead of ball partitions. In other words, we carefully carve petals (as opposed of balls), while making sure that we never violate the conditions of prioritized distortion.

1.2 Bibliographic Note In an independent and concurrent work, [FGK20] also showed an isometric embedding of a tree into ℓ_∞ with prioritized dimension $O(\log j)$. They have several additional interesting results on prioritized embedding and distance labeling, e.g., an embedding of any metric space into ℓ_∞ with prioritized distortion $2 \log j$ and prioritized dimension $2j$, and also several lower bounds. See [FGK20] for more details.

2 Prioritized Embedding of Trees Into ℓ_∞

2.1 Path Folding For a graph $G = (V, E)$, *identifying* two vertices $x, y \in V$ means replacing them by a single vertex z , which is connected to every neighbor of x and of y (keeping only the shorter edge among parallel edges). We often abuse notation and denote the vertex set of the new graph obtained after this operation by V as well, that is, we may refer to the new vertex z as

either x or y .

Given a graph $G = (V, E)$, two vertices $u, v \in V$, and some shortest path P from u to v of length t , for each $0 \leq \alpha \leq 1$ let x_α be the (possibly imaginary) vertex at distance $\alpha \cdot t$ from u on the path P . (So that $x_0 = u$ and $x_1 = v$.) Define the *folding* of P as the operation that produces a graph G' out of G by identifying, for every $0 \leq \alpha \leq 1/2$, the vertex x_α with the vertex $x_{1-\alpha}$. The vertex $x_{1/2}$ is called the *folding point*. See [Figure 1](#).

CLAIM 2.1. *Let T be a tree. Then for any path P in T , if T' is obtained from T by folding P , then T' is a tree as well.*

Proof. It is easy to see that T' is connected (identifying vertices cannot disconnect a graph). Seeking contradiction, let C' be a cycle in T' . Clearly C' must contain some of the vertices $\{x_\alpha\}$, as otherwise C' would be a cycle in T . We now create a cycle in T , and derive a contradiction. For every (non-imaginary) vertex x_α in C' : if both edges of C' touching it are edges of T that touch x_α (respectively, $x_{1-\alpha}$), we do nothing. Otherwise, we replace x_α by the subpath of P from x_α to $x_{1-\alpha}$. After all these replacements, we get a cycle in T . \square

We also consider multiple foldings: given a set $K \subseteq V(T)$ of designated vertices in the tree T , a *K-folding* generates a tree \hat{T} in which all vertices of K are identified as the same vertex. This is done by iteratively selecting two (arbitrary) vertices in K that are not identified yet, and folding them along the unique path between them, so that they become identified. Observe that [Claim 2.1](#) implies that *K-folding* indeed generates a tree, but that tree is not unique.

For $u, v \in T$ let P_{uv} be the unique path between them in T . The purpose of the next claim is to show that folding a path in a tree can change distances only for pairs u, v whose path P_{uv} in T crosses (non-trivially) the folding point. For a path P with a folding point $x = x_{1/2}$, we denote by P^- the first half of the path: the interior of the path from x_0 to x . Denote by P^+ the interior of the path from x to x_1 . We say that a pair of vertices u, v crosses x , if P_{uv} intersects both P^- and P^+ .

CLAIM 2.2. *Let T be a tree, and T' the tree obtained by folding some path P . If $x = x_{1/2}$ is the folding point, then for any $u, v \in T$ that do not cross x , we have*

$$d_T(u, v) = d_{T'}(u, v)$$

Proof. We first observe that identifying vertices is equivalent (distance-wise) to adding a zero-weight edge between them. So if G' is the graph obtained from T

by adding zero weight edges between every vertex x_α to $x_{1-\alpha}$, it follows that $d_{T'}(y, z) = d_{G'}(y, z) \leq d_T(y, z)$ for every $y, z \in V$. Clearly, if P_{uv} does not intersect P then $d_{T'}(y, z) = d_T(y, z)$. Otherwise, let $\alpha_1 \leq \alpha_2 \leq 1$ be the first (resp., last) index of a vertex of P_{uv} on P . By our assumption, either both $\alpha_1, \alpha_2 \leq 1/2$ or both $\alpha_1, \alpha_2 \geq 1/2$. The main observation is that the zero-weight edges added to T will not shorten the distance between x_{α_1} to x_{α_2} (i.e. $d_T(x_{\alpha_1}, x_{\alpha_2}) = d_{T'}(x_{\alpha_1}, x_{\alpha_2})$). This is because the folded path is a shortest path. The claim follows. \square

COROLLARY 2.1. *If T is a tree and \hat{T} a K -folding of T that uses $k - 1$ folding points $X = \{x^{(1)}, \dots, x^{(k-1)}\}$, then for any pair $u, v \in T$ that do not cross any point in X , we have $d_{\hat{T}}(u, v) = d_T(u, v)$.*

2.2 Prioritized Embedding We now define an isometric embedding of T into ℓ_∞ with cohesive prioritized dimension $O(\log j)$. That is, given a ranking (v_1, \dots, v_n) of the vertices of T , vertex v_j will be nonzero in at most $O(\log j)$ coordinates, and moreover, these coordinates will be the first $O(\log j)$ coordinates of the image of v_j . The main technical lemma provides, for a given set of terminals $K \subseteq V$, two basic constructions: 1) A *folding* of the tree T along $|K| - 1$ paths, so that in the tree \hat{T} created by these foldings, all terminals are identified as a single vertex. 2) An embedding which is isometric on every pair whose shortest path in T crosses a folding point.

We will use the following well-known claim (see, e.g., [\[LLR95\]](#)).

CLAIM 2.3. *Let $T = (V, E)$ be a tree, and $K \subseteq V$ a set of vertices. Then there exist subtrees T_1 and T_2 of T that share a single vertex and together cover T , such that each of T_1, T_2 contains at most $\lceil 2|K|/3 \rceil$ vertices of K .*

LEMMA 2.4. *Let $T = (V, E)$ be a tree, and $K \subseteq V$ a set of k terminals. Then there exists a non-expanding embedding $f : T \rightarrow \ell_\infty^{O(\log |K|)}$ and a tree \hat{T} which is a K -folding of T , so that for every $u, v \in V$ at least one of the following holds:*

1. $\|f(u) - f(v)\|_\infty = d_T(u, v)$.
2. $d_{\hat{T}}(u, v) = d_T(u, v)$.

Proof. Let X be the set of $k - 1$ folding points in some fixed K -folding of T . By [Corollary 2.1](#) it suffices to show a non-expansive embedding $f : T \rightarrow \ell_\infty^{O(\log k)}$ that preserves exactly distances between pairs $u, v \in V$ whose shortest path crosses a folding point in X . We assume w.l.o.g. that the path in T between any




Figure 1: An example of path folding. In the left is a tree with the $u - v$ path marked, on the right is the tree obtained by folding this path, with $x_{1/2}$ the folding point.

two distinct points $x, y \in X$ contains a vertex in its interior (if it does not, add an imaginary vertex on the edge (x, y)). Consider the $2(k - 1)$ half paths $P_1^-, P_1^+, \dots, P_{k-1}^-, P_{k-1}^+$ in the K -folding of T . We prove by induction on b , that if H is a subtree of T with $b = |H \cap X|$, then it admits a non-expansive embedding into $\mathbb{R}^{4(\log b+1)}$ with distortion 1 (in the ℓ_∞ norm) for all pairs whose path in H crosses (at least) one of these b vertices. Clearly a path in H cannot cross a vertex which is not in H .

The base case is when $b = 1$ (when $b = 0$ we can just map all vertices to 0). Let $x \in X$ be the single point in $H \cap X$, and let P^- and P^+ be the two parts of the path P for which x is the middle point. We note that for any vertex $u \in H$, the unique path from x to u can intersect at most one of P^- or P^+ (since both of these paths emanate from x). Let $R \subseteq H$ be the set of vertices $u \in H$ such that the path from x to u intersects P^- , and define the embedding $f : H \rightarrow \mathbb{R}$ by

$$f(u) = \begin{cases} d_H(x, u) & u \in R \\ -d_H(x, u) & \text{otherwise} \end{cases}$$

We first observe that if $u \in R$ and $v \notin R$, then the path between u, v must contain x . This is because the very first edge e on the path from x to u is on P^- , so if the path from u to v does not pass through x , then the path from x to v would contain e , a contradiction to the fact that $v \notin R$. This observation implies that f is non-expansive: for pairs with both $u, v \in R$ or both $u, v \notin R$ by the triangle inequality, and for pairs $u \in R, v \notin R$ we have $d_H(u, v) = d_H(u, x) + d_H(x, v) = |f(u) - f(v)|$. The latter proves that pairs u, v whose path crosses x

(recall the path must intersect both P^- and P^+) attain distortion 1, as required.

For the induction step, let H_1 and H_2 be the subtrees of H obtained by applying Claim 2.3 on the tree H with respect to the terminal set $H \cap X$ (we stress that this terminal set is not the given terminal set K), and let $s \in H$ be the unique vertex common to both subtrees. Let $b_1 = |X \cap H_1|$ and $b_2 = |X \cap H_2|$, by Claim 2.3 we have that both $b_1, b_2 \leq b' = \lceil 2b/3 \rceil$. (Note that in the case $b = 2$ we have $\lceil 2b/3 \rceil = 2$ as well, but since we ensured that there is at least one vertex on the path between any two points of X , this vertex can serve as a separator, and each H_1, H_2 will contain only one point of X .) Let $f_1 : H_1 \rightarrow \mathbb{R}^{4 \log b'+1}$ and $f_2 : H_2 \rightarrow \mathbb{R}^{4 \log b'+1}$ be the embeddings guaranteed by the induction hypothesis. We may shift these embeddings so that $f_1(s) = f_2(s) = 0$, without affecting the distortion. Note that for $b > 2$ we have that $4 \log \lceil 2b/3 \rceil + 1 \leq 4 \log b$ (and for $b = 2$, we have that both H_1, H_2 contain a single vertex of X , so by the base case, f_1 and f_2 are embeddings into \mathbb{R}). Thus we can define $f : H \rightarrow \mathbb{R}^{4 \log b+1}$ by adding a single coordinate $h : H \rightarrow \mathbb{R}$ with

$$h(u) = \begin{cases} d_H(s, u) & u \in H_1 \\ -d_H(s, u) & u \in H_2 \end{cases}$$

Define the embedding for $u \in H_i, i \in \{1, 2\}$, by $f(u) = f_i(u) \oplus h(u)$. We claim that f is indeed non-expansive. To see this, consider first a pair $u, v \in H_i$, then by the inductive hypothesis f_i is non-expansive, and the additional coordinate h is also non-expansive by the triangle inequality. Consider now a pair $u \in H_1$ and $v \in H_2$. As before we have that $d_H(u, v) =$

$d_{H_1}(u, s) + d_{H_2}(s, v)$. Since s belongs to both trees and $f_1(s) = f_2(s) = 0$, it follows that

$$\begin{aligned} & \|f_1(u) - f_2(v)\|_\infty \\ &= \|f_1(u) - f_1(s) + f_2(s) - f_2(v)\|_\infty \\ &\leq \|f_1(u) - f_1(s)\|_\infty + \|f_2(s) - f_2(v)\|_\infty \\ &\leq d_{H_1}(u, s) + d_{H_2}(s, v) = d_H(u, v). \end{aligned}$$

For the additional coordinate h we have $|h(u) - h(v)| = d_H(u, s) + d_H(s, v) = d_H(u, v)$. So f is indeed an isometry for such u, v . By the induction hypothesis, every pair $u, v \in H_i$, $i \in \{1, 2\}$, for which the path from u to v crosses some $x \in H \cap X$, has distortion 1 by f_i , and thus also by f . \square

THEOREM 2.1. *Let T be a tree on n vertices with priority ordering (v_1, \dots, v_n) . Then there is an isometric embedding of T into ℓ_∞ with cohesive prioritized dimension $O(\log j)$.*

Proof. We partition the vertices of T into $\lceil \log \log n \rceil$ sets $K_1, \dots, K_{\lceil \log \log n \rceil}$ by taking $K_i = \{v_j : 2^{2^{i-1}} < j \leq 2^{2^i}\}$ (put v_1, v_2 into the first set, and the last set contains vertices only until v_n). Initially, let $T_1 = T$, $z = v_1$, and in the i -th step, apply the embedding of Lemma 2.4 on the tree T_i with the set $K_i \cup \{z\}$. Define $T_{i+1} = \hat{T}_i$ the tree of the next iteration as the $K_i \cup \{z\}$ -folding of T_i . Let $f_i : T_i \rightarrow \mathbb{R}^{O(2^i)}$ be the embedding of T_i , we may assume by an appropriate shift that $f_i(z) = 0$. Note that $z \in T_i$ is the unique vertex corresponding to the terminals $v_1, \dots, v_{2^{2^i-1}}$. The embedding is defined as $f = \bigoplus_i f_i$.

Isometry. To see that f is an isometry, fix any $u, v \in T$. Since $T_{\lceil \log \log n \rceil + 1}$ is a single vertex, surely there is an index i such that $d_{T_i}(u, v) \neq d_{T_{i+1}}(u, v)$, and by the first property of Lemma 2.4 we get that $\|f_i(u) - f_i(v)\|_\infty = d_{T_i}(u, v)$ for the first such index i . Note that each f_j is non-expanding, therefore so is f , and the isometry follows.

Cohesive Prioritized dimension. Recall that the embedding f_i applied with the terminal set $K_i \cup \{z\}$ has $O(\log |K_i|) = O(2^i)$ coordinates. Consider the vertex with priority j , $v_j \in T$, and let i be such that $v_j \in K_i$. We have that $i \leq \log \log j + 1$. Observe that for every $i' > i$, since $v_j \in K_i$ we have $f_{i'}(v) = f_{i'}(z) = 0$. In particular, the total number of nonzero coordinates for v_j in f is at most

$$\sum_{i'=0}^i O(2^{i'}) = O(2^i) = O(\log j),$$

and these are the first coordinates. So the embedding is indeed a cohesive prioritized embedding. \square

3 Prioritized Embeddings of General Metrics into ℓ_∞

In this section we devise prioritized counterparts of the embeddings of Matousek [Mat96] into ℓ_∞ . Matousek showed that for any $k \geq 1$, any n -point metric embeds into ℓ_∞ with distortion $2k - 1$ and dimension $O(k \cdot n^{1/k} \cdot \log n)$. We will show an adaptation of this embedding that prioritizes the distortion, and a more involved construction that prioritizes the dimension (and slightly the distortion).

We note that at the cost of increasing the distortion by a constant factor, one can use the framework introduced in [EFN18] to obtain prioritized distortion $8 \left\lceil \frac{k \cdot \log j}{\log n} \right\rceil - 4$ and dimension $O(k^2 \cdot n^{1/k} \cdot \log n)$. This is achieved by defining $K_i = \{x_j : j \leq n^{i/k}\}$, and applying as a black-box the embedding of Matousek on each K_i with distortion $2i - 1$ and dimension $O(k \cdot (n^{i/k})^{1/i} \cdot \log n^{i/k}) \leq O(k \cdot n^{1/k} \cdot \log n)$ for any $x_j \in K_i$. Then extend the embedding of each K_i to the entire set X , increasing the distortion by a factor of 4 to $8i - 4 \leq 8 \left\lceil \frac{k \cdot \log j}{\log n} \right\rceil - 4$ for pairs containing $x_j \in K_i$. Observe, however, that in our result, with this distortion the dimension is just $O(kn^{1/4k} \log n)$, i.e., *polynomially* smaller than the above bound. We remark that we know a simple extension that doubles the number of coordinates, achieving distortion $4 \left\lceil \frac{k \cdot \log j}{\log n} \right\rceil - 1$, to be described in the full version.

Alternatively, one may use the equivalence between coarse scaling distortion and prioritized distortion [BFN16], and a result from [ABN11], to derive prioritized distortion $c \cdot \left\lceil \frac{k \cdot \log j}{\log n} \right\rceil$ with improved dimension $O(n^{1/k} \cdot \log n)$, but the leading constant c in the distortion is large.⁸ Again, our embedding below can have this distortion with dimension $O(k \cdot n^{\frac{2}{c \cdot k}} \cdot \log n)$, which is significantly better for $k = o\left(\frac{\log n}{\log \log n}\right)$.

We next provide a lossless prioritized result, that does not increase the worst case distortion *at all*. Moreover, the aforementioned techniques cannot provide prioritized dimension, while our approach provides prioritized dimension as well. We note also that our result can also be used in conjunction with the aforementioned equivalence of [BFN16] to derive an embedding of general metrics into ℓ_∞ with a coarse scaling distortion $c' \cdot k \cdot \left\lceil \frac{\log 2/\epsilon}{\log n} \right\rceil$, for a small fixed constant c' (obtained from the reduction of [BFN16]) and dimension $O(kn^{1/k} \log n)$. This compares favorably with the coarse scaling embedding of [ABN11] for $k = o\left(\frac{\log n}{\log \log n}\right)$. (In

⁸It is unspecified in [ABN11], but we believe that it is at least, say, 50.

the latter embedding the dimension is $O(n^{1/k} \log n)$, but the coarse scaling distortion is $c \cdot k \cdot \lceil \frac{\log 2/\epsilon}{\log n} \rceil$, for a much larger constant c .)

3.1 Prioritized Distortion Let (X, d) be a metric space with $|X| = n$, and let (x_1, \dots, x_n) be the given priority. Let $k \geq 1$ be an integer parameter. For $1 \leq i \leq k$ define $S_i = \{x_j : n^{(i-1)/k} < j \leq n^{i/k}\}$ (and put x_1 in S_1).

Construction. We introduce many "copies" of high priority points, specifically, each point in S_i will have $(2^k \cdot n)^{1-i/k}$ copies. This gives us a (semi)-metric (X', d') , by defining the distance between a point to all its copies as 0. Since there are at most $n^{i/k}$ points in S_i , the total number of points in X' , denoted $N = |X'|$, is at most

$$(3.2) \quad N \leq \sum_{i=1}^k (2^k \cdot n)^{1-i/k} \cdot n^{i/k} = 2^k \cdot n \cdot \sum_{i=1}^k 2^{-i} \leq 2^k \cdot n.$$

For each $1 \leq i \leq k$ we sample $m = c \cdot N^{1/k} \cdot \ln N$ (Frechet) sets $A_1^{(i)}, \dots, A_m^{(i)}$, where c is a constant to be determined later. For every $1 \leq h \leq m$, each element of X' is sampled to $A_h^{(i)}$ independently with probability $N^{-i/k}$. The embedding $f : X \rightarrow \mathbb{R}^{km}$ is defined by

$$f(x) = \bigoplus_{i=1}^k \bigoplus_{h=1}^m d(x, A_h^{(i)}).$$

Analysis. By the triangle inequality, we have that the embedding is non-expansive in each coordinate, thus in ℓ_∞ it is non-expansive. It remains to show that, with high probability, for every $1 \leq j \leq n$, any pair containing x_j has distortion at most $2 \lceil \frac{k \cdot \log j}{\log n} \rceil - 1$. Fix such a pair $x_j, y \in X$. Let i be such that $x_j \in S_i$, and define $r = \frac{d(x_j, y)}{2^{i-1}}$. Consider the balls $B_0 = B(x_j, 0)$, $B_1 = B^\circ(y, r)$, $B_2 = B(x_j, 2r)$, and so on until B_i . The balls of even index are closed and centered at x_j , while those of odd index are open (denoted by B°) and centered at y . The radius of B_b , $0 \leq b \leq i$, is $b \cdot r$. By the definition of r , $B_{b-1} \cap B_b = \emptyset$.

Let $1 \leq b \leq i$ be an integer such that $|B_{b-1}| \geq N^{1-(i-b+1)/k}$ and $|B_b| \leq N^{1-(i-b)/k}$. Such an index b must exist since $x_j \in S_i$ has $(2^k \cdot n)^{1-i/k}$ copies, and

$$|B_0| = (2^k \cdot n)^{1-i/k} \stackrel{(3.2)}{\geq} N^{1-i/k},$$

so if no such b exists, it follows that $|B_1| > N^{1-(i-1)/k}$, and so on until $|B_i| > N^{1-(i-i)/k} = N$, a contradiction.

For each $1 \leq h \leq m$, let \mathcal{E}_h^{hit} be the event that $B_{b-1} \cap A_h^{(k-i+b)} \neq \emptyset$, and let \mathcal{E}_h^{miss} be the event that $B_b \cap A_h^{(k-i+b)} = \emptyset$. Since the balls are disjoint and each

point joins the $\{A_h^{(i)}\}$ sets independently, these events are independent. We calculate:

$$\begin{aligned} \Pr[\mathcal{E}_h^{hit}] &= 1 - \left(1 - N^{-(k-i+b)/k}\right)^{|B_{b-1}|} \\ &\geq 1 - e^{-N^{-(k-i+b)/k} \cdot N^{1-(i-b+1)/k}} \\ &\geq N^{-1/k}/2. \end{aligned}$$

$$\begin{aligned} \Pr[\mathcal{E}_h^{miss}] &= \left(1 - N^{-(k-i+b)/k}\right)^{|B_b|} \\ &\geq \left(1 - N^{-1+(i-b)/k}\right)^{N^{1-(i-b)/k}} \\ &\geq 1/4. \end{aligned}$$

Thus, the probability that there is no $1 \leq h \leq m$ for which both events occur, is bounded by

$$\left(1 - N^{-1/k}/8\right)^m \leq e^{-N^{-1/k}/8 \cdot c \cdot N^{1/k} \ln N} = 1/N^2,$$

whenever $c \geq 16$. If there is an $1 \leq h \leq m$ such that both events occur, then for odd b we have $d(x_j, A_h^{(k-i+b)}) \leq (b-1)r$ and $d(y, A_h^{(k-i+b)}) \geq br$ (for even b replace the roles of x_j, y). It follows that

$$\begin{aligned} |d(y, A_h^{(k-i+b)}) - d(x_j, A_h^{(k-i+b)})| &\geq br - (b-1)r \\ &= r = \frac{d(x_j, y)}{2^{i-1}}. \end{aligned}$$

Recall that $x_j \in S_i$ means that $n^{(i-1)/k} < j \leq n^{i/k}$, or equivalently, $i-1 < k \cdot \frac{\log j}{\log n} \leq i$. So the prioritized distortion is indeed at most $2 \lceil \frac{k \cdot \log j}{\log n} \rceil - 1$. The dimension we obtain is $k \cdot m = O(k \cdot N^{1/k} \cdot \ln N) = O(k \cdot n^{1/k} \cdot \ln n)$. This is since $N^{1/k} = (2^k \cdot n)^{1/k} = 2n^{1/k}$, and $\ln N = \ln(2^k \cdot n) = O(k + \ln n) = O(\ln n)$, using that it makes no sense to consider $k > \log n$.

We proved the following.

THEOREM 3.1. *For any parameter $k \geq 1$, any n -point metric embeds into ℓ_∞ with prioritized distortion $2 \lceil \frac{k \cdot \log j}{\log n} \rceil - 1$ and dimension $O(k \cdot n^{1/k} \cdot \log n)$.*

3.2 Prioritized Dimension Note that in the construction of Section 3.1, a point in S_i has $(2^k \cdot n)^{1-i/k}$ copies, which intuitively means it is very unlikely that all its copies are missed by the a Frechet set $A_h^{(b)}$ for $b < k - i$ (that includes each point independently with probability $N^{-b/k}$). If a Frechet set hits at least one copy, it means that such a point will have a value of zero in the corresponding coordinate (the distance to that Frechet set is zero). However, it might get a high number of nonzeros from the sets with larger value of b .

Recall that we repeat $O(N^{1/k} \ln N)$ times the random choices for each density $1 \leq b \leq k$. In particular, the last density level $b = k$ is very likely to incur high dimension for every point: all the copies of x_1 , say, are expected to be missed by $\Omega(N^{1/k})$ of the sets in $\{A_h^{(k)}\}_{h=1}^m$ (and thus have many nonzero coordinates).

To alleviate this issue, we refine the sampling probabilities, and enforce exponentially smaller gaps between neighboring densities, as the exponent becomes closer to -1 . Specifically, we will define the priority sets S_i for every $0 \leq i < \log \log n$ by

$$S_i = \{x_j : 2^{2^i} < j \leq 2^{2^{i+1}}\}.$$

Note that x_1, x_2 do not belong to any of these sets. To handle this, simply introduce additional 2 coordinates, and map each $x \in X$ to the vector $(d(x, x_1), d(x, x_2))$. Clearly this is non-expanding, and has distortion 1 for pairs containing x_j with $1 \leq j \leq 2$. So in all that follows we only care for pairs containing x_j for $j > 2$.

Now, each point $x \in S_i$ will have $C(i) = \frac{n(\log \log n + 1)^2}{2^{2^{i+1}} \cdot (i+2)^2}$ copies. Observe that $C(i) \geq 1$ for all $i < \log \log n$ (assuming $\log \log n$ is an integer).⁹ Since $|S_i| \leq 2^{2^{i+1}}$, we have that the total number of points in the new metric (X', d') is

$$\begin{aligned} N &\leq \sum_{i=0}^{\log \log n} C(i) \cdot 2^{2^{i+1}} \\ &= n(\log \log n + 1)^2 \cdot \sum_{i=0}^{\log \log n} \frac{1}{(i+2)^2} \\ &\leq n(\log \log n + 1)^2. \end{aligned}$$

For every $0 \leq i \leq \log \log n$ we define $R(i) = c \cdot 2^{(2^i+2)/k} \cdot \ln n$ as the number of samples (Frechet sets) to be taken with the appropriate density corresponding to i . Note that when $i = \log \log n$ we have $R(i) = \Theta(n^{1/k} \cdot \ln n)$. This is the number of Frechet sets on every density level in our construction from Section 3.1 (and in Matousek's original construction). On the other hand, for small values of i , e.g. $i < \log k$, we have $R(i) \leq 2c \cdot \ln n$.

We now define the embedding. For every $0 \leq i < \log \log n$ and $1 \leq s \leq k$ sample $R(i)$ Frechet sets $\{A_h^{(s,i)}\}_{1 \leq h \leq R(i)}$, so that every element in X' is included in $A_h^{(s,i)}$ independently with probability $\min\{\frac{2^{2^i(1+s/k)-2+2s/k} \cdot (i+2)^2}{N}, 1\}$. To compensate for the increased number of points N , we also introduce additional $c \cdot \ln n$ Frechet sets $\{E_g\}$ for every $1 \leq g \leq c \cdot \ln n$,

⁹Ultimately, our dependence on n is just logarithmic, i.e., non-integrality of $\log \log n$ may increase the dimension by a multiplicative constant factor.

where each set E_g includes every element of X' independently with probability $\frac{1}{N}$. The embedding $f : X \rightarrow \mathbb{R}^{O(k \cdot n^{1/k} \cdot \ln n)}$ is defined for every $x \in X$ as

$$f(x) = \bigoplus_{s=1}^k \bigoplus_{i=0}^{\log \log n - 1} \bigoplus_{h=1}^{R(i)} d(x, A_h^{(s,i)}) \oplus \bigoplus_{g=1}^{c \cdot \ln n} d(x, E_g).$$

Distortion bound. Note that f is non-expansive, and we bound its contraction. Fix a pair $x_j, y \in X$ for $2 < j \leq n$, and let $0 \leq i < \log \log n$ be such that $x_j \in S_i$. Define $\alpha = 2ki + 1$, $q = (\alpha + 1)/2 = ki + 1$, and $r = \frac{d(x_j, y)}{\alpha}$. Define $B_0 = B(x_j, 0)$, $B_1 = B^o(y, r)$, $B_2 = B(x_j, 2r)$, and so on until B_q , as before (like in Section 3.1). Note that since $x_j \in S_i$, we have that $|B_0| = C(i) = \frac{n(\log \log n + 1)^2}{2^{2^{i+1}} \cdot (i+2)^2} \geq \frac{N}{2^{2^{i+1}} \cdot (i+2)^2}$.

$$(3.3) \quad |B_{b-1}| \geq \frac{N \cdot 2^{(2^i+2) \cdot (b-1)/k}}{2^{2^{i+1}} \cdot (i+2)^2} \quad \text{and} \quad |B_b| \leq \frac{N \cdot 2^{(2^i+2) \cdot b/k}}{2^{2^{i+1}} \cdot (i+2)^2},$$

then we will stop here. Otherwise, we have

$$|B_k| > \frac{N \cdot 2^{(2^i+2) \cdot k/k}}{2^{2^{i+1}} \cdot (i+2)^2} \geq \frac{N}{2^{2^i} \cdot (i+1)^2}$$

(using that $\frac{4}{(i+2)^2} \geq \frac{1}{(i+1)^2}$ holds for all $i \geq 0$). We then continue to ask whether there is a $1 \leq b \leq k$ so that

$$\begin{aligned} |B_{k+b-1}| &\geq \frac{N \cdot 2^{(2^{i-1}+2) \cdot (b-1)/k}}{2^{2^i} \cdot (i+1)^2} \\ \text{and} \quad |B_{k+b}| &\leq \frac{N \cdot 2^{(2^{i-1}+2) \cdot b/k}}{2^{2^i} \cdot (i+1)^2}, \end{aligned}$$

if there is no such b , we get that

$$|B_{2k}| > \frac{N \cdot 2^{(2^{i-1}+2) \cdot k/k}}{2^{2^i} \cdot (i+1)^2} \geq \frac{N}{2^{2^{i-1}} \cdot i^2}.$$

In general, after ℓ such phases we have that $|B_{\ell k}| > \frac{N}{2^{2^{i-\ell+1}} \cdot (i-\ell+2)^2}$. If all these i phases fail (to satisfy a condition similar to (3.3)), then $|B_{ik}| \geq N/16$. (Recall that $q > ki$, so all these balls do exist.)

Let us first see what happens if there is an $0 \leq \ell < i$ and $1 \leq b \leq k$ such that

$$\begin{aligned} |B_{\ell k+b-1}| &\geq \frac{N \cdot 2^{(2^{i-\ell}+2) \cdot (b-1)/k}}{2^{2^{i-\ell+1}} \cdot (i-\ell+2)^2} \\ &= \frac{N}{2^{2^{i-\ell}(2-(b-1)/k)-2(b-1)/k} \cdot (i-\ell+2)^2} \\ \text{and} \\ |B_{\ell k+b}| &\leq \frac{N \cdot 2^{(2^{i-\ell}+2) \cdot b/k}}{2^{2^{i-\ell+1}} \cdot (i-\ell+2)^2} \\ &= \frac{N}{2^{2^{i-\ell}(2-b/k)-2b/k} \cdot (i-\ell+2)^2}. \end{aligned}$$

For each $1 \leq h \leq R(i - \ell)$, let \mathcal{E}_h^{hit} be the event that $B_{\ell k + b - 1} \cap A_h^{(k-b, i-\ell)} \neq \emptyset$, and let \mathcal{E}_h^{miss} be the event that $B_{\ell k + b} \cap A_h^{(k-b, i-\ell)} = \emptyset$. Recall that $A_h^{(k-b, i-\ell)}$ contains each element of X' independently with probability

$$\begin{aligned} & \frac{2^{2^i - \ell(1+(k-b)/k) - 2 + 2(k-b)/k \cdot (i-\ell+2)^2}}{N} \\ &= \frac{2^{2^i - \ell(2-b/k) - 2b/k \cdot (i-\ell+2)^2}}{N}. \end{aligned}$$

It follows that

$$\begin{aligned} & \Pr[\mathcal{E}_h^{hit}] \\ &= 1 - \left(1 - \frac{2^{2^i - \ell(2-b/k) - 2b/k \cdot (i-\ell+2)^2}}{N}\right)^{|B_{\ell k + b - 1}|} \\ &\geq 1 - e^{-2^{-(2^i - \ell + 2)/k}} \geq 2^{-(2^i - \ell + 2)/k - 1}. \\ & \Pr[\mathcal{E}_h^{miss}] \\ &= \left(1 - \frac{2^{2^i - \ell(2-b/k) - 2b/k \cdot (i-\ell+2)^2}}{N}\right)^{|B_{\ell k + b}|} \\ &\geq 1/4. \end{aligned}$$

The two events are independent because the balls $B_{\ell k + b - 1}$ and $B_{\ell k + b}$ are disjoint. Hence $\Pr(\mathcal{E}_h^{hit} \wedge \mathcal{E}_h^{miss}) \geq 2^{-(2^i - \ell + 2)/k - 3}$.

Recall that $R(i - \ell) = c \cdot 2^{(2^i - \ell + 2)/k} \cdot \ln n$, so the probability that none of the $1 \leq h \leq R(i - \ell)$ have that both events occur is at most

$$(1 - 2^{-(2^i - \ell + 2)/k - 3})^{c \cdot 2^{(2^i - \ell + 2)/k} \cdot \ln n} \leq 1/n^2,$$

whenever $c \geq 16$. Note that if there is an h for which both events occur, it implies that

$$|d(x_j, A_h^{(k-b, i-\ell)}) - d(y, A_h^{(k-b, i-\ell)})| \geq r,$$

and thus the distortion for the pair x_j, y is at most α .

We now consider the case that no such ℓ and b were found, so we have the guarantee that $|B_{ki}| \geq N/16$. Recall that $B_q = B_{ki+1}$, and $|B_q| \leq N$. For each $1 \leq g \leq c \ln n$, let \mathcal{F}_g^{hit} be the event that $B_{ki} \cap E_g \neq \emptyset$, and let \mathcal{F}_g^{miss} be the event that $B_{ki+1} \cap E_g = \emptyset$. Since the balls are disjoint these events are independent, and

$$\begin{aligned} \Pr[\mathcal{F}_g^{hit}] &= 1 - \left(1 - \frac{1}{N}\right)^{|B_{ki}|} \geq 1 - e^{-N/(16N)} \geq \frac{1}{32}. \\ \Pr[\mathcal{F}_g^{miss}] &= \left(1 - \frac{1}{N}\right)^{|B_{ki+1}|} \geq 1/4. \end{aligned}$$

Thus, the probability that there is no $1 \leq g \leq c \ln n$ for which both events occur, is bounded by

$$\left(1 - \frac{1}{128}\right)^{c \ln n} \leq 1/n^2,$$

whenever c is large enough. If there is an $1 \leq g \leq c \ln n$ such that both events occur, then the distortion of the pair x_j, y is bounded by α in this case as well. By the union bound, there is a constant probability (which can easily be made polynomially close to 1 by increasing c), that all pairs have such bounded distortion. As $x_j \in S_i$, we have that $i \leq \log \log j$, so the distortion for pairs containing x_j is at most $\alpha = 2k \log \log j + 1$.

Bounding the dimension. We now turn to proving a bound on the number of nonzero coordinates of x_j . Let $0 \leq i < \log \log n$ be such that $x_j \in S_i$. Recall that x_j has $C(i) = \frac{n(\log \log n + 1)^2}{2^{2^{i+1}} \cdot (i+2)^2} \geq \frac{N}{2^{2^{i+1}} \cdot (i+2)^2}$ copies, and each element is chosen to be in $A_h^{(s,t)}$, $1 \leq s \leq k$, $0 \leq t \leq \log \log n$, independently with probability $\frac{2^{2^t(1+s/k) - 2 + 2s/k \cdot (t+2)^2}}{N}$, where $1 \leq h \leq R(t)$ and $R(t) = O(2^{(2^t+2)/k} \cdot \ln n)$.

For the first $k(i+1)$ collections of sets $\{A_h^{(s,t)}\}_h$ with $0 \leq t \leq i$ and $1 \leq s \leq k$, we will not try to prove that they contain a copy of x_j . The total number of such sets is

$$\begin{aligned} k \cdot \sum_{t=0}^i R(t) &= O(k \cdot \ln n \cdot (2^{(2^i+2)/k} + \log k)) \\ (3.4) \quad &= O(k \cdot \ln n \cdot (j^{1/k} + \log k)). \end{aligned}$$

(The additive term of $\log k$ comes from the first $\log k$ terms in the summation. The rest is dominated by the last term.) We also have an additional $c \cdot \ln n$ coordinates corresponding to the sets $\{E_g\}$. So the total number of non-zero coordinates that we have accounted for so far is still as in (3.4).

Consider the set $A_h^{(s,t)}$ for some $t \geq i + 1$, and let $Y_h^{(s,t)}$ be an indicator random variable for the event that $A_h^{(s,t)}$ does not contain any copy of x_j . Observe that x_j will have nonzero value in the coordinate of $A_h^{(s,t)}$ iff $Y_h^{(s,t)} = 1$. Denote $Y^{(s,t)} = \sum_{h=1}^{R(t)} Y_h^{(s,t)}$, and $Y = \sum_{t \geq i+1} \sum_{s=1}^k Y^{(s,t)}$. The number of nonzero coordinates of x_j from the sets $A_h^{(s,t)}$, for $t \geq i + 1$, is equal to Y .

We will soon show that $\mathbb{E}[Y] \leq c' \cdot k \cdot \ln n$ for some constant c' . Thus, by Chernoff bound (note the random variables $Y_h^{(s,t)}$ are independent)

$$\Pr[Y > 5c'k \cdot \ln n] \leq e^{-2k \ln n} \leq 1/n^2.$$

We will conclude that with constant probability, every x_j has at most $O(k \cdot (j^{1/k} + \log k) \cdot \ln n)$ nonzero coordinates.

For $t \geq i + 1$ we have that

$$\begin{aligned} \mathbb{E}[Y_h^{(s,t)}] &\leq \\ &\left(1 - \frac{2^{2^t(1+s/k)-2+2s/k} \cdot (t+2)^2}{N}\right)^{\frac{N}{2^{2^{t+1} \cdot (i+2)^2}}} \\ &\leq e^{-2^{2^t/k+2^t-2^{i+1}-2}}. \end{aligned}$$

Then

$$\begin{aligned} \mathbb{E}[Y^{(s,t)}] &= \sum_{h=1}^{R(t)} \mathbb{E}[Y_h^{(s,t)}] \\ &\leq O(2^{(2^t+2)/k} \cdot \ln n) \cdot e^{-2^{2^t/k+2^t-2^{i+1}-2}}, \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}[Y] &= \sum_{t \geq i+1} \sum_{s=1}^k \mathbb{E}[Y^{(s,t)}] \\ &\leq O(k \cdot \ln n) \cdot \sum_{t \geq i+1} 2^{(2^t+2)/k} \cdot e^{-2^{2^t/k+2^t-2^{i+1}-2}} \\ &= O(k \cdot \ln n), \end{aligned}$$

as the sum converges to a small constant.

We proved the following.

THEOREM 3.2. *For any parameter $k \geq 1$, any n -point metric embeds into ℓ_∞ with prioritized distortion $2k \log \log j + 1$ and prioritized dimension $O(k \cdot (j^{1/k} + \log k) \cdot \ln n)$.*

Prioritized Cohesive Dimension. We have shown that whp the number of nonzero coordinates is small, but did not guarantee which coordinates these will be. If one desires a cohesive prioritized dimension (recall that this means that only the first $\beta(j)$ coordinates are allowed to be nonzero), then this can be obtained as well, at the small cost of increasing the prioritized dimension to $O(k \cdot \ln n \cdot (\max\{j, \log n\})^{2/k})$. To achieve this, recall that $x_j \in S_j$. Set $l = \max\{i + 1, \log \log \log n + 1\}$, and we will show that whp only the first $k(l + 1)$ collections of Frechet sets $\{A_h^{(s,t)}\}_h$ with $0 \leq t \leq l$ and $1 \leq s \leq k$ may have nonzero values for x_j in the corresponding coordinates. Indeed, for $t > l$ it holds that

$$\begin{aligned} \Pr[Y_h^{(s,t)} = 1] &\leq \\ &\left(1 - \frac{2^{2^t(1+s/k)-2+2s/k} \cdot (t+2)^2}{N}\right)^{\frac{N}{2^{2^{t+1} \cdot (i+2)^2}}} \\ &\leq e^{-2^{2^t/k+2^t-2^{i+1}}} \\ &\leq e^{-2^{2^t-1}} \leq 1/n^4. \end{aligned}$$

The third inequality uses that $t \geq i + 2$, and the fourth that $t > \log \log \log n + 1$. By the union bound over all $1 \leq h \leq R(t)$ and $1 \leq j \leq n$, we get that each x_j has at most

$$k \cdot \sum_{t=0}^l R(t) \leq 2k \cdot R(l) \leq O(k \cdot \ln n \cdot (\max\{j, \log n\})^{2/k})$$

nonzero coordinates corresponding to the first $k \cdot \sum_{t=0}^l R(t)$ sets.

THEOREM 3.3. *For any parameter $k \geq 1$, any n -point metric embeds into ℓ_∞ with prioritized distortion $2k \log \log j + 1$ and prioritized cohesive dimension $O(k \cdot \max\{(j^{2/k} + \log k), \log^{2/k} n\} \cdot \ln n)$.*

4 Prioritized Embedding into a Single Tree

Define Φ to be the family of functions $\alpha : \mathbb{N} \rightarrow \mathbb{R}_+$ that satisfy the following properties:

- α is non-decreasing.
- $\sum_{i=1}^{\infty} 1/\alpha(i) < 1$.

For instance, one can take $\alpha(j) = c \cdot j \cdot \log j \cdot (\log \log j)^{1.1}$ for a suitable constant c .

Recall the result of [EFN18], who proved that every metric space embeds into a single tree with prioritized distortion $O(\alpha(j))$, for any function $\alpha \in \Phi$. It was also shown in [EFN18] that for every $\alpha \notin \Phi$, there exists a metric that does not admit a prioritized embedding into a single tree with distortion less than $\alpha/8$.

Here we strengthen the result of [EFN18], by showing a prioritized embedding of any metric into an ultrametric (which is a special (and useful) kind of tree metric; its definition appears below in the sequel), and of a graph into one of its spanning trees, with prioritized distortion $O(\alpha(j))$ for any $\alpha \in \Phi$.

4.1 Single Ultrametric An ultrametric (U, d) is a metric space satisfying a strong form of the triangle inequality, that is, for all $x, y, z \in U$, $d(x, z) \leq \max\{d(x, y), d(y, z)\}$. The following definition is known to be an equivalent one (see [BLMN05]).

DEFINITION 1. *An ultrametric U is a metric space (U, d) whose elements are the leaves of a rooted labeled tree T . Each $z \in T$ is associated with a label $\ell(z) \geq 0$ such that if $q \in T$ is a descendant of z then $\ell(q) \leq \ell(z)$ and $\ell(q) = 0$ iff q is a leaf. The distance between leaves $z, q \in U$ is defined as $d_T(z, q) = \ell(\text{lca}(z, q))$ where $\text{lca}(z, q)$ is the least common ancestor of z and q in T .*

THEOREM 4.1. *For any finite metric space (X, d) and any $\alpha \in \Phi$, there is a (non-contractive) embedding of X into an ultrametric with priority distortion $2\alpha(j)$.*

Proof. We create the ultrametric tree T for X by partitioning X into X_1 and X_2 , creating a root for T labeled by $\Delta = \text{diam}(X)$, and placing as its two children the trees T_1 and T_2 created recursively for X_1 and X_2 respectively. Clearly this embedding is non-contractive (as the distance between two points in T is the diameter of the cluster in which they are separated). It remains to see how to generate a partition of X that will have the required prioritized expansion. Note that a pair separated by the initial partition will have their distance exactly Δ in the ultrametric. This means that for each $j \in [n]$, we cannot separate x_j from x_i whenever

$$(4.5) \quad 2 \min\{\alpha(i), \alpha(j)\} \cdot d(x_j, x_i) < \Delta .$$

Call a pair x_j, x_i *bad* with respect to X_1 if $|\{x_j, x_i\} \cap X_1| = 1$ and (4.5) holds. We are now ready to define the partition of X : Pick any pair $u, v \in X$ so that $d(u, v) = \Delta$. Initialize $r = 0$, and repeat:

1. Set $X_1 \leftarrow B(u, r)$.
2. Take the minimal j such that there exists a bad pair (with respect to X_1) containing x_j . If there is no bad pair, stop.
3. Set $r \leftarrow r + \Delta/(2\alpha(j))$. Return to 1.

We claim that at the end of the process $v \notin X_1$. To see this, we first argue that for each $1 \leq j \leq n$, there can be at most two iterations in which r increases by $\Delta/(2\alpha(j))$. This is because after the first such iteration, we claim that $x_j \in X_1$. This is because either $x_j \in X_1$ even before the radius increase, or otherwise, it must be that $x_i \in X_1$, and the radius increases by $\Delta/(2\alpha(j)) > d(x_i, x_j)$, so by the triangle inequality x_j will be in X_1 . After the second increase we have that $B(x_j, \Delta/(2\alpha(j))) \subseteq X_1$ by the minimality of j , so there are no more bad pairs x_j, x_i with $j < i$ (since α is monotone). Using that $\alpha \in \Phi$, we obtain that the total increase in the radius is at most

$$(4.6) \quad \sum_{j=1}^n \frac{\Delta}{\alpha(j)} < \Delta .$$

We conclude that the partition of X to X_1 and $X_2 = X \setminus X_1$ is non-trivial ($u \in X_1$ and $v \in X_2$). Furthermore, for every pair $x, y \in X$ separated by the partition, $x \in X_1$ and $y \in X_2$ must abide the required distortion, since this pair is not bad, thus (4.5) does not hold.

The proof that the recursive construction satisfies the distortion constraints is by induction on $n = |X|$. Indeed both $|X_1|, |X_2| < |X|$, so there are trees T_1 and T_2 that preserve all the distortion constraints among the pairs in $\binom{X_1}{2} \cup \binom{X_2}{2}$ (note that α is monotone, so using

the induced priority ranking on X_1 and X_2 will only yield improved distortion bounds). Finally, all pairs in $X_1 \times X_2$ suffer appropriate distortion by the discussion above. \square

4.2 Single Spanning Tree In this section we extend the partition technique used to prove [Theorem 4.1](#), and obtain an embedding of any graph into a single spanning tree with priority distortion.

THEOREM 4.2. *For any weighted graph $G = (V, E)$ and any $\alpha \in \Phi$, there is an embedding of G into a spanning tree with priority distortion $O(\alpha(j))$.*

Petal Decomposition We will make use of a construction of a spanning tree by a *Petal-Decomposition*, introduced in [\[AN19\]](#). This decomposition generates a hierarchical partition of the vertex set of a graph, which can be viewed as a laminar family over subsets of V . Every cluster X in the family is also associated with a center $x \in X$ and a target $t \in X$. The cluster X is partitioned to X_0, X_1, \dots, X_s (for some integer s), and the algorithm also specifies s edges that connect these clusters in a tree manner. Each X_i for $1 \leq i \leq s$ is called a *petal*, and is generated in the graph induced on $Y_{i-1} = X \setminus (\cup_{j=1}^{i-1} X_j)$, by picking an arbitrary target t_i (except perhaps for t_1 , which is picked on the shortest path from x to t), of distance at least $3\text{rad}(X)/4$ from x_0 ,¹⁰ picking an arbitrary radius $0 \leq r_i \leq \text{rad}(X)/8$, and setting $X_i = P(t_i, r_i)$. (The definition of $P(t_i, r_i)$ appears below.) The final cluster $X_0 = Y_s$, is simply the remaining vertices when there are no vertices sufficiently far from the center x .

We now explain how a petal $P(t, r)$ is defined. Given a weighted undirected graph $G = (V, E, w)$ with a center x and a desired target t for the new petal, let P_{tx} be the shortest path from t to x . Let $\tilde{G} = (V, A, \bar{w})$ be the weighted directed graph created by adding the two directed edges $(u, v), (v, u) \in A$ for each $\{u, v\} \in E$, and setting $\bar{w}(u, v) = d_G(u, v) - (d_G(v, x) - d_G(u, x))$. However, the weight of each directed edge (u, v) on the path P_{tx} , where u is closer to t , is set to be $\bar{w}(u, v) = w(u, v)/2$ (note that by our definition it should have been $\bar{w}(u, v) = 2w(u, v)$). The petal $P(t, r)$ is the ball around t of radius $r/2$ in \tilde{G} (the center of the new petal is defined as the farthest point from t on P_{tx} which lies in $P(t, r)$). The crucial properties of the Petal-Decomposition are (see [\[AN19\]](#) for a proof):

1. For each cluster X , the tree $T[X]$ spanning X has radius at most $4 \cdot \text{rad}(X)$.

¹⁰The radius of a cluster X with center x is defined as $\text{rad}(X) = \max_{y \in X} d(x, y)$. Observe that $\text{diam}(X)/2 \leq \text{rad}(X) \leq \text{diam}(X)$.

2. For any $y \in V$ within distance δ from (some vertex in) the petal $P(t, r)$, it holds that $y \in P(t, r + 4\delta)$.

The first property mean that like in the ultrametric case, a separated pair in a cluster of radius Δ will be at distance $O(\Delta)$ in the final tree, and the second property suggests that petals are similar to balls.

Proof. [Proof of Theorem 4.2] We generate a tree T by applying Petal-Decomposition on the graph. For each cluster X generated in the process, that is partitioned into X_0, X_1, \dots, X_s , we must provide for every $1 \leq i \leq s$ a radius from the range $[0, \text{rad}(X)/8]$. Note that if a pair $x, y \in X$ is separated when creating X_i (that is $x \in X_i$ and $y \in Y_i$), then by the first property their final distance in the tree will be at most $\text{diam}(T[X]) \leq 8 \cdot \text{rad}(X)$. It follows that to satisfy the distortion constraints we can use essentially the same definition of bad pairs and algorithm for choosing r as in Section 4.1. The only difference is that we will use larger constants: a pair $x_j, x_{j'}$ is called *bad* if they are separated when partitioning X and $128 \min\{\alpha(j), \alpha(j')\} \cdot d(x_j, x_{j'}) < \text{rad}(X)$. The process to find the radius r_i for carving the next petal X_i with a given target t_i out of the graph induced on Y_{i-1} is the following. Initialize $r_i = 0$, and repeat:

1. Set $X_i \leftarrow P(t_i, r_i)$.
2. Take the minimal j such that there exists a bad pair containing x_j . If there are no bad pairs, stop.
3. Set $r_i \leftarrow r_i + \text{rad}(X)/(16\alpha(j))$. Return to 1.

Using the second property of petals, we obtain that j can appear as the minimal index at most twice. This is because after the first time, as the radius increased by $\text{diam}(X)/(16\alpha(j))$, so vertices within distance $\text{rad}(X)/(64\alpha(j)) \geq d(x_{j'}, x_j)$ will be contained in X_i (where $j' > j$ is index of the other vertex in the bad pair). In particular, it follows that after the first time it must be that $x_j \in X_i$. After the second time j is the minimal we get that $B(x_j, \text{rad}(X)/(64\alpha(j))) \subseteq X_i$. Thus (by a similar argument to the section above) x_j will never be the minimal-index vertex of a bad pair again (until r_i is set). By the calculation done as in (4.6), the total increase in r_i is at most $2\text{rad}(X)/16 = \text{rad}(X)/8$, as required.

Finally, we bound the distortion of any pair $x_j, x_{j'}$ which is not bad. If this pair is separated while partitioning X , their distance in T will be at most $8\text{rad}(X)$, so by definition of a bad pair, the distortion they suffer is at most $8 \cdot 128\alpha(j) = O(\alpha(j))$.

□

References

- [AB17] Amir Abboud and Greg Bodwin. The 4/3 additive spanner exponent is tight. *J. ACM*, 64(4):28:1–28:20, 2017.
- [ABC⁺05] Ittai Abraham, Yair Bartal, Hubert T.-H. Chan, Kedar Dhamdhere, Anupam Gupta, Jon M. Kleinberg, Ofer Neiman, and Aleksandrs Slivkins. Metric embeddings with relaxed guarantees. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, pages 83–100. IEEE Computer Society, 2005.
- [ABN08] Ittai Abraham, Yair Bartal, and Ofer Neiman. Embedding metric spaces in their intrinsic dimension. In *Proceedings of the 19th annual ACM-SIAM symposium on Discrete algorithms*, SODA '08, pages 363–372, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [ABN11] Ittai Abraham, Yair Bartal, and Ofer Neiman. Advances in metric embedding theory. *Advances in Mathematics*, 228(6):3026 – 3126, 2011.
- [AFGN18] Ittai Abraham, Arnold Filtser, Anupam Gupta, and Ofer Neiman. Metric embedding via shortest path decompositions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 952–963, 2018.
- [AKPW95] Noga Alon, Richard M. Karp, David Peleg, and Douglas West. A graph-theoretic game and its application to the k-server problem. *SIAM Journal on Computing*, 24(1):78–100, February 1995.
- [AN19] Ittai Abraham and Ofer Neiman. Using petal-decompositions to build a low stretch spanning tree. *SIAM J. Comput.*, 48(2):227–248, 2019.
- [Ass83] P. Assouad. Plongements lipschitziens dans \mathbb{R}^n . *Bull. Soc. Math. France*, 111(4):429–448, 1983.
- [Bar96] Yair Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, FOCS '96, pages 184–193, Washington, DC, USA, 1996. IEEE Computer Society.
- [Bar98] Yair Bartal. On approximating arbitrary metrics by tree metrics. In *Proceedings of the 30th annual ACM symposium on Theory of computing*, STOC '98, pages 161–168, New York, NY, USA, 1998. ACM.
- [BFN16] Yair Bartal, Arnold Filtser, and Ofer Neiman. On notions of distortion and an almost minimum spanning tree with constant average distortion. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 873–882, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics.
- [BLMN05] Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. On metric ramsey-type phenomena. *Annals of Mathematics*, 162(2):643–709, 2005.
- [Bou85] Jean Bourgain. On Lipschitz embedding of finite

- metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985.
- [CE06] Don Coppersmith and Michael Elkin. Sparse source-wise and pairwise distance preservers. *SIAM J. Discrete Math.*, 20(2):463–501, 2006.
- [EEST08] Michael Elkin, Yuval Emek, Daniel A. Spielman, and Shang-Hua Teng. Lower-stretch spanning trees. *SIAM Journal on Computing*, 38(2):608–628, 2008.
- [EFN17] Michael Elkin, Arnold Filtser, and Ofer Neiman. Terminal embeddings. *Theor. Comput. Sci.*, 697:1–36, 2017.
- [EFN18] Michael Elkin, Arnold Filtser, and Ofer Neiman. Prioritized metric structures and embedding. *SIAM J. Comput.*, 47(3):829–858, 2018.
- [EN18] Michael Elkin and Ofer Neiman. Near isometric terminal embeddings for doubling metrics. In *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, pages 36:1–36:15, 2018.
- [EN19] Michael Elkin and Ofer Neiman. Efficient algorithms for constructing very sparse spanners and emulators. *ACM Trans. Algorithms*, 15(1):4:1–4:29, 2019.
- [EP04] Michael Elkin and David Peleg. (1+epsilon, beta)-spanner constructions for general graphs. *SIAM J. Comput.*, 33(3):608–631, 2004.
- [FGK20] Arnold Filtser, Lee-Ad Gottlieb, and Robert Krauthgamer. Labelings vs. embeddings: On distributed representations of distances. To appear in SODA, 2020.
- [FRT04] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 69(3):485–497, 2004.
- [GKL03] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science, FOCS '03*, pages 534–543, Washington, DC, USA, 2003. IEEE Computer Society.
- [JL84] William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*, pages 189–206. American Mathematical Society, Providence, RI, 1984.
- [KSW09] Jon Kleinberg, Aleksandrs Slivkins, and Tom Wexler. Triangulation and embedding using small sets of beacons. *Journal of the ACM*, 56(6):32:1–32:37, September 2009.
- [LLR95] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [LMN04] James R. Lee, Manor Mendel, and Assaf Naor. Metric structures in ℓ_1 : Dimension, snowflakes, and average distortion. In *LATIN*, pages 401–412, 2004.
- [Mat96] Jyri Matousek. On the distortion required for embedding finite metric spaces into normed spaces. *Israel J. Math.*, 93:333–344, 1996.
- [Mat02] Jiri Matoušek. *Lectures on discrete geometry*. Springer-Verlag, New York, 2002.
- [MMMR18] Sepideh Mahabadi, Konstantin Makarychev, Yury Makarychev, and Ilya P. Razenshteyn. Nonlinear dimension reduction via outer bi-lipschitz extensions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 1088–1101, 2018.
- [NN19] Shyam Narayanan and Jelani Nelson. Optimal terminal dimensionality reduction in euclidean space. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019.*, pages 1064–1069, 2019.
- [Pel99] David Peleg. Proximity-preserving labeling schemes and their applications. In *Graph-Theoretic Concepts in Computer Science, 25th International Workshop, WG'99, Ascona, Switzerland, June 17-19, 1999, Proceedings*, pages 30–41, 1999.
- [PS89] David Peleg and Alejandro A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.
- [RR98] Yuri Rabinovich and Ran Raz. Lower bounds on the distortion of embedding finite metric spaces in graphs. *Discrete & Computational Geometry*, 19(1):79–94, 1998.
- [Tho01] Mikkel Thorup. Compact oracles for reachability and approximate distances in planar digraphs. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 242–251, 2001.
- [TZ05] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *J. ACM*, 52(1):1–24, 2005.