

Nonlinear Dimension Reduction via Outer Bi-Lipschitz Extensions

Sepideh Mahabadi
Columbia University
New York, NY, United States
mahabadi@mit.edu

Yury Makarychev*
TTIC
Chicago, IL, United States
yury@ttic.edu

Konstantin Makarychev
Northwestern University
Chicago, IL, United States
konstantin@northwestern.edu

Ilya Razenshteyn
Microsoft Research Redmond
Redmond, WA, United States
ilyaraz@mit.edu

ABSTRACT

We introduce and study the notion of an *outer bi-Lipschitz extension* of a map between Euclidean spaces. The notion is a natural analogue of the notion of a *Lipschitz extension* of a Lipschitz map. We show that for every map f there exists an outer bi-Lipschitz extension f' whose distortion is greater than that of f by at most a constant factor. This result can be seen as a counterpart of the classic Kirszbraun theorem for outer bi-Lipschitz extensions. We also study outer bi-Lipschitz extensions of near-isometric maps and show upper and lower bounds for them. Then, we present applications of our results to prioritized and terminal dimension reduction problems, described next.

We prove a *prioritized* variant of the Johnson–Lindenstrauss lemma: given a set of points $X \subset \mathbb{R}^d$ of size N and a permutation (“priority ranking”) of X , there exists an embedding f of X into $\mathbb{R}^{O(\log N)}$ with distortion $O(\log \log N)$ such that the point of rank j has only $O(\log^{3+\varepsilon} j)$ non-zero coordinates – more specifically, all but the first $O(\log^{3+\varepsilon} j)$ coordinates are equal to 0; the distortion of f restricted to the first j points (according to the ranking) is at most $O(\log \log j)$. The result makes a progress towards answering an open question by Elkin, Filtser, and Neiman about prioritized dimension reductions.

We prove that given a set X of N points in $\mathbb{R}^d$, there exists a *terminal* dimension reduction embedding of $\mathbb{R}^d$ into $\mathbb{R}^{d'}$, where $d' = O(\frac{\log N}{\varepsilon^4})$, which preserves distances $\|x - y\|$ between points $x \in X$ and $y \in \mathbb{R}^d$, up to a multiplicative factor of $1 \pm \varepsilon$. This improves a recent result by Elkin, Filtser, and Neiman.

The dimension reductions that we obtain are nonlinear, and this nonlinearity is necessary.

CCS CONCEPTS

• **Mathematics of computing** → **Dimensionality reduction**; • **Theory of computation** → **Random projections and metric embeddings**; **Computational geometry**; *Design and analysis of algorithms*;

KEYWORDS

Dimension Reduction, Bi-Lipschitz Extension, Prioritized Johnson–Lindenstrauss, Near-Isometric Maps, Metric Embedding

ACM Reference Format:

Sepideh Mahabadi, Konstantin Makarychev, Yury Makarychev, and Ilya Razenshteyn. 2018. Nonlinear Dimension Reduction via Outer Bi-Lipschitz Extensions. In *Proceedings of 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC’18)*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3188745.3188828>

1 INTRODUCTION

In this paper, we introduce and study the notion of an *outer bi-Lipschitz extension*. The notion is a natural analogue of the notion of a Lipschitz extension, which is widely used in mathematics and theoretical computer science. Recall that a map $f : X \rightarrow Y$ is C -Lipschitz if for any two points $x, y \in X$ we have $d_Y(f(x), f(y)) \leq C \cdot d_X(x, y)$; the Lipschitz constant of f is the minimum C such that f is C -Lipschitz. In the Lipschitz extension problem, given a Lipschitz map f from a subset A of X to Y and a superset $A' \supset A$, the goal is to find an extension map f' from A' to Y such that the Lipschitz constant of f' is equal to or not significantly larger than the Lipschitz constant of f . This problem has found numerous applications in mathematics and theoretical computer science (see e.g., [11, 29, 31, 35, 37–39, 41–43]). One of the most important results in the field is the Kirszbraun theorem, which states that any map $f : A \rightarrow \mathbb{R}^m$ from a subset A of Euclidean space $\mathbb{R}^n$ to Euclidean space $\mathbb{R}^m$ can be extended to a map $f' : \mathbb{R}^n \rightarrow \mathbb{R}^m$ so that the Lipschitz constant of f' equals that of f [31] (see Theorem 1.13 in Section 1.2; see also [4]).

Outer bi-Lipschitz extension. In this paper, we prove several analogues of the Kirszbraun theorem for *bi-Lipschitz* maps. The bi-Lipschitz constant of a map $f : X \rightarrow Y$ is the minimum D such that for some $\lambda > 0$ and every $x, y \in X$, $\lambda \cdot d_X(x, y) \leq d_Y(f(x), f(y)) \leq \lambda \cdot D \cdot d_X(x, y)$. If there is no such number D , we say that the map

*Supported by NSF awards CAREER CCF-1150062, IIS-1302662, and CCF-1718820.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC’18, June 25–29, 2018, Los Angeles, CA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5559-9/18/06...\$15.00

<https://doi.org/10.1145/3188745.3188828>

is not bi-Lipschitz. Bi-Lipschitz maps are also known as embeddings with distortion D . Low distortion metric embeddings have numerous applications in approximation and online algorithms (see e.g. [1, 2, 13–16, 20, 23–25, 36, 40]); hardness of approximation (see e.g. [30]); computational geometry (see e.g. [27, 40] and references therein); and sketching, streaming, and similarity search algorithms (see e.g. [3, 9, 10, 12, 17, 19, 21, 28, 44–46]).

Since bi-Lipschitz maps are widely used in mathematics and theoretical computer science, it is natural to ask whether there is a counterpart of the Kirszbraun theorem for bi-Lipschitz maps.

Given a bi-Lipschitz map f from a subset of $\mathbb{R}^n$ to $\mathbb{R}^m$, can we extend it to a bi-Lipschitz map from the whole space $\mathbb{R}^n$ to $\mathbb{R}^m$?

This question has been extensively studied in the literature (see e.g. [5, 6, 26, 32, 47, 48]). It turns out that the answer to this question depends on the geometry of the set A . In general, the answer is “no”. For instance, consider a map that maps points 0, 1, 2 to 0, -1 , 2, respectively. There is no continuous one-to-one extension of this map to $\mathbb{R}$, let alone a bi-Lipschitz extension. The reason is that in one dimension we cannot connect points 0 and -1 and points -1 and 2 with non-intersecting paths. However, we can easily do this in $\mathbb{R}^2$. This observation suggests the following idea. Let $A \subset \mathbb{R}^n$ and $f : A \rightarrow \mathbb{R}^m$ be a bi-Lipschitz map. Let us allow extension f' of f to use additional dimensions or, in other words, allow f' to map points $x \in \mathbb{R}^n \setminus A$ to points in some higher-dimensional (ambient) space $\mathbb{R}^{m'}$ that contains $\mathbb{R}^m$. We get the following definition.

Definition 1.1 (Outer extension). A map $f' : A' \rightarrow \mathbb{R}^{m'}$ (where $m' \geq m$) is an outer extension of f if $f(a) = f'(a)$ for all $a \in A$; we assume that $\mathbb{R}^m$ is the subspace of $\mathbb{R}^{m'}$ spanned by the first m standard basis vectors; that is, we identify points $(x_1, \dots, x_m) \in \mathbb{R}^m$ and $(x_1, \dots, x_m, 0, \dots, 0) \in \mathbb{R}^{m'}$. We say that the extension is proper if $m = m'$.

Note that the exact dimension of the image is not very important in many applications in computer science, as long as the dimension is comparable to m and n . Therefore, outer extensions seem to be as useful as proper (standard) extensions. However, in stark contrast with proper bi-Lipschitz extensions, outer bi-Lipschitz extensions always exist – for every bi-Lipschitz map $f : A \rightarrow \mathbb{R}^m$ there exists an outer bi-Lipschitz extension $f' : \mathbb{R}^n \rightarrow \mathbb{R}^{m'}$, as we prove in this paper.

1.1 Results

Outer bi-Lipschitz Extensions. One of the main results of this paper is an analogue of the Kirszbraun theorem for bi-Lipschitz maps.

THEOREM 1.2. *Let $X \subset \mathbb{R}^n$ and $f : X \rightarrow \mathbb{R}^m$ be a bi-Lipschitz map with distortion at most D . There exists an outer extension $f' : \mathbb{R}^n \rightarrow \mathbb{R}^{m'}$ of f with the distortion at most $3D$ and $m' = n + m$.*

The main difference between the outer bi-Lipschitz extension from Theorem 1.2 and the Lipschitz extension from the Kirszbraun theorem – aside from the difference we discussed above (that Theorem 1.2 gives an outer extension and not a proper extension) – is that while the Lipschitz extension preserves the Lipschitz constant

of the map exactly, the bi-Lipschitz extension preserves the distortion only up to a constant factor. This limitation is unavoidable; it is easy to see that even in the example we considered – extending the map f that sends 0, 1, 2 to 0, -1 , 2, respectively – the distortion of any outer extension of f is greater than the distortion of f . Thus, for arbitrary bi-Lipschitz maps we cannot get a result stronger than Theorem 1.2 (except that factor 3 in the statement of the theorem can be potentially replaced with a smaller factor $c > 1$).

We then focus on an important class of near-isometric maps, maps with distortion $D = 1 + \varepsilon$. Observe that if the distortion of f is exactly 1 (i.e., f is an isometric embedding), it can be extended to an isometric embedding of the whole space $\mathbb{R}^n$ into $\mathbb{R}^{m'}$. In this case, we can extend f without increasing its distortion. What happens if the distortion of f is close to 1 but not 1? Let $\varphi(\varepsilon)$ be the smallest ε' such that the following holds: for every map $f : A \rightarrow \mathbb{R}^m$ with distortion at most $D = 1 + \varepsilon$, there exists an outer extension $f' : \mathbb{R}^n \rightarrow \mathbb{R}^{m'}$ with distortion at most $D' = 1 + \varepsilon'$. Note that $\varphi(0) = 0$, as discussed above.

OPEN PROBLEM 1. *Find the asymptotic behavior of $\varphi(\varepsilon)$ as $\varepsilon \rightarrow 0$. Does $\varphi(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$?*

We study this problem and get partial results for it. First, we show that $\varphi(\varepsilon) \geq \Omega(1/\log^2(1/\varepsilon))$.

THEOREM 1.3. *There exists a map $f : X \rightarrow \mathbb{R}$, where $X \subset \mathbb{R}$, with the distortion $1 + \varepsilon$, such that every outer extension $f' : \mathbb{R} \rightarrow \mathbb{R}^m$ of f has distortion at least $1 + \Omega(\frac{1}{\log^2(1/\varepsilon)})$.*

Note that $1/\log^2(1/\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$, but the dependence of $1/\log^2(1/\varepsilon)$ on ε is not polynomial and, in our opinion, highly unusual. This result rules out the possibility that $\varphi(\varepsilon) = O(\varepsilon^{1/k})$ for any k . Further, we provide some evidence that $\varphi(\varepsilon)$ might, in fact, be equal to $1 + \Theta(\frac{1}{\log^2(1/\varepsilon)})$. Namely, we prove the following result for 1-dimensional case: for every map from $X \subset \mathbb{R}$ to $\mathbb{R}$, there is an outer extension with $D' = 1 + O(\frac{1}{\log^2(1/\varepsilon)})$. By Theorem 1.3, this bound is asymptotically optimal.

THEOREM 1.4. *Let $X \subset \mathbb{R}$ and $f : X \rightarrow \mathbb{R}$ be a map with the distortion at most $1 + \varepsilon$. There exists an outer extension $f' : \mathbb{R} \rightarrow \mathbb{R}^2$ of f with the distortion at most $1 + O(\frac{1}{\log^2(1/\varepsilon)})$.*

We also consider a simpler problem of extending a near-isometric map by one point. We prove the following result.

THEOREM 1.5. *Let f be a $(1 + \varepsilon)$ -bi-Lipschitz map from a subset X of $\mathbb{R}^n$ to $\mathbb{R}^m$ and $u \in \mathbb{R}^n$. There exists an outer extension $f' : X \cup \{u\} \rightarrow \mathbb{R}^{m+1}$ of f with the distortion at most $1 + O(\sqrt{\varepsilon})$.*

The bound in this theorem is asymptotically tight – there exist a map f from a subset of $\mathbb{R}$ to $\mathbb{R}$ and a point $u \in \mathbb{R}$ such that every outer extension of f to u has distortion $1 + \Omega(\sqrt{\varepsilon})$.

Computability. Given sets $A \subset A' \subset \mathbb{R}^n$ and a map $f : A \rightarrow \mathbb{R}^m$, we can compute an outer extension $f' : A' \rightarrow \mathbb{R}^n$ with the least possible distortion using semidefinite programming (SDP). The running time is polynomial in $|A'|$ and $\log 1/\delta$, where δ is the desired precision. In particular, we can find outer extensions f' , whose existence is guaranteed by Theorems 1.2 and 1.5.

Applications. Using our extension results, we obtain *prioritized* and *terminal* dimension reductions [22, 23]. Recall the statement of the Johnson–Lindenstrauss lemma [29].

THEOREM 1.6 (THE JOHNSON–LINDENSTRAUSS LEMMA [29]). *For every $0 < \varepsilon < 1/2$ and every set $X \subset \mathbb{R}^d$ of size N , there exists an embedding $f : X \rightarrow \mathbb{R}^{d'}$, where $d' = O\left(\frac{\log N}{\varepsilon^2}\right)$, such that for every $p, q \in X$: $\|p - q\|_2 \leq \|f(p) - f(q)\|_2 \leq (1 + \varepsilon)\|p - q\|_2$.*

Prioritized metric structures and embeddings were introduced and studied by Elkin, Filtser, and Neiman [22]. Among several very interesting results obtained in [22], one is a construction of prioritized embeddings. We give a definition of a prioritized dimension reduction in the spirit of [22].

Definition 1.7 (Prioritized dimension reduction). Consider a set of points $X \subset \mathbb{R}^d$ of size N . Let π be a bijection from $[N] = \{1, \dots, N\}$ to X , which defines a *priority ranking* of X : $\pi(1), \dots, \pi(N)$. An embedding $f : X \rightarrow \mathbb{R}^{d'}$ is an (α, β) -prioritized dimension reduction, where $\alpha : [N] \rightarrow \mathbb{R}$ and $\beta : [N] \rightarrow \mathbb{N}$, if

- for every $j \in [N]$, the distortion of f restricted to points $\pi(1), \dots, \pi(j)$ is at most $\alpha(j)$.
- for every $j \in [N]$, $\pi(j)$ is mapped to a point $f(\pi(j))$ in $\mathbb{R}^{\beta(j)}$; that is, all but the first $\beta(j)$ coordinates of $f(\pi(j))$ are equal to 0.

Note that points $f(\pi(1)), \dots, f(\pi(j))$ lie in Euclidean space of dimension $\beta(j)$ and $\beta(j)$ may potentially be much smaller than $\log N$ (when $j \ll N$). The definition requires that the distortion of the distance between points $\pi(i)$ and $\pi(j)$ be at most $\alpha(\max(i, j))$ (note that this condition is weaker than a similar condition in the definition of a *prioritized embedding* in [22], which requires that the distortion be at most $\alpha(\min(i, j))$).

Ideally, we want to have a dimension reduction with parameters $(1 + \varepsilon, \text{polylog } j)$.

OPEN PROBLEM 2 ([22, TALK AND PERS. COMM.]). *Is there a prioritized dimension reduction with parameters $(1 + \varepsilon, \text{polylog } j)$?*

Very little is known about prioritized dimension reductions. The only known result follows from Theorem 15 in [22]. (The theorem is a prioritized variant of Bourgain's theorem [18] and is more general than its corollary stated below.)

THEOREM 1.8 ([22]). *For every set $X \subset \mathbb{R}^d$ and $\varepsilon > 0$, there is a $(c_1 \log^4 j, c_2 \log^{4+\varepsilon} j)$ -prioritized dimension reduction $f : X \rightarrow \mathbb{R}^{O(\log^2 |X|)}$ (where c_1, c_2 depend only on ε).*

We make further progress towards solving Open Problem 2.

THEOREM 1.9. *For every set $X \subset \mathbb{R}^d$, $\varepsilon > 0$, and $N = |X|$, there exist*

- a $(c_1 \log_2 \log_2 j, c_2 \log_2^{3+\varepsilon} j)$ -prioritized dimension reduction $f : X \rightarrow \mathbb{R}^{O(\log N)}$, where $c_1 = 3 + \varepsilon$ and $c_2 = O(1/\varepsilon^2)$,
- a $((3 + \varepsilon)^k, c_1 \log_2 j \log^{1/k} N)$ -prioritized dimension reduction $f : X \rightarrow \mathbb{R}^{O(\log N)}$ for every integer parameter $k > 1$, where $c_1 = O(1/\varepsilon^2)$.

The dimension reductions can be computed in polynomial time.

The first result gives a prioritized dimension reduction with a reasonably small distortion $O(\log \log j)$ and desired polylogarithmic dimension. The second result gives a constant distortion and maps the first j points to a subspace of dimension $O(\log_2 j \log^{1/k} N)$.

Now we switch to another problem introduced by Elkin, Filtser, and Neiman [23].

Definition 1.10 (Terminal dimension reduction). Suppose that we are given a set of points (which we call *terminals*) $X \subset \mathbb{R}^d$. We say that a map $f : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ is a *terminal dimension reduction* with distortion D if for every terminal $x \in X$ and point $p \in \mathbb{R}^d$ (p may be a terminal), we have

$$\|p - x\| \leq \|f(p) - f(x)\| \leq D \|p - x\|.$$

Elkin, Filtser, and Neiman [23] proved that there exists a terminal dimension reduction with distortion $O(1)$ and dimension $d' = O(\log |X|)$. We show how to obtain the distortion of $1 + \varepsilon$.

THEOREM 1.11. *For every set $X \subset \mathbb{R}^d$ of size N and parameter $0 < \varepsilon < 1/2$, there exists a terminal dimension reduction $f : X \rightarrow \mathbb{R}^{d'}$ with distortion $1 + \varepsilon$, where $d' = O\left(\frac{\log N}{\varepsilon^4}\right)$. The dimension reduction can be computed in polynomial time.*

It is an interesting question if the dimension $O\left(\frac{\log N}{\varepsilon^4}\right)$ can be lowered. Since f is also a (standard) dimension reduction for X , d' must be at least $\Omega\left(\frac{\log N}{\varepsilon^2}\right)$ as was shown by Larsen and Nelson [34] (see also [7, 8, 33]).

OPEN PROBLEM 3. *Is it possible to decrease the dimension to $O\left(\frac{\log N}{\varepsilon^2}\right)$ in Theorem 1.11.*

It is interesting that while most dimension reduction constructions described in the literature are given by linear transformations, prioritized and terminal dimension reductions must be non-linear (see the full version of the paper for details). In particular, all dimension reductions presented in this paper are non-linear.

In Section 2, we prove Theorem 1.2. In Section 3, we obtain an optimal bound on one-point outer bi-Lipschitz extensions (prove Theorem 1.5 and show its optimality). Then, in Section 4, we present applications of our results. Finally, in Section 5, we present the proof of Theorem 1.4, as well as a matching lower bound.

1.2 Preliminaries

In this paper, $\mathbb{R}^n$ denotes n -dimensional Euclidean space, equipped with the standard Euclidean norm $\|\cdot\|$. For $m < m'$, we identify $\mathbb{R}^m$ with the m -dimensional subspace of $\mathbb{R}^{m'}$ spanned by the first m standard basis vectors (in other words, we identify vectors $(x_1, \dots, x_m) \in \mathbb{R}^m$ and $(x_1, \dots, x_m, 0, \dots, 0) \in \mathbb{R}^{m'}$).

Definition 1.12 (Lipschitz constant and distortion). Let (X, d_X) and (Y, d_Y) be metric spaces, and let $f : X \rightarrow Y$ be a map. Define the Lipschitz constant of f as $\|f\|_{\text{Lip}} = \sup_{x, y \in X} \frac{d_Y(f(x), f(y))}{d_X(x, y)}$. We say that the map f is Lipschitz if $\|f\|_{\text{Lip}} < \infty$. A map f is non-expanding if $\|f\|_{\text{Lip}} \leq 1$. The distortion or bi-Lipschitz constant of an injective map f is $D = D(f) = \|f\|_{\text{Lip}} \cdot \|f^{-1}\|_{\text{Lip}}$. If a map is not injective, its distortion is infinite. A map f is bi-Lipschitz if $D(f) < \infty$.

THEOREM 1.13 (KIRSZBRAUN EXTENSION THEOREM). *Consider Euclidean spaces $\mathbb{R}^n$ and $\mathbb{R}^m$, and an arbitrary non-empty subset X of $\mathbb{R}^n$. Let $f : X \rightarrow \mathbb{R}^m$ be a Lipschitz map. There exists a proper extension $f' : \mathbb{R}^n \rightarrow \mathbb{R}^m$ of f with the same Lipschitz constant as f : $\|f'\|_{\text{Lip}} = \|f\|_{\text{Lip}}$.*

We note that Makarychev and Makarychev [38] previously introduced a notion of external bi-Lipschitz extension, but that notion is significantly different and less natural from the notion of the outer bi-Lipschitz extension studied in this paper.

2 OUTER BI-LIPSCHITZ EXTENSION

In this section, we prove Theorem 1.2 that states that any bi-Lipschitz map f from a subset X of $\mathbb{R}^n$ to $\mathbb{R}^m$ can be extended to a bi-Lipschitz map $f' : \mathbb{R}^n \rightarrow \mathbb{R}^m$ for some $m' > m$. The result can be seen as a counterpart of the Kirszbraun theorem.

Informal overview of the proof idea. For simplicity, let us assume for now that f is near-isometric (it approximately preserves distances). We want to construct a map $f' : \mathbb{R}^n \rightarrow \mathbb{R}^{m'}$ that satisfies the following conditions:

- (1) f' is an outer extension of f ; that is, $f'(x) = f(x)$ for every $x \in X$;
- (2) $\|f'(x) - f'(y)\| \leq O(\|x - y\|)$ for all $x, y \in \mathbb{R}^n$;
- (3) $\|f'(x) - f'(y)\| \geq \Omega(\|x - y\|)$ for all $x, y \in \mathbb{R}^n$.

First, using the Kirszbraun theorem, we find a Lipschitz extension $\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$. If we were to let $f' = \tilde{f}$, then f' would satisfy conditions (1) and (2) but not necessarily (3); namely, for some points $x, y \in \mathbb{R}^n$, the distance between $f'(x)$ and $f'(y)$ would potentially be considerably smaller than that between x and y ; in fact, it could happen that $\tilde{f}(x) = \tilde{f}(y)$ for some $x \neq y$. Instead, we are going to let $f'(x) = \tilde{f}(x) \oplus h(x) \in \mathbb{R}^{n+m}$ for some map h from $\mathbb{R}^n$ to $\mathbb{R}^n$. We will choose h which satisfies the following conditions:

- (1') For $x \in X$, $h(x) = 0$. This condition is necessary to ensure that f' is an outer extension of f .
- (2') For all $x, y \in \mathbb{R}^n$, $\|h(x) - h(y)\| \leq O(\|x - y\|)$ and thus $\|f'(x) - f'(y)\| \leq \|\tilde{f}(x) - \tilde{f}(y)\| + \|h(x) - h(y)\| \leq O(\|x - y\|)$.
- (3') If $\|\tilde{f}(x) - \tilde{f}(y)\| \ll \|x - y\|$ for some $x, y \in \mathbb{R}^n$, then $\|h(x) - h(y)\| = \Omega(\|x - y\|)$ and thus $\|f'(x) - f'(y)\| \geq \|h(x) - h(y)\| \geq \Omega(\|x - y\|)$.

As we see, if h satisfies conditions (1'), (2'), and (3'), then $f' = \tilde{f} \oplus h$ satisfies conditions (1), (2), and (3). Now we proceed with a formal proof. Our main task will be to define h appropriately.

PROOF. As above, let $\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a Lipschitz extension of f with $\|\tilde{f}\|_{\text{Lip}} = \|f\|_{\text{Lip}}$. Further, let $g = f^{-1} : f(X) \rightarrow X$ be the inverse map of f and $\tilde{g} : \mathbb{R}^m \rightarrow \mathbb{R}^n$ be its Lipschitz extension given by the Kirszbraun theorem. Denote $\alpha = \|\tilde{g}\|_{\text{Lip}}$. Since the distortion of f is at most D ,

$$\|f\|_{\text{Lip}} \leq D/\alpha, \quad \|\tilde{f}\|_{\text{Lip}} \leq D/\alpha,$$

$$\|g\|_{\text{Lip}} \leq \alpha, \quad \|\tilde{g}\|_{\text{Lip}} \leq \alpha, \quad \|\tilde{g} \circ \tilde{f}\|_{\text{Lip}} \leq D.$$

Let $h(x) = \frac{\tilde{g}(\tilde{f}(x)) - x}{\sqrt{2}\alpha}$ and $f'(x) = \tilde{f}(x) \oplus h(x) = \tilde{f}(x) \oplus \frac{\tilde{g}(\tilde{f}(x)) - x}{\sqrt{2}\alpha} \in \mathbb{R}^{n+m}$. We verify that f' satisfies conditions (1), (2), and (3) described in the proof overview above.

Condition (1). We prove that f' is an outer extension of f ; i.e., for every $x \in X$, we have

$$f'(x) = f(x) \oplus \frac{\tilde{g}(\tilde{f}(x)) - x}{\sqrt{2}\alpha} = f(x) \oplus \frac{g(f(x)) - x}{\sqrt{2}\alpha} = f(x) \oplus 0 = f(x).$$

Condition (2). For every $x, y \in \mathbb{R}^n$, we have

$$\begin{aligned} \sqrt{2}\alpha \cdot \|h(x) - h(y)\| &= \|(x - \tilde{g} \circ \tilde{f}(x)) - (y - \tilde{g} \circ \tilde{f}(y))\| \\ &\leq \|x - y\| + \|\tilde{g} \circ \tilde{f}(x) - \tilde{g} \circ \tilde{f}(y)\| \\ &\leq (1 + D)\|x - y\|. \end{aligned}$$

Thus,

$$\begin{aligned} \|f'(x) - f'(y)\|^2 &\leq \|\tilde{f}(x) - \tilde{f}(y)\|^2 + \|h(x) - h(y)\|^2 \\ &\leq \left(\left(\frac{D}{\alpha} \right)^2 + \frac{(1 + D)^2}{2\alpha^2} \right) \|x - y\|^2 \\ &\leq \frac{3D^2}{\alpha^2}. \end{aligned}$$

Therefore, $\|f'\|_{\text{Lip}} \leq \sqrt{3}D/\alpha$.

Condition (3). Finally, we prove that the Lipschitz constant of the inverse map f'^{-1} is at most $\sqrt{3}\alpha$. Consider two distinct points $x, y \in \mathbb{R}$. Let $\rho = \frac{\alpha \|\tilde{f}(x) - \tilde{f}(y)\|}{\|x - y\|}$. If $\rho \geq 1$, then $\|f'(x) - f'(y)\| \geq \|\tilde{f}(x) - \tilde{f}(y)\| \geq \|x - y\|/\alpha$. Otherwise, $\|\tilde{g}(\tilde{f}(x)) - \tilde{g}(\tilde{f}(y))\| \leq \rho\|x - y\| < \|x - y\|$, and

$$\begin{aligned} \|f'(x) - f'(y)\|^2 &= \|\tilde{f}(x) - \tilde{f}(y)\|^2 \\ &\quad + \frac{1}{2\alpha^2} \|(x - y) - (\tilde{g}(\tilde{f}(x)) - \tilde{g}(\tilde{f}(y)))\|^2 \\ &\geq \frac{\rho^2}{\alpha^2} \|x - y\|^2 + \frac{(1 - \rho)^2 \|x - y\|^2}{2\alpha^2} \\ &= \frac{(1 - 2\rho + 3\rho^2) \|x - y\|^2}{2\alpha^2} \\ &\geq \frac{\|x - y\|^2}{3\alpha^2}. \end{aligned}$$

Here we used that the minimum of the quadratic polynomial $1 - 2\rho + 3\rho^2$ equals $2/3$. In both cases, we have $\|f'(x) - f'(y)\| \geq \|x - y\|/\sqrt{3}\alpha$. Therefore, $\|f'^{-1}\|_{\text{Lip}} \leq \sqrt{3}\alpha$. We conclude that the distortion of f' is at most $3D$. $\square$

3 ONE-POINT EXTENSION OF NEAR-ISOMETRIC MAPS

3.1 Upper Bound

In this section, we prove Theorem 1.5. The theorem states that every near-isometric map can be extended to an extra point so that the extended map is also near isometric.

PROOF OF THEOREM 1.5. Without loss of generality, we can make several simplifying assumptions. First, it is sufficient to prove the theorem only for finite subsets X of $\mathbb{R}^n$; the statement for infinite subsets follows from a simple compactness argument. We will assume that $\varepsilon \in (0, 1)$, if $\varepsilon > 1$, the theorem follows from Theorem 1.2. Further, by rescaling f , if necessary, we may assume that

$\|v - w\| \leq \|f(v) - f(w)\| \leq (1 + \varepsilon)\|v - w\|$ for every $v, w \in X$. In particular,

$$\|v - w\|^2 \leq \|f(v) - f(w)\|^2 \leq (1 + 3\varepsilon)\|v - w\|^2. \quad (1)$$

If $u \in X$ then there is nothing to prove, so we assume that $u \notin X$. Let v_0 be the point closest to u in X (or one of the closest points to u if there is more than one such point). To simplify notation, we assume that $v_0 = 0$, $f(v_0) = 0$, and $\|v_0 - u\| = 1$. Then $\|u\| = 1$ and $\|u - v\| \geq 1$ for every $v \in X$. The theorem will follow from the following lemma.

LEMMA 3.1. *There exists a vector $u' \in \mathbb{R}^m$ such that*

- (1) $\|u'\| \leq 1$,
- (2) $|\langle u', f(v) \rangle - \langle u, v \rangle| \leq 3\sqrt{\varepsilon}(\|v\|^2 + 1)$ for every $v \in X$.

PROOF. Let $\Lambda = \{\lambda \in \mathbb{R}^X : \|\lambda\|_1 \leq 1\}$ be the unit ℓ_1 -ball in the space of functions $\lambda : X \rightarrow \mathbb{R}$ and $B = \{y \in \mathbb{R}^m : \|y\|_2 \leq 1\}$ be the unit ℓ_2 -ball in $\mathbb{R}^m$. Define

$$\Phi(y, \lambda) = \sum_{v \in X} \left(\lambda(v)(\langle u, v \rangle - \langle y, f(v) \rangle) - 3|\lambda(v)|\sqrt{\varepsilon}(\|v\|^2 + 1) \right).$$

We shall prove that there exists $u' \in B$ such that for every $\lambda \in \Lambda$, $\Phi(u', \lambda) \leq 0$. Observe that this u' will satisfy the statement of the lemma for the following reason. First, $\|u'\| \leq 1$. Second, let $I_v \in \Lambda$ be the indicator function of $v \in X$; then $\Phi(u', I_v) \leq 0$ and $\Phi(u', -I_v) \leq 0$. Therefore, $|\langle u, v \rangle - \langle u', f(v) \rangle| \leq 3\sqrt{\varepsilon}(\|v\|^2 + 1)$, as required.

To prove that such u' exists, we show that

$$\min_{y \in B} \max_{\lambda \in \Lambda} \Phi(y, \lambda) \leq 0.$$

Note that Λ and B are compact convex sets, Φ is linear in y and concave in λ ; thus, by the von Neumann minimax theorem [49],

$$\min_{y \in B} \max_{\lambda \in \Lambda} \Phi(y, \lambda) = \max_{\lambda \in \Lambda} \min_{y \in B} \Phi(y, \lambda).$$

Let $\hat{\lambda} \in \Lambda$ be the λ that maximizes the expression on the right. We need to prove that there is $\hat{y} \in B$ s.t. $\Phi(\hat{y}, \hat{\lambda}) \leq 0$. Define the point $P = \sum_{v \in V} \hat{\lambda}(v)v$ and $P' = \sum_{v \in V} \hat{\lambda}(v)f(v)$. For every $y \in B$, we have

$$\Phi(y, \hat{\lambda}) = \langle y, P \rangle - \langle y, P' \rangle - 3\sqrt{\varepsilon} \sum_{v \in X} |\hat{\lambda}(v)|\|v\|^2 - 3\sqrt{\varepsilon}\|\hat{\lambda}\|_1.$$

Now, $\langle u, P \rangle \leq \|P\|$ since $\|u\| \leq 1$. Let $\hat{y} = P'/\|P'\| \in B$. We have,

$$\Phi(\hat{y}, \hat{\lambda}) \leq \|P\| - \|P'\| - 3\sqrt{\varepsilon} \sum_{v \in X} |\hat{\lambda}(v)|\|v\|^2 - 3\sqrt{\varepsilon}\|\hat{\lambda}\|_1.$$

If $\|P\| \leq \|P'\|$ then $\Phi(\hat{y}, \hat{\lambda}) \leq 0$ and we are done. Similarly, if $\|P\| \leq 3\sqrt{\varepsilon} \sum_{v \in X} |\hat{\lambda}(v)|\|v\|^2$, we are done. We assume below that $\|P\| > \|P'\|$ and $\|P\| > 3\sqrt{\varepsilon} \sum_{v \in X} |\hat{\lambda}(v)|\|v\|^2$. Then,

$$\begin{aligned} \|P\| - \|P'\| &= \frac{\|P\|^2 - \|P'\|^2}{\|P\| + \|P'\|} \leq \frac{\|P\|^2 - \|P'\|^2}{\|P\|} \\ &= \frac{1}{\|P\|} \sum_{v, w \in X} \hat{\lambda}(v)\hat{\lambda}(w)(\langle v, w \rangle - \langle f(v), f(w) \rangle). \end{aligned}$$

Since f satisfies bi-Lipschitz condition (1) and $\|v - w\|^2 \leq 2(\|v\|^2 + \|w\|^2)$, we have

$$\begin{aligned} |\langle v, w \rangle - \langle f(v), f(w) \rangle| &= \frac{1}{2} |\|f(v) - f(w)\|^2 - \|f(v)\|^2 - \|f(w)\|^2| \\ &\leq \frac{3\varepsilon}{2} \max(\|v - w\|^2, \|v\|^2 + \|w\|^2) \\ &\leq 3\varepsilon(\|v\|^2 + \|w\|^2). \end{aligned}$$

Finally, we use that $\sum_{v \in V} |\hat{\lambda}(v)| = \|\hat{\lambda}\|_1$ and $\|P\| > 3\sqrt{\varepsilon} \sum_{v \in X} |\hat{\lambda}(v)|\|v\|^2$, and obtain

$$\begin{aligned} \|P\| - \|P'\| &\leq \frac{3\varepsilon}{\|P\|} \sum_{v, w \in X} |\hat{\lambda}(v)\hat{\lambda}(w)|(\|v\|^2 + \|w\|^2) \\ &= \frac{6\varepsilon\|\hat{\lambda}\|_1}{\|P\|} \sum_{v \in X} |\hat{\lambda}(v)|\|v\|^2 \\ &\leq \frac{6\varepsilon\|\hat{\lambda}\|_1}{3\sqrt{\varepsilon}} \leq 2\sqrt{\varepsilon}. \end{aligned}$$

Therefore, $\Phi(\hat{y}, \hat{\lambda}) < 0$. $\square$

Now we proceed with the proof of Theorem 1.5. Let $u' \in \mathbb{R}^m$ as in Lemma 3.1 and $w' = \sqrt{1 - \|u'\|^2}e_{m+1}$ (where e_{m+1} is a standard basis vector for $\mathbb{R}^{m+1}$). Note that w' is orthogonal to all vectors $f(v) \in \mathbb{R}^m$. Extend f to f' by letting $f'(u) = u' + w'$. Then, $\|f'(u)\|^2 = \|u'\|^2 + \|w'\|^2 = 1$. For every $v \in X$, we have

$$\begin{aligned} \|f'(v) - f'(u)\|^2 &= \|w'\|^2 + \|f(v) - u'\|^2 \\ &= (\|w'\|^2 + \|u'\|^2) + \|f(v)\|^2 - 2\langle f(v), u' \rangle \quad (2) \\ &= 1 + \|f(v)\|^2 - 2\langle f(v), u' \rangle, \quad (3) \\ \|v - u\|^2 &= 1 + \|v\|^2 - 2\langle v, u \rangle. \quad (4) \end{aligned}$$

From bounds $\|v - u\|^2 \geq 1$ and $\|v - u\|^2 \geq (\|v\| - 1)^2$, it easily follows that $\|v - u\|^2 \geq (\|v\|^2 + 1)/5$. By (3), (4), and the bound on $|\langle f(v), u' \rangle - \langle v, u \rangle|$ from Lemma 3.1, we have

$$\begin{aligned} \left| \|f'(v) - f'(u)\|^2 - \|v - u\|^2 \right| &\leq 3\varepsilon\|v\|^2 + 6\sqrt{\varepsilon}(\|v\|^2 + 1) \\ &\leq 9\sqrt{\varepsilon}(\|v\|^2 + 1) \\ &\leq 45\sqrt{\varepsilon}\|v - u\|^2. \end{aligned}$$

This implies that f' has distortion $1 + O(\sqrt{\varepsilon})$. $\square$

3.2 Lower Bound

In this section, we show that the bound in Theorem 1.5 is tight (up to a constant factor in the O -notation) – extending a map with distortion $1 + \varepsilon$ by one point might require blowing up the distortion to $1 + \Omega(\sqrt{\varepsilon})$, even when $n = m = 1$ (the extension f' may use extra dimensions).

The construction is as follows. Consider points: $A = 0$, $B = \varepsilon$, $B' = -\varepsilon$, and $C = 1$. Let $X = \{A, B, C\}$. Consider map $f : X \rightarrow \mathbb{R}$ that maps A, B, C to points A, B', C , respectively. Clearly f has distortion $\frac{C-B'}{C-B} = \frac{1+\varepsilon}{1-\varepsilon} \leq 1 + 3\varepsilon$ for $\varepsilon \leq 1/3$. Our goal is to extend f to the fourth point $D = \sqrt{\varepsilon}$. Note that we can assume that the extension uses at most one additional dimension.

CLAIM 3.2. *Any outer extension of the map f to the point D has distortion at least $(1 + \sqrt{\varepsilon}/2)$.*

PROOF. Let $f(D) = (x, y) \in \mathbb{R}^2$, and suppose that the distortion is less than $(1 + \sqrt{\varepsilon}/2)$. Then we must have

- $\|f(D) - f(A)\| \geq (1 - \sqrt{\varepsilon}/2) \|D - A\|$, so $x^2 + y^2 \geq (\sqrt{\varepsilon} - \varepsilon/2)^2$.
- $\|f(D) - f(B)\| \leq (1 + \sqrt{\varepsilon}/2) \|D - B\|$, so $(x + \varepsilon)^2 + y^2 \leq ((1 + \sqrt{\varepsilon}/2)(\sqrt{\varepsilon} - \varepsilon))^2 = (\sqrt{\varepsilon} - \varepsilon/2)^2$.

We get that $x^2 + y^2 \geq (x + \varepsilon)^2 + y^2$. Thus, $x \leq -\varepsilon/2$. Then $\frac{\|f(D) - f(C)\|}{\|D - C\|} \geq \frac{1}{1 - \sqrt{\varepsilon}} \geq 1 + \sqrt{\varepsilon}$, which is a contradiction. $\square$

4 APPLICATIONS – PRIORITIZED AND TERMINAL DIMENSION REDUCTIONS

In this section, we prove Theorems 1.9 and 1.11.

PROOF OF THEOREM 1.9. First, what we do is to we construct a $(c_1 \log \log j, c_2 \log^{3+\varepsilon} j)$ -prioritized dimension reduction. Denote $C = 3 + \varepsilon$. We define an increasing family of $T = \lceil \log_C \log_2 \log_2 N \rceil$ subsets $S_0, S_1, \dots, S_T$ of $X \subset \mathbb{R}^d$: S_i consists of the first $\min(2^{2^{C^i}}, N)$ points according to the priority ranking π .

For each set S_i , we construct an embedding $f_i : S_i \rightarrow \mathbb{R}^{d_i}$ with distortion at most C^i for $d_i = O(\log |S_i|)$ in such a way that each f_i is an outer extension of f_{i-1} . We start with S_0 – we let f_0 be an isometric embedding of S_0 (which consists of 4 points) into $\mathbb{R}^3$. Then we iteratively construct mapping f_i . At iteration i , we take map f_{i-1} and extend it to map f_i as follows. Using Theorem 1.2, we find an outer-bi-Lipschitz extension $h : S \rightarrow \mathbb{R}^{d'}$ of f_{i-1} to S_i . The extension h is not yet what we want:

- while, by Theorem 1.2, its distortion is at most $3 \cdot (3 + \varepsilon)^{i-1}$, which is less than C^i (the desired upper bound on the distortion),
- the dimension d' is possibly greater than $\Omega(\log |S_i|)$.

To reduce the dimension, we write $h(x) = h_1(x) \oplus h_2(x) \in \mathbb{R}^{d_{i-1}} \oplus \mathbb{R}^{d' - d_{i-1}}$, here $h_1(x)$ is the vector consisting of the first d_{i-1} coordinates of $h(x)$ and $h_2(x)$ is the vector consisting of the remaining coordinates of $h(x)$. Since h is an extension of f_{i-1} , we have $h_1(x) = f_{i-1}(x)$ and $h_2(x) = 0$ for $x \in S_{i-1}$. Now, we use the Johnson–Lindenstrauss lemma to find a dimension reduction g from $h_2(S_i)$ to $\mathbb{R}^{d''}$ with distortion at most $1 + \varepsilon/3$, where $d'' = c_{JL} \log |S_i|/\varepsilon^2$ for some absolute constant c_{JL} . We assume that $g(0) = 0$ (if necessary, we redefine g as $g'(x) = g(x) - g(0)$). Finally, we let $f_i = (id \oplus g) \circ h$; in other words, $f_i(x) = h_1(x) \oplus g(h_2(x))$.

Note that $f_i(x)$ is an outer extension of f_{i-1} , since $f_i(x) = h_1(x) \oplus g(h_2(x)) = f_{i-1}(x) \oplus g(0) = f_{i-1}(x)$ for $x \in S_{i-1}$. The distortion of $id \oplus g$ is at most the distortion of g , which is at most $1 + \varepsilon/3$; therefore, the distortion of f_i is at most $(1 + \varepsilon/3) \times 3 \cdot (3 + \varepsilon)^{i-1} = C^i$. We bound the dimension

$$\begin{aligned} d_i &= d_{i-1} + d'' = 4 + \sum_{t=1}^i c_{JL} \log |S_t|/\varepsilon^2 \\ &\leq 4 + \sum_{t=1}^{i-1} c_{JL} 2^{C^t}/\varepsilon^2 + c_{JL} \log |S_i|/\varepsilon^2 \\ &= O(\log |S_i|). \end{aligned}$$

The constant in the big- O notation is proportional to $1/\varepsilon^2$.

Finally, let $f = f_T$. We verify f is $(c_1 \log \log j, c_2 \log^{3+\varepsilon} j)$ -prioritized dimension reduction. Fix some $j \in \{1, \dots, N\}$. Let S_i be the smallest of the sets $S_0, \dots, S_T$ that contains $\pi(j)$; i.e., $i = \lceil \log_C \log_2 \log_2 j \rceil$ if $j > 4$, and $i = 0$ otherwise. Then f restricted to $\pi(1), \dots, \pi(j)$ coincides with f_i . The distortion of f_i is at most (for $j \geq 4$)

$$C^i \leq C^{1+\log_C \log_2 \log_2 j} \leq C \log_2 \log_2 j = (3 + \varepsilon) \log_2 \log_2 j.$$

Further, $f(\pi(j)) = f_i(\pi(j)) \in \mathbb{R}^{d_i}$. Hence, in the vector $f(\pi(j))$ all but the first d_i coordinates are equal to 0; we upper bound d_i as follows (for $j \geq 4$):

$$d_i \leq O(\log |S_i|) \leq O(2^{C^i}) \leq O\left((2^{C^{i-1}})^C\right) \leq O(\log j)^C,$$

as required. Note that the image of X under f lies in space $\mathbb{R}^{d_T}$ of dimension $d_T = O(\log |S_T|) = O(\log N)$.

By setting the parameters differently, we can obtain different trade-offs between the distortion and dimension. Fix a parameter $k \in \mathbb{N}$, $1 < k < \log \log \log N$. Let $T = k$ and S_i be the set consisting of the first $2^{\log_2^{1/k} N}$ points in X , according to the priority ordering π . Construct maps f_i as described above. The distortion of f is at most $C^T = (3 + \varepsilon)^k$. The vector $f(\pi(j))$ lies in the space $\mathbb{R}^{d_i}$, where $i = \lceil k \log_2 \log_2 j \rceil$ and

$$\begin{aligned} d_i &\leq \frac{c_{JL}}{\varepsilon^2} \sum_{t=0}^i \log |S_t| = O\left(\sum_{t=0}^i \log_2^{t/k} N\right) \leq O(\log_2^{i/k} N) \\ &= O(\underbrace{\log_2^{(i-1)/k} N}_{\text{less than } \log_2 j} \cdot \log_2^{1/k} N) \leq O(\log_2 j \log_2^{1/k} N). \end{aligned}$$

We can compute map f in polynomial time, since, at each iteration, we can compute the outer extension h and dimension reduction g in polynomial time. $\square$

Now we prove Theorem 1.11.

PROOF OF THEOREM 1.11. First we apply the Johnson–Lindenstrauss lemma to X with $\varepsilon' = \varepsilon^2$. We get an embedding $g : X \rightarrow \mathbb{R}^{d'}$ with the distortion at most $1 + \varepsilon^2$ and $d' = O(\log N/\varepsilon^4)$; we rescale it so that $\lambda \|x - y\| \leq \|g(x) - g(y)\| \leq \lambda(1 + \varepsilon^2) \|x - y\|$, where $\lambda = 1 + c\varepsilon$ (we will specify c later).

For every point $p \in \mathbb{R}^d$, we extend g to a map $g_p : X \cup \{p\} \rightarrow \mathbb{R}^{d'+1}$ using Theorem 1.5; for $p \in X$, $g_p = g$. The distortion of g_p is $1 + O(\sqrt{\varepsilon^2}) = 1 + O(\varepsilon)$. Finally, we let $f(p) = g_p(p)$. The image of f lies in $\mathbb{R}^{d'+1}$, as required. For every $x \in X$ and $p \in \mathbb{R}^d$, we have $g_x(x) = g(x) = g_p(x)$ and

$$\begin{aligned} \|f(p) - f(x)\| &= \|g_p(p) - g_x(x)\| = \|g_p(p) - g_p(x)\| \\ &\in [(1 + c\varepsilon)(1 - O(\varepsilon))\|p - x\|, (1 + c\varepsilon)(1 + O(\varepsilon))\|p - x\|]. \end{aligned}$$

We choose c so that the $(1 + c\varepsilon)(1 - O(\varepsilon))$ term is 1; then $(1 + c\varepsilon)(1 + O(\varepsilon)) = 1 + O(\varepsilon)$.

Note that we can compute $f(x)$ in polynomial time, since we can compute each map g_p in polynomial time. $\square$

5 OUTER EXTENSION OF A MAP FROM $\mathbb{R}$ TO $\mathbb{R}$

In this section, we consider the case of map $f: X \rightarrow \mathbb{R}$ with distortion $(1 + \varepsilon)$, where $X \subset \mathbb{R}$. We show that such a map is very structured, which allows us to extend it to $\tilde{f}: \mathbb{R} \rightarrow \mathbb{R}^2$ with the distortion $1 + O(1/\log^2(1/\varepsilon))$. Here we provide an informal overview to illustrate the main steps.

First, suppose that X consists of three points $0, \varepsilon, 1$ that f maps to $0, -\varepsilon, 1$, respectively. It turns out that this simple case is in fact very important. We extend f to the whole segment $[0; 1]$ as follows¹. For $0 \leq x \leq \varepsilon$, we map x to $(-x, 0)$, and for $\varepsilon \leq x \leq 1$, we map x to point $g(x) = (r(x), \varphi(x))$ in polar coordinates, where the radius is $r(x) = x$ and the angle is $\varphi(x) = \frac{\pi \ln(1/x)}{\ln(1/\varepsilon)}$, see Figure 1. First, the map is continuous (i.e., $g(\varepsilon) = -\varepsilon$ and $g(1) = 1$). Second, for every x , $\|g(x)\| = |x|$, which implies that g is non-contractive. We refer to this map as the “spiral”. We prove that its distortion is $1 + O(1/\ln^2(1/\varepsilon))$, and in fact this is the optimal distortion one can achieve for this specific choice of X and f (see Section 5.3 for the proof).

For the general case, we decompose f into “flips” and use this decomposition to assemble the extension from the above spirals on various distance scales.

For a set X and map f , consider how f changes the relative ordering of points X ; denote the corresponding permutation by $\pi_f \in S_{|X|}$. For instance, if $X = \{x_1, x_2, x_3\}$, where $x_1 < x_2 < x_3$, and $f(x_1) < f(x_3) < f(x_2)$, we set $\pi_f = (1 \ 3 \ 2)$. We show that a permutation can arise as π_f for some f iff it excludes $(3 \ 1 \ 4 \ 2)$ and $(2 \ 4 \ 1 \ 3)$ as a subpermutation. Furthermore, we show that π_f can be decomposed into a laminar sequence of *flips*. We start with the identity permutation, and then iteratively choose a substring and reverse its order (this is one flip). We do this so that every two flips are either disjoint, or the later is strictly contained in the earlier one. For example, if $\pi_f = (3 \ 1 \ 2 \ 4 \ 6 \ 5)$, then the decomposition is as follows: $(1 \ 2 \ 3 \ 4 \ 5 \ 6)$, $(3 \ 2 \ 1 \ 4 \ 5 \ 6)$, $(3 \ 1 \ 2 \ 4 \ 5 \ 6)$, $(3 \ 1 \ 2 \ 4 \ 6 \ 5)$.

We use this decomposition to build the desired extension. For each flip, we add two spirals. We show that the points that participate in a given flip are well-separated from others. For example if the permutation is $(1 \ 3 \ 2)$, then the distance between 2 and 3 should be much smaller by a factor of ε than the distance from 1 to either of them – both in the domain and in the image. We show that this separation is sufficient for these spirals not to interfere much with each other, and the bound of $1 + O(1/\log^2(1/\varepsilon))$ on the distortion holds for the overall construction. See Figure 1, for the construction for the case $\pi_f = (3 \ 1 \ 2 \ 4 \ 6 \ 5)$.

5.1 Extension to the Whole Line

In this section we prove the following theorem.

THEOREM 5.1 (THEOREM 1.4). *Let $X \subset \mathbb{R}$ be an arbitrary set. Suppose that $f: X \rightarrow \mathbb{R}$ is a map such that for every $x_1, x_2 \in X$, we have:*

$$|f(x_1) - f(x_2)| \in (1 \pm \varepsilon) \cdot |x_1 - x_2|. \quad (5)$$

Then there exists a map $h: \mathbb{R} \rightarrow \mathbb{R}^2$ such that:

- For every $x \in X$, we have $h(x) = (f(x), 0)$;

¹Extending f to the whole $\mathbb{R}$ requires a bit more work.




Figure 1: Top: one possible extension for the map $0 \mapsto 0$, $\varepsilon \mapsto -\varepsilon$, $1 \mapsto 1$. It has distortion $1 + O(1/\log^2(1/\varepsilon))$, which is tight for this example. Bottom: an extension built from the spirals recursively for the map f with $\pi_f = (3 \ 1 \ 2 \ 4 \ 6 \ 5)$. The picture is intentionally out of proportion.

- For every $u, v \in \mathbb{R}$, we have

$$\|h(u) - h(v)\| \in \left(1 \pm O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) \cdot |u - v|.$$

By a standard compactness argument, it is enough to handle the case of a finite X . From now on, we denote $n = |X|$.

5.1.1 Characterizing near-isometric maps. To prove the main theorem, we will first prove the necessary conditions f needs to satisfy in order to be a near-isometric mapping. In the rest, we will denote the initial point set by $X = \{x_1, x_2, \dots, x_n\}$ and without loss of generality we may assume that $x_1 < x_2 < \dots < x_n$. Let $\pi_f \in S_n$ be the permutation defined by our mapping f such that $f(x_{\pi_f(1)}) < f(x_{\pi_f(2)}) < \dots < f(x_{\pi_f(n)})$. The following lemma characterizes the properties of π_f .

Definition 5.2 (Sub-permutation). Given a permutation σ of $[k]$, and a permutation π of $[n]$, where $n \geq k$, we say that π contains σ as a sub-permutation iff there exists $i_1 < \dots < i_k \in [n]$ such that for any $j, j' \in [k]$, if $\sigma(j) < \sigma(j')$, then $\pi(i_j) < \pi(i_{j'})$.

LEMMA 5.3. *If $\varepsilon > 0$ is sufficiently small, then π_f does not have $(3 \ 1 \ 4 \ 2)$ or $(2 \ 4 \ 1 \ 3)$ as sub-permutations.*

PROOF. Let us prove the statement for (3 1 4 2), the proof for (2 4 1 3) is the same. Assume the contrary. Then, there exists $1 \leq i < j < k < l \leq n$ such that

$$f(x_k) < f(x_i) < f(x_l) < f(x_j). \quad (6)$$

Denote $\Delta = x_l - x_i > 0$. Then,

$$\begin{aligned} \Delta &= x_l - x_i \geq (x_l - x_k) + (x_j - x_i) \\ &\geq (1 - O(\varepsilon)) \cdot ((f(x_l) - f(x_k)) + (f(x_j) - f(x_i))) \\ &\geq (2 - O(\varepsilon)) \cdot (f(x_l) - f(x_i)) \\ &\geq (2 - O(\varepsilon)) \cdot (x_l - x_i) \\ &= (2 - O(\varepsilon)) \cdot \Delta, \end{aligned}$$

where the first step follows from $x_i < x_j < x_k < x_l$ (which in turn follows from $i < j < k < l$), the second step follows from f having distortion $(1 + \varepsilon)$ and from (6), and the fourth step again follows from f being a near-isometry. Thus, if $\varepsilon > 0$ is sufficiently small, we get a contradiction. $\square$

5.1.2 Permutation decomposition.

LEMMA 5.4. If $\varepsilon > 0$ is sufficiently small, then π_f can be decomposed as follows. We start with π_0 which is the identity permutation. Then, we perform $T \geq 0$ flips as follows. Each flip $1 \leq t \leq T$ is defined by two numbers $1 \leq a_t < b_t \leq n$, naturally defining a segment in the permutation. We obtain π_t from π_{t-1} as follows.

$$\pi_t(k) = \begin{cases} \pi_{t-1}(a_t + b_t - k), & \text{if } a_t \leq k \leq b_t, \\ \pi_{t-1}(k), & \text{otherwise.} \end{cases}$$

In words, we obtain π_t from π_{t-1} by reversing the segment $[a_t, b_t]$. Moreover, the segments form a laminar family: for every $1 \leq t_1 < t_2 \leq T$ the segments $[a_{t_1}, b_{t_1}]$ and $[a_{t_2}, b_{t_2}]$ are either disjoint or $[a_{t_1}, b_{t_1}] \supset [a_{t_2}, b_{t_2}]$. The permutation π_f is equal to the final permutation π_T .

PROOF. The proof is by induction over n . If $n = 1$, the statement is trivial. Denote $1 \leq u \leq n$ such that $\pi_f(u) = 1$ (the position where 1 is mapped to), and $1 \leq v \leq n$ such that $\pi_f(v) = n$ (the position where n is mapped to). Suppose that $u < v$. If $u = 1$, then the statement follows from using the induction assumption on π_f without the first element. Assume that $u > 1$. Then, define $A = \{\pi_f(j) \mid j \leq u\}$, to be the set of numbers that are mapped to the left of 1. Let $z < u$ be such that $\pi_f(z) = \max A$, i.e., the maximum number mapped to the left of 1. Define $w = \min\{k \mid \pi_f(k) > \max A\}$. Clearly, $w \leq v$. We claim that the sequence $(\pi_f(1) \pi_f(2) \dots \pi_f(w-1))$ is a permutation of the numbers from 1 to z . Assume not. Then, there exists $w' > w$ such that $\pi_f(w') < z$. Then, considering positions z, u, w , and w' , we obtain a sub-permutation (3 1 4 2), which can not be the case by Lemma 5.3. Now we can apply the inductive assumption on the first $w-1$ numbers, and on the last $n-w+1$ numbers, and merge the resulting sequences of flips. If $u > v$, then we add a flip with $a = 1$ and $b = n$ and reduce to the case, when $u < v$. $\square$

It is not hard to show that the above condition is also a sufficient condition, but we will not need it in our construction.




Figure 2: Illustration to the proof of Lemma 5.8

5.1.3 Well-separateness and the portals. First, for each flip $1 \leq t \leq T$, we define the set of points F_t that are affected by it, the set of points to the left of F_t , denoted L_t , and the points to the right, R_t . Formally, we have the following.

Definition 5.5. For an iteration $1 \leq t \leq T$, we define

- $L_t = \{\pi_t(1), \pi_t(2), \dots, \pi_t(a_t - 1)\}$;
- $F_t = \{\pi_t(a_t), \pi_t(a_t + 1), \dots, \pi_t(b_t)\}$;
- $R_t = \{\pi_t(b_t + 1), \pi_t(b_t + 2), \dots, \pi_t(n)\}$.

LEMMA 5.6. F_t is the set of $|F_t| = b_t - a_t + 1$ consecutive integers. Moreover, the sequence $\pi_t(a_t), \pi_t(a_t + 1), \dots, \pi_t(b_t)$ is either increasing or decreasing.

PROOF. Follows trivially from Lemma 5.4. $\square$

Definition 5.7. For an iteration $t \leq T$, we define $u_t = \pi_{t-1}(a_t)$ and $v_t = \pi_{t-1}(b_t)$. We also define $\Delta_t = x_{v_t} - x_{u_t}$. It can be either positive or negative.

The quantity Δ_t can be seen as the signed diameter of the flipped points. The following lemma is a key to the overall analysis. We show that the flipped points F_t are very well-separated from the remainder: by the amount $\Omega(|\Delta_t|/\varepsilon)$.

LEMMA 5.8. For every $k \in F_t$, and every $p \in L_t \cup R_t$, we have $|x_k - x_p| \geq \Omega\left(\frac{|\Delta_t|}{\varepsilon}\right)$.

PROOF. Wlog, we can assume that t is the first flip that separates p and k and for which $k \in F_t$, but $p \notin F_t$. Indeed, if $\tilde{t} < t$ is the first such flip, then $|\Delta_{\tilde{t}}| > |\Delta_t|$, and the required statement follows from that about $\tilde{t}$. Suppose that $p \in L_t$, the case $p \in R_t$ is similar. Then, we have $f(x_p) < f(x_{v_t}) < f(x_{u_t})$ (here we use crucially the fact that t is the first flip that separates p and k). Indeed, t is the last flip, which affects the relative order of $f(x_p)$, $f(x_{v_t})$ and $f(x_{u_t})$, since the flips that are not disjoint are nested. At the same time, either $x_p < x_{u_t} \leq x_k \leq x_{v_t}$ or $x_p > x_{u_t} \geq x_k \geq x_{v_t}$. Let us show how to handle the first case, the second case is similar. Let us denote $s = x_{u_t} - x_p$. See Figure 2 for the clarification. Then,

$$\begin{aligned} s(1 + \varepsilon) &\geq f(x_{u_t}) - f(x_p) \\ &= (f(x_{u_t}) - f(x_{v_t})) + (f(x_{v_t}) - f(x_p)) \\ &\geq (1 - \varepsilon)(x_{v_t} - x_{u_t}) + (1 - \varepsilon)(x_{v_t} - x_p) \\ &= (1 - \varepsilon)(s + 2\Delta_t). \end{aligned}$$

Thus, $\Delta_t = O(\varepsilon \cdot s)$. Finally, $|x_k - x_p| \geq s = \Omega(\Delta_t/\varepsilon)$. $\square$

Definition 5.9 (Portals). For every $1 \leq t \leq T$, we define portals as follows (see Figure 3). We set:

$$\begin{aligned} \bullet \alpha_t &= x_{u_t} - \frac{\Delta_t}{\varepsilon^{2/3}}; \beta_t = x_{u_t} - \frac{\Delta_t}{\varepsilon^{1/3}}; \gamma_t = x_{v_t} + \frac{\Delta_t}{\varepsilon^{1/3}}; \delta_t = \\ &= x_{v_t} + \frac{\Delta_t}{\varepsilon^{2/3}}; \end{aligned}$$




Figure 3: Portals. Note that the scales of the relative distances are not correct.

$$\begin{aligned} \bullet \quad \alpha'_t &= f(x_{v_t}) - \frac{\Delta_t}{\varepsilon^{2/3}}; \beta'_t = f(x_{v_t}) - \frac{\Delta_t}{\varepsilon^{1/3}}; \gamma'_t = f(x_{u_t}) + \frac{\Delta_t}{\varepsilon^{1/3}}; \\ \delta'_t &= f(x_{u_t}) + \frac{\Delta_t}{\varepsilon^{2/3}}. \end{aligned}$$

We will use the portals in our construction to make sure that the spirals at different levels do not interfere with each other.

5.1.4 Construction of the final map. Now we are ready to define the final map $h: \mathbb{R} \rightarrow \mathbb{R}^2$. First, for every $1 \leq k \leq n$, we set $h(x_k) = (f(x_k), 0)$. Second, for every $1 \leq t \leq T$, we define h between α_t and β_t and between γ_t and δ_t according to the Corollary 5.19 (note that we only take the part of the map which corresponds to these two intervals, see Figure 3 for the illustration). In particular, $h(\alpha_t) = (\alpha'_t, 0)$, $h(\beta_t) = (\beta'_t, 0)$, $h(\gamma_t) = (\gamma'_t, 0)$ and $h(\delta_t) = (\delta'_t, 0)$. After we are done with constructing the spirals for all iterations t , on the remaining bounded intervals on the real line, we define h to be linear and consistent with the values at the endpoints. For the two unbounded intervals, we define the map to be appropriate shifts.

Let us now show that for every $x, y \in \mathbb{R}$, we have:

$$\|h(x) - h(y)\| \in \left(1 \pm O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) \cdot |x - y|.$$

For a point $t \in \mathbb{R}$, there are two cases: either it is mapped using the map g from Corollary 5.19, or it is mapped using a linear extension. In the former case, we say that t is of “type A”, while in the latter case it is said to be of “type B”. Note that the type A points are mapped on a spiral curve in $\mathbb{R}^2$, and the type B points are mapped on a segment in $\mathbb{R}$.

CLAIM 5.10. *If we extend the original map f to the portals (such that $\alpha_t \mapsto \alpha'_t$, $\beta_t \mapsto \beta'_t$, $\gamma_t \mapsto \gamma'_t$ and $\delta_t \mapsto \delta'_t$), then the resulting map is a $(1 \pm O(\varepsilon^{1/3}))$ -isometry.*

PROOF. It is immediate to check that the worst case is achieved when we consider distances between portals α_t and β_t or γ_t and δ_t . In this case, the distortion is $1 + \Theta(\varepsilon^{1/3})$ (this follows from the definition of the portals). $\square$

CLAIM 5.11. *If $t \in \mathbb{R}$ is type B, and h is smooth at t , then $\|\nabla h(t)\|_2 = 1 \pm O(\varepsilon^{1/3})$.*

PROOF. This is a direct corollary of Claim 5.10. $\square$

CLAIM 5.12. *If both $x, y \in \mathbb{R}$ are type B, then*

$$\|h(x) - h(y)\| \in \left(1 \pm O\left(\varepsilon^{1/3}\right)\right) \cdot |x - y|.$$

PROOF. If $x = y$, then there is nothing to prove. If $x \neq y$ by a small perturbation we can assume wlog that h is smooth in both x and y . By Claim 5.11, $\|\nabla h(x)\|_2, \|\nabla h(y)\|_2 \in 1 \pm O(\varepsilon^{1/3})$. If the signs of $(\nabla h(x))_1$ and $(\nabla h(x))_2$ are the same, then the claim follows from Claim 5.10 and Claim 5.11.

Now consider the case of the different signs of the derivatives. Then consider an extension of f to the portals as stated in Claim 5.10. Abusing notation, let us denote this map f as well. Since the extended map has distortion $1 \pm O(\varepsilon^{1/3})$, we decompose it as per Lemma 5.4, and we get that Lemma 5.8 holds.

Let us denote $p_x < x < q_x$ the portals of elements which are closest to x , similarly, we denote $p_y < y < q_y$. Wlog, $q_x \leq p_y$. If a decomposition for f has a flip containing p_y and q_y , but not p_x and q_x , then $p_y - q_x \geq \Omega\left(\frac{q_y - p_y}{\varepsilon^{1/3}}\right)$. Similarly, if there is a flip containing p_x and q_x , but not p_y and q_y , then $p_y - q_x \geq \Omega\left(\frac{q_x - p_x}{\varepsilon^{1/3}}\right)$. Note that if neither of these two cases hold, then their gradients could not have different signs. Combining these observations with Claim 5.10 and Claim 5.11, we get the required result. $\square$

CLAIM 5.13. *If both $x, y \in \mathbb{R}$ are type A, then*

$$\|h(x) - h(y)\| \in \left(1 \pm O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) \cdot |x - y|.$$

PROOF. Define t_x to be the flip $1 \leq t \leq T$, such that x lies between α_t and β_t or γ_t and δ_t . We define t_y similarly.

If $t_x = t_y$, then the claim follows from Corollary 5.19.

First, suppose that $[a_{t_x}, b_{t_x}]$ and $[a_{t_y}, b_{t_y}]$ are disjoint. Assume wlog that $|\Delta_{t_x}| \geq |\Delta_{t_y}|$. Then,

$$\begin{aligned} \|h(x) - h(y)\| &= \|h(\alpha_{t_x}) - h(\alpha_{t_y})\| \pm O(|\Delta_{t_x}|/\varepsilon^{2/3}) \\ &\in (1 \pm O(\varepsilon^{1/3}))|\alpha_{t_x} - \alpha_{t_y}| \pm O(|\Delta_{t_x}|/\varepsilon^{2/3}) \\ &\in (1 \pm O(\varepsilon^{1/3}))|x - y| \pm O(|\Delta_{t_x}|/\varepsilon^{2/3}) \\ &\in (1 \pm O(\varepsilon^{1/3}))|x - y|, \end{aligned}$$

where the first step follows from Corollary 5.19, the second step follows from Lemma 5.8, the third step follows from the definition of the terminals, and the last step follows from Lemma 5.8.

Now assume that $[a_{t_x}, b_{t_x}] \supseteq [a_{t_y}, b_{t_y}]$, but $t_x \neq t_y$. Then, we have $|x - y| \geq \Omega(|\Delta_{t_x}|/\varepsilon^{1/3})$, $|\Delta_{t_x}| = \Omega(|\Delta_{t_y}|/\varepsilon)$ and:

$$\begin{aligned} \|h(x) - h(y)\| &= |h(x) - h(\alpha_{t_y})| \pm O(|\Delta_{t_y}|/\varepsilon^{2/3}) \\ &\in \left(1 \pm O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) |x - \alpha_{t_y}| \pm O(|\Delta_{t_y}|/\varepsilon^{2/3}) \\ &\in \left(1 \pm O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) |x - y| \pm O(|\Delta_{t_y}|/\varepsilon^{2/3}) \\ &\in \left(1 \pm O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) |x - y|, \end{aligned}$$

where the first step is due to the definition of the portals and Corollary 5.19, the second step is due to Corollary 5.19, the third step is again due to the definition of the portals, and the last step is due to $|x - y| \geq \Omega(|\Delta_{t_x}|/\varepsilon^{1/3}) \geq \Omega(|\Delta_{t_y}|/\varepsilon^{4/3})$. $\square$




Figure 4: Illustration for the proof of Claim 5.14.

CLAIM 5.14. If $x \in \mathbb{R}$ is type A and $y \in \mathbb{R}$ is type B, then

$$\|h(x) - h(y)\| \in \left(1 \pm O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) \cdot |x - y|.$$

PROOF. Denote $1 \leq t_x \leq T$ to be the flip such that x lies within α_{t_x} and β_{t_x} or between γ_{t_x} and δ_{t_x} . Wlog, let us assume that x lies between α_{t_x} and β_{t_x} . Then, y can lie between β_{t_x} and γ_{t_x} or outside of the segment connecting α_{t_x} and δ_{t_x} . Let us assume the former, and the latter can be handled similarly. By Corollary 5.19, we have:

$$\|h(y) - h(x)\| \in O\left(1 \pm O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) \cdot |x - \tilde{y}|, \quad (7)$$

where $\tilde{y} - \beta_t = \gamma'_t - h(y)_1$ (see Figure 4). By Claim 5.12,

$$|\beta_t - \tilde{y}| = \|h(\beta_t) - h(y)\| \in (1 \pm O(\varepsilon^{1/3})) \cdot |\beta_t - y|.$$

Thus,

$$|y - \tilde{y}| \in O(\varepsilon^{1/3}) \cdot |\beta_t - y| \leq O(\varepsilon^{1/3}) \cdot |x - y|. \quad (8)$$

Combining (7) and (8), we are done. $\square$

5.2 An Auxiliary Map: the Spiral

LEMMA 5.15. Let $\varepsilon > 0$ be a small positive parameter. Let $g: \mathbb{R} \rightarrow \mathbb{R}^2$ be the map defined as follows.

$$g(t) = \begin{cases} (t, 0), & \text{if } |t| > 1, \\ (-t, 0), & \text{if } |t| < \varepsilon \\ (t \cos \varphi(t), t \sin \varphi(t)), & \text{where } \varphi(t) = \frac{\pi \ln(1/|t|)}{\ln(1/\varepsilon)} \text{ otherwise.} \end{cases}$$

Where the third term can be viewed in the polar coordinates as $(r(t) = t, \varphi(t) = \frac{\pi \ln(1/|t|)}{\ln(1/\varepsilon)})$. Then we have the following properties,

- Distortion: for every $t_1, t_2 \in \mathbb{R}$, one has:

$$\|g(t_1) - g(t_2)\| \in \left(1 + O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) \cdot |t_1 - t_2|;$$

- Total movement: for every $t \in \mathbb{R}$, one has:

$$\|g(t) - (t, 0)\| \leq O(1).$$

PROOF. First of all note that the function is continuous as

$$\begin{aligned} g(\varepsilon) &= (\varepsilon \cos \varphi(\varepsilon), \varepsilon \sin \varphi(\varepsilon)) = (-\varepsilon, 0), \\ g(-\varepsilon) &= (\varepsilon, 0), \\ g(1) &= (\cos \varphi(1), \sin \varphi(1)) = (1, 0), \text{ and} \\ g(-1) &= (-1, 0). \end{aligned}$$

Next we show that the distortion is bounded as desired. First, we prove that g does not increase the distance by more than a multiplicative factor of $1 + O(\frac{1}{\ln^2(1/\varepsilon)})$, and second in Claim 5.17, we prove that the distances do not decrease by more than the same factor. These two prove the bound on the distortion as desired. Finally in Claim 5.18, we show the total movement property.

CLAIM 5.16. For $\varepsilon \leq t_1 < t_2 \leq 1$, we have $\|g(t_1) - g(t_2)\| \leq (1 + O(\frac{1}{\log^2(1/\varepsilon)})) \cdot |t_1 - t_2|$.

PROOF. The distance between $g(t_1)$ and $g(t_2)$ is at most the length of the curve between them which is given by the following formula

$$\begin{aligned} & \int_{t=t_1}^{t_2} \sqrt{\left(\frac{d(t \cos \varphi(t))}{dt}\right)^2 + \left(\frac{d(t \sin \varphi(t))}{dt}\right)^2} dt = \\ & \int_{t=t_1}^{t_2} \sqrt{\left(\cos \varphi(t) + \frac{\pi \sin \varphi(t)}{\ln(1/\varepsilon)}\right)^2 + \left(\sin \varphi(t) - \frac{\pi \cos \varphi(t)}{\ln(1/\varepsilon)}\right)^2} dt \\ & = \int_{t=t_1}^{t_2} \sqrt{1 + \left(\frac{\pi}{\ln(1/\varepsilon)}\right)^2} dt \\ & = (t_2 - t_1) \sqrt{1 + \left(\frac{\pi}{\ln(1/\varepsilon)}\right)^2} \\ & \leq (t_2 - t_1) \left(1 + \frac{\pi^2}{2 \ln^2(1/\varepsilon)}\right) \end{aligned}$$

$\square$

The above claim, together with the fact that the function is symmetric around the origin, and the definition of the function for $|t| \geq 1$ and $|t| \leq \varepsilon$, and triangle inequality, proves that for any $t_1, t_2 \in \mathbb{R}$, the distance between the images, $g(t_1)$ and $g(t_2)$ is increased by at most $\mathcal{D} = 1 + O(\frac{1}{\ln^2(1/\varepsilon)})$. Next we prove that the distances do not decrease too much either.

CLAIM 5.17. Given $t_1 < t_2$, we have $\|g(t_1) - g(t_2)\| \geq \frac{|t_1 - t_2|}{\mathcal{D}}$.

PROOF. The claim is trivial if both $|t_1|, |t_2| \geq 1$ or $|t_1|, |t_2| \leq \varepsilon$. Also if $t_2 \geq 1$ and $-\varepsilon \leq t_1 \leq \varepsilon$, the claim holds as $\frac{t_2 - t_1}{g(t_2) - g(t_1)} \leq \frac{t_2 + \varepsilon}{t_2 - \varepsilon} \leq \frac{1 + \varepsilon}{1 - \varepsilon} \leq 1 + 3\varepsilon$ for sufficiently small ε . Also if $\varepsilon \leq t_1 < t_2 \leq 1$, then by triangle inequality, $\|g(t_2) - g(t_1)\| \geq \|g(t_2)\| - \|g(t_1)\| = t_2 - t_1$. The remaining cases are discussed below or implied by symmetry.

Case 1. If $\varepsilon \leq t_2 \leq 1$ and $-1 \leq t_1 \leq -\varepsilon$, by symmetry we can assume that $t_2 \geq |t_1|$, and thus suppose that $t_1 = -\alpha t_2$, where $0 \leq \alpha \leq 1$. First, note that if $\alpha \leq 1/\ln^2(1/\varepsilon)$, then since the distances from the origin to the points remain unchanged, we have that

$$\begin{aligned} \frac{\|g(t_1) - g(t_2)\|}{|t_1 - t_2|} & \geq \frac{t_2 + t_1}{t_2 - t_1} \geq \frac{1 - \alpha}{1 + \alpha} \geq 1 - O(\alpha) \\ & \geq 1 - O(1/\log^2(1/\varepsilon)) \end{aligned}$$

which proves the claim. Therefore, we can assume that $\alpha \geq 1/\ln^2(1/\varepsilon)$. We should show that $\|g(t_1) - g(t_2)\|/|t_1 - t_2| \geq 1/\mathcal{D} \geq 1 - O(\frac{1}{\ln^2(1/\varepsilon)})$, or equivalently, $\|g(t_1) - g(t_2)\|^2/|t_1 - t_2|^2 \geq 1 - O(\frac{1}{\ln^2(1/\varepsilon)})$.

$$\begin{aligned}
\frac{\|g(t_1) - g(t_2)\|^2}{|t_1 - t_2|^2} &= \frac{t_1^2 + t_2^2 - 2t_1t_2 \cos(\varphi(t_1) - \varphi(t_2))}{(t_1 - t_2)^2} \\
&= 1 + \frac{2t_1t_2(1 - \cos(\varphi(t_1) - \varphi(t_2)))}{(t_1 - t_2)^2} \\
&= 1 - \frac{2t_2^2\alpha(1 - \cos(\varphi(t_1) - \varphi(t_2)))}{t_2^2(1 + \alpha)^2} \\
&= 1 - \frac{2\alpha(1 - \cos(\varphi(t_1) - \varphi(t_2)))}{(1 + \alpha)^2} \\
&= 1 - O(\alpha(1 - \cos(\varphi(t_1) - \varphi(t_2))))
\end{aligned}$$

Therefore, we just need to show that $\alpha(1 - \cos(\varphi(t_1) - \varphi(t_2))) = O(1/\ln^2(1/\varepsilon))$. Note that

$$\begin{aligned}
\varphi(t_1) - \varphi(t_2) &= \frac{\pi \ln(1/(\alpha t_2))}{\ln(1/\varepsilon)} - \frac{\pi \ln(1/t_2)}{\ln(1/\varepsilon)} = \frac{\pi \ln(1/\alpha)}{\ln(1/\varepsilon)} \\
&\leq \frac{2\pi \ln \ln(1/\varepsilon)}{\ln(1/\varepsilon)},
\end{aligned}$$

and therefore, we can use the Taylor expansion for cosine and get that

$$\begin{aligned}
\alpha(1 - \cos(\varphi(t_1) - \varphi(t_2))) &\leq \alpha \left(1 - \left[1 - \frac{\pi^2 \ln^2(1/\alpha)}{2 \ln^2(1/\varepsilon)} \right] \right) \\
&\leq O\left(\frac{\alpha \ln^2(1/\alpha)}{\ln^2(1/\varepsilon)} \right)
\end{aligned}$$

which is at most $O(1/\ln^2(1/\varepsilon))$ as $\alpha \ln^2(1/\alpha)$ is at most e for $0 \leq \alpha \leq 1$. This completes the proof for this case.

Case 2. If $t_2 \geq 1$ and $\varepsilon \leq |t_1| \leq 1$, then let us again write the term we need to bound

$$\begin{aligned}
\frac{\|g(t_2) - g(t_1)\|^2}{|t_2 - t_1|^2} &= \frac{(t_2 - t_1 \cos \varphi(t_1))^2 + t_1^2 \sin^2 \varphi(t_1)}{(t_2 - t_1)^2} \\
&= \frac{t_1^2 + t_2^2 - 2t_1t_2 \cos \varphi(t_1)}{(t_1 - t_2)^2}
\end{aligned}$$

Now if t_1 is positive, i.e., $\varepsilon \leq t_1 \leq 1$, then clearly, since $\cos \varphi(t_1) \leq 1$, we have that $-2t_1t_2 \cos \varphi(t_1) \geq -2t_1t_2$, and therefore the above fraction is at least 1. Thus, we now consider the case where $-1 \leq t_1 \leq -\varepsilon$, and need to show that $-2t_1t_2(1 - \cos \varphi(t_1))/(t_2 - t_1)^2 \leq O(1/\ln^2(1/\varepsilon))$. Again, we let $t_1 = -\alpha t_2$ where $0 < \alpha \leq 1$, and we get that

$$\frac{-2t_1t_2(1 - \cos \varphi(t_1))}{(t_2 - t_1)^2} = O(\alpha(1 - \cos \varphi(t_1)))$$

Again, if $\alpha \leq 1/\ln^2(1/\varepsilon)$, we have that $\alpha(1 - \cos \varphi(t_1)) \leq O(1/\ln^2(1/\varepsilon))$ as $(1 - \cos \varphi(t_1)) \leq 2$. Otherwise, as $t_2 \geq 1$, we have $|t_1| \geq 1/\ln^2(1/\varepsilon)$, and therefore, $\varphi(t_1) \leq \frac{2\pi \ln \ln(1/\varepsilon)}{\ln(1/\varepsilon)}$. Thus, similar to Case 1, we can write that

$$\begin{aligned}
\alpha(1 - \cos \varphi(t_1)) &\leq \alpha \left(1 - \left[1 - \frac{\pi^2 \ln^2(1/\alpha)}{2 \ln^2(1/\varepsilon)} \right] \right) \\
&= O\left(\frac{\alpha \ln^2(1/\alpha)}{\ln^2(1/\varepsilon)} \right) \\
&= O(1/\ln^2(1/\varepsilon))
\end{aligned}$$

where the above holds for similar reasons as Case 1.

Case 3. If $\varepsilon \leq t_2 \leq 1$ and $-\varepsilon \leq t_1 \leq \varepsilon$, then we have

$$\begin{aligned}
\frac{\|g(t_2) - g(t_1)\|^2}{|t_2 - t_1|^2} &= \frac{t_1^2 + t_2^2 + 2t_1t_2 \cos \varphi(t_2)}{(t_2 - t_1)^2} \\
&= 1 + \frac{2t_1t_2(1 + \cos \varphi(t_2))}{(t_2 - t_1)^2}
\end{aligned}$$

Now, if $t_1 > 0$, then the above term is at least $1 \geq 1 - O(1/\ln^2(1/\varepsilon))$ and the claim holds. So we assume that $-\varepsilon \leq t_1 \leq 0$, and let $t_1 = -\alpha t_2$ where $0 \leq \alpha \leq 1$. Our goal is to prove that $-2t_1t_2(1 + \cos \varphi(t_2))/(t_2 - t_1)^2 \leq O(1/\ln^2(1/\varepsilon))$. Again we can write

$$\frac{-2t_1t_2(1 + \cos \varphi(t_2))}{(t_2 - t_1)^2} = O(\alpha(1 + \cos \varphi(t_2)))$$

Now if $\alpha \leq 1/\ln^2(1/\varepsilon)$, we are done as $1 + \cos \varphi(t_2) \leq 2$. But then if $\alpha \geq 1/\ln^2(1/\varepsilon)$, we have that $t_2 = -t_1/\alpha \leq \varepsilon/\alpha \leq \varepsilon \ln^2(1/\varepsilon)$, and therefore,

$$\varphi(t_2) \geq \frac{\pi \ln(1/(\varepsilon \ln^2(1/\varepsilon)))}{\ln(1/\varepsilon)} = \pi - \frac{2\pi \ln \ln(1/\varepsilon)}{\ln(1/\varepsilon)}.$$

Since $\frac{2\pi \ln \ln(1/\varepsilon)}{\ln(1/\varepsilon)}$ is small, we can write the Taylor expansion and get that

$$\begin{aligned}
\alpha(1 + \cos \varphi(t_2)) &= \alpha \left(1 + \cos \left(\pi - \frac{2\pi \ln \alpha}{\ln(1/\varepsilon)} \right) \right) \\
&= \alpha \left(1 - \cos \left(\frac{2\pi \ln \alpha}{\ln(1/\varepsilon)} \right) \right) \\
&\leq \alpha \left(1 - \left[1 - \left(\frac{2\pi \ln \alpha}{\ln(1/\varepsilon)} \right)^2 \right] \right) \\
&\leq O\left(\frac{\alpha \ln^2 \alpha}{\ln^2(1/\varepsilon)} \right) \\
&= O\left(\frac{1}{\ln^2(1/\varepsilon)} \right),
\end{aligned}$$

as desired. This completes the proof of this case. $\square$

Finally, we need to prove the total movement condition as follows².

CLAIM 5.18. *For every point $t \in \mathbb{R}$, one has $\|g(t) - (t, 0)\| \leq O(1)$.*

PROOF. The claim is clearly true for $|t| \geq 1$. Also for $-\varepsilon \leq t \leq \varepsilon$, the claim holds since those points move by at most 2ε . Finally for points that are on the curve, i.e., $\varepsilon \leq |t| \leq 1$, we know that their distances to the origin is preserved. Therefore, by triangle inequality, $\|g(t) - (t, 0)\| \leq 2|t| = O(1)$. $\square$

This concludes the proof of the lemma. $\square$

COROLLARY 5.19. *Let ε be a sufficiently small constant and let $p, q, p', q' \in \mathbb{R}$, such that $p' - q' \in (1 \pm \varepsilon)\Delta$, where $\Delta = q - p$ which can be positive or negative. Denote $\alpha = p - \frac{\Delta}{\varepsilon^{2/3}}$, $\beta = p - \frac{\Delta}{\varepsilon^{1/3}}$, $\gamma = q + \frac{\Delta}{\varepsilon^{1/3}}$, $\delta = q + \frac{\Delta}{\varepsilon^{2/3}}$, $\alpha' = q' - \frac{\Delta}{\varepsilon^{2/3}}$, $\beta' = q' - \frac{\Delta}{\varepsilon^{1/3}}$, $\gamma' = q' + \frac{\Delta}{\varepsilon^{1/3}}$, $\delta' = q' + \frac{\Delta}{\varepsilon^{2/3}}$. Then, there exists a map $g: \mathbb{R} \rightarrow \mathbb{R}^2$ such that:*

- $g(\alpha) = (\alpha', 0)$; $g(\beta) = (\gamma', 0)$; $g(\gamma) = (\beta', 0)$; $g(\delta) = (\delta', 0)$.

²We remark that a stronger bound for the total movement can be achieved but for our purposes the above bound suffices.

- $\forall t_1, t_2 \in \mathbb{R}$, one has $\|g(t_1) - g(t_2)\| \in \left(1 + O\left(\frac{1}{\log^2(1/\varepsilon)}\right)\right) \cdot |t_1 - t_2|$.

PROOF. Without loss of generality, we will assume that Δ is positive. Thus we have that $\alpha \leq \beta \leq p \leq q \leq \gamma \leq \delta$, and $\alpha' \leq \beta' \leq q' \leq p' \leq \gamma' \leq \delta'$. The other case is symmetric. Let $m = \frac{p+q}{2}$ which is also equal to $\frac{\alpha+\delta}{2} = \frac{\beta+\gamma}{2}$, and let $m' = \frac{q'+p'}{2}$ which is also equal to $\frac{\alpha'+\delta'}{2} = \frac{\beta'+\gamma'}{2}$.

Let $\eta = \frac{\|p'-q'\|}{\|q-p\|}$ which clearly lies in $(1 \pm \varepsilon)$. First we define the map $h: \mathbb{R} \rightarrow \mathbb{R}$ as follows:

$$h(t) = \begin{cases} \alpha' + t - \alpha, & \text{if } t \in (-\infty; \beta], \\ \gamma' + t - \gamma, & \text{if } t \in [\gamma, \infty), \\ m' + \eta(t - m), & \text{otherwise,} \end{cases}$$

which trivially maps the points from $(-\infty; \beta]$ to $(-\infty; \beta']$, and the points from $[\gamma; \infty)$ to $[\gamma'; \infty)$ by translation, and linearly maps $[\beta, \gamma]$ to $[\beta', \gamma']$ by scaling and translating the points. It is clear that the map is continuous and its distortion is at most $\max\{\eta, 1/\eta\}$ which is at most $1 + O(\varepsilon)$.

Now let $g_0: \mathbb{R} \rightarrow \mathbb{R}^2$ be the map of Lemma 5.15 with $\varepsilon' = \frac{m'-\beta'}{m'-\alpha'} = \frac{\gamma'-m'}{\delta'-m'} = \frac{\Delta(\eta/2+1/\varepsilon^{1/3})}{\Delta(\eta/2+1/\varepsilon^{2/3})} \leq O(\varepsilon^{1/3})$ which has distortion $1 + O(1/\log^2(1/\varepsilon')) = 1 + O(1/\log^2(1/\varepsilon))$, and define the scale parameter $\lambda = (m' - \alpha') = (\delta' - m')$. Our final map is just defined as $g(t) = m' + \lambda g_0(\frac{h(t)-m'}{\lambda})$ and it is clear that its distortion $\mathcal{D}_g \leq \mathcal{D}_h \cdot \mathcal{D}_{g_0} \leq (1 + \varepsilon)(1 + O(1/\log^2(1/\varepsilon))) \leq (1 + O(1/\log^2(1/\varepsilon)))$. This proves the second property. For the first property we have the following.

- $g(\alpha)$: We have that $h(\alpha) = \alpha'$ and thus $g(\alpha) = m' + (m' - \alpha')g_0(\frac{\alpha'-m'}{m'-\alpha'}) = m' + (m' - \alpha')g_0(-1) = m' - (m' - \alpha') = \alpha'$.
- $g(\beta)$: We have that $h(\beta) = \beta'$ and thus $g(\beta) = m' + (\delta' - m')g_0(\frac{\beta'-m'}{m'-\alpha'}) = m' + (m' - \alpha')g_0(-\varepsilon') = m' + (\delta' - m')\varepsilon' = \gamma'$.
- $g(\gamma)$: We have that $h(\gamma) = \gamma'$ and thus $g(\gamma) = m' + (m' - \alpha')g_0(\frac{\gamma'-m'}{\delta'-m'}) = m' + (m' - \alpha')g_0(\varepsilon') = m' - \varepsilon'(m' - \alpha') = \beta'$.
- $g(\delta)$: We have that $h(\delta) = \delta'$ and thus $g(\delta) = m' + (\delta' - m')g_0(1) = m' + (\delta' - m') = \delta'$.

□

5.3 Lower Bound

In this section, we show that there exist maps with distortion $1 + \varepsilon$ such that every outer extension of it has distortion at least $1 + \Omega(1/\log_2(1/\varepsilon))^2$.

THEOREM 5.20 (THEOREM 1.3). *There exist $X \subset \mathbb{R}$ and a map $f: X \rightarrow \mathbb{R}$ with distortion $1 + O(\varepsilon)$ such that every outer bi-Lipschitz extension $f': \mathbb{R} \rightarrow \mathbb{R}^m$ has distortion at least $1 + \Omega(1/\log^2(1/\varepsilon))$.*

PROOF. Consider a map f that maps three points $-\varepsilon, 0$, and 1 to points $\varepsilon, 0$, and 1 , respectively. The map has distortion $\frac{1+\varepsilon}{1-\varepsilon} = 1 + 2\varepsilon + O(\varepsilon^2)$. We show that any bi-Lipschitz extension $f': [-\varepsilon, 1] \rightarrow \mathbb{R}^n$ of f has distortion at least

$$1 + \left(\frac{\pi}{2 \log_2(1/\varepsilon)}\right)^2$$

asymptotically.




Figure 5: Points $x'_0, \dots, x'_k$ and angles $\alpha_1, \dots, \alpha_k$.

Consider a bi-Lipschitz extension $f': [-\varepsilon, 1] \rightarrow \mathbb{R}^n$ of f . Without loss of generality, we assume that $\varepsilon = 1/2^k$. Let $x_i = 1/2^i$ for $i \in \{0, 1, \dots, k\}$, and $x'_i = f'(x_i)$. We will need the following claim.

CLAIM 5.21. *Consider three points a, b, c on a line such that b lies exactly in the middle between a and c ; i.e., $b = (a + c)/2$. Assume that they are mapped to points a', b', c' in $\mathbb{R}^m$. Let α be the angle between segments $[a', b']$ and $[a', c']$. Then the distortion D of the map is at least $1/\cos \alpha$ if $\alpha \leq \pi/4$ and $\sqrt{2}$, otherwise. In particular,*

$$D \geq \min(1/\cos \alpha, \sqrt{2}).$$

PROOF. First, assume that $\alpha \leq \pi/4$. We now show that $\|a' - b'\| \geq \frac{\|a' - c'\|}{2 \cos \alpha}$ or $\|b' - c'\| \geq \frac{\|a' - c'\|}{2 \cos \alpha}$. Let $\rho = \|a' - b'\|/\|a' - c'\|$. If $\rho \geq \frac{1}{2 \cos \alpha}$, we are done. Otherwise,

$$\begin{aligned} \|b' - c'\|^2 &= \|a' - b'\|^2 + \|a' - c'\|^2 \\ &\quad - 2 \cos \alpha \cdot \|a' - b'\| \|a' - c'\| \\ &= \|a' - c'\|^2 (\rho^2 - 2 \cos \alpha \cdot \rho + 1). \end{aligned}$$

Now, the polynomial $t^2 - 2 \cos \alpha \cdot t + 1$ attains its minimum on $[0, 1/(2 \cos \alpha)]$ at point $t = 1/(2 \cos \alpha)$, where it equals $1/(2 \cos \alpha)^2$ (here we use that $\alpha \leq \pi/4$ and hence $1/(2 \cos \alpha) < \cos \alpha$). Therefore, $\|b' - c'\| \geq \|a' - c'\|/(2 \cos \alpha)$, as required. Note that the distortion is at least

$$\frac{\|a' - b'\|}{\|a' - c'\|} \Big/ \frac{\|a - b\|}{\|a - c\|} \quad \text{and} \quad \frac{\|b' - c'\|}{\|a' - c'\|} \Big/ \frac{\|b - c\|}{\|a - c\|}.$$

One of these two ratios is at least $1/\cos \alpha$.

Now, assume that $\alpha \in (\pi/4, \pi/2)$. The distance from c' to the line passing through a' and b' is $\sin \alpha \|a' - c'\| \geq \|a' - c'\|/\sqrt{2}$; in particular, $\|b' - c'\| \geq \|a' - c'\|/\sqrt{2}$. As in the previous case, this implies that the distortion is at least $\sqrt{2}$. Finally, assume that $\alpha \geq \pi/2$, then the angle at vertex a' in the triangle $a'b'c'$ is obtuse, therefore $b'c'$ is the longest side of $a'b'c'$. In particular, $\|b' - c'\| \geq \|a' - c'\|$. We get that the distortion is at least 2. □

Now we are ready to prove Theorem 5.20. Let α_i be the angle between segments $[0, x'_{i-1}]$ and $[0, x'_i]$ (see Figure 5). Consider point $z = (\varepsilon, \bar{0}) = f'(-\varepsilon)$. Let β be the largest among the following angles:

- the angle between $[x'_k, z]$ and $[x'_k, 0]$,
- the angle between $[z, x'_k]$ and $[z, 0]$.

Finally, let γ be the angle between $[0, x'_k]$ and $[0, z]$.

First, we apply Claim 5.21 to points $0, x_i, x_{i-1}$. We get that

$$D \geq \min\left(\frac{1}{\cos \alpha_i}, \sqrt{2}\right).$$

Second, we apply Claim 5.21 to points x_k , 0 , $-\varepsilon$ and to $-\varepsilon$, 0 , x_k (see Figure 6). We get that

$$D \geq \min\left(\frac{1}{\cos \beta}, \sqrt{2}\right).$$

Now, we write an upper bound for γ (which follows from the




Figure 6: Points $-\varepsilon, 0, x_k$ and their images $z = f'(-\varepsilon), 0 = f'(0), x'_k = f'(x_k)$.

triangle inequality in spherical geometry)

$$\gamma \leq \sum_{i=1}^k \alpha_i.$$

Consider the triangle with vertices $0, z, x'_k$. One of the angles of this triangle is γ and the largest of the other two angles is β . Therefore, $\gamma + 2\beta \geq \pi$ and thus,

$$2\beta + \sum_{i=1}^k \alpha_i \geq \pi.$$

Consequently, either $\beta \geq \pi/(k+2)$ or some $\alpha_i \geq \pi/(k+2)$ (or both). We conclude that the distortion is at least

$$D \geq \min\left(\frac{1}{\cos \frac{\pi}{\log_2(1/\varepsilon)+2}}, \sqrt{2}\right) = 1 + (1 - o(1)) \frac{\pi^2}{2 \log_2^2(1/\varepsilon)}$$

when $\varepsilon \rightarrow 0$. $\square$

REFERENCES

- [1] Ittai Abraham, Yair Bartal, T-H. Hubert Chan, Kedar Dhamdhere, Anupam Gupta, Jon Kleinberg, Ofer Neiman, and Aleksandrs Slivkins. 2005. Metric embeddings with relaxed guarantees. In *Proceedings of the Foundations of Computer Science*. 83–100.
- [2] Ittai Abraham, Yair Bartal, and Ofer Neiman. 2006. Advances in metric embedding theory. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*. ACM, 271–286.
- [3] Dimitris Achlioptas. 2003. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *J. Comput. System Sci.* 66, 4 (2003), 671–687.
- [4] Arseniy V. Akopyan and Aleksey S. Tarasov. 2008. A constructive proof of Kirschbraun's theorem. *Mathematical Notes* 84, 5-6 (2008), 725–728.
- [5] P. Alestalo and D. A. Trotsenko. 2009. Plane sets allowing bilipschitz extensions. *Math. Scand.* (2009), 134–146.
- [6] P. Alestalo, D. A. Trotsenko, and J. Väisälä. 2003. Linear bilipschitz extension property. *Siberian Mathematical Journal* 44, 6 (2003), 959–968.
- [7] Noga Alon. 2009. Perturbed identity matrices have high rank: Proof and applications. *Combinatorics, Probability and Computing* 18, 1-2 (2009), 3–15.
- [8] Noga Alon and Bo'az Klartag. 2017. Optimal compression of approximate inner products and dimension reduction. In *Proceedings of the Symposium on Foundations of Computer Science*.
- [9] Alexandr Andoni, Piotr Indyk, and Robert Krauthgamer. 2008. Earth Mover Distance over High-Dimensional Spaces. In *Proceedings of the Symposium on Discrete Algorithms*. 343–352.
- [10] Alexandr Andoni, Piotr Indyk, and Robert Krauthgamer. 2009. Overcoming the ℓ_1 Non-Embeddability Barrier: Algorithms for Product Metrics. In *Proceedings of the Symposium on Discrete Algorithms*. 865–874.
- [11] Alexandr Andoni, Robert Krauthgamer, and Ilya Razenshteyn. 2015. Sketching and Embedding are Equivalent for Norms. In *Proceedings of the Symposium on Theory of Computing*. 479–488.
- [12] Alexandr Andoni, Huy L. Nguyen, Aleksandar Nikolov, Ilya Razenshteyn, and Erik Waingarten. 2017. Approximate Near Neighbors for General Symmetric Norms. In *Proceedings of the 49th ACM Symposium on the Theory of Computing (STOC '17)*. Available as arXiv:1611.06222.
- [13] Sanjeev Arora, James Lee, and Assaf Naor. 2008. Euclidean distortion and the sparsest cut. *Journal of the American Mathematical Society* 21, 1 (2008), 1–21.
- [14] Yonatan Aumann and Yuval Rabani. 1998. An $O(\log k)$ Approximate Min-Cut Max-Flow Theorem and Approximation Algorithm. *SIAM J. Comput.* 27, 1 (1998), 291–301.
- [15] Yair Bartal. 1998. On approximating arbitrary metrics by tree metrics. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*. ACM, 161–168.
- [16] Yair Bartal, Béla Bollobás, and Manor Mendel. 2006. Ramsey-type theorems for metric spaces with applications to online problems. *J. Comput. System Sci.* 72, 5 (2006), 890–921.
- [17] Tugkan Batu, Funda Ergun, and Cenk Sahinalp. 2006. Oblivious String Embeddings and Edit Distance Approximation. In *Proceedings of the Symposium on Discrete Algorithms*. 792–801.
- [18] Jean Bourgain. 1985. On lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics* 52, 1-2 (1985), 46–52.
- [19] Moses Charikar and Robert Krauthgamer. 2006. Embedding the Ulam Metric into ℓ_1 . *Theory of Computing* 2, 11 (2006), 207–224.
- [20] Eden Chlamtac, Konstantin Makarychev, and Yuri Makarychev. 2006. How to Play Unique Games Using Embeddings. In *Proceedings of the Symposium on Foundations of Computer Science*. 687–696.
- [21] Graham Cormode, S. Muthukrishnan, and Cenk Sahinalp. 2001. Permutation Editing and Matching via Embeddings. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*. 481–492.
- [22] Michael Elkin, Arnold Filtser, and Ofer Neiman. 2015. Prioritized Metric Structures and Embedding. In *Proceedings of the Symposium on Theory of Computing*. 489–498.
- [23] Michael Elkin, Arnold Filtser, and Ofer Neiman. 2017. Terminal embeddings. *Theoretical Computer Science* 697 (2017), 1–36.
- [24] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. 2008. *Approximating Metric Spaces by Tree Metrics*. 1–99.
- [25] Uriel Feige. 1998. Approximating the bandwidth via volume respecting embeddings. In *Proceedings of the Symposium on Theory of Computing*. ACM, 90–99.
- [26] Manouchehr Ghamsari. 1993. Sobolev and Quasiconformal Extension Domains. *Proc. Amer. Math. Soc.* 119, 4 (1993), 1179–1188.
- [27] Piotr Indyk and Jiri Matousek. 2004. Low-Distortion Embeddings of Finite Metric Spaces. In *Handbook of Discrete and Computational Geometry*. Citeseer.
- [28] Piotr Indyk and Rajeev Motwani. 1998. Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality. In *Proceedings of the Symposium on Theory of Computing*. 604–613.
- [29] William Johnson and Joram Lindenstrauss. 1984. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability, New Haven, Connecticut*. Contemporary Mathematics, Vol. 26. 189–206.
- [30] Subhash Khot and Nisheeth Vishnoi. 2015. The Unique Games Conjecture, Integrality Gap for Cut Problems and Embeddability of Negative-Type Metrics into ℓ_1 . *J. ACM* 62:1, 8 (2015), 8:1–8:39.
- [31] Mojżesz D. Kirschbraun. 1934. Über die zusammenziehende und Lipschitzsche Transformationen. *Fundamenta Mathematicae* 22 (1934), 77–108.
- [32] Leonid V. Kovalev. 2017. Symmetrization and extension of planar bi-Lipschitz maps. *arXiv preprint arXiv:1706.00102* (2017).
- [33] Kasper G. Larsen and Jelani Nelson. 2016. The Johnson-Lindenstrauss lemma is optimal for linear dimensionality reduction. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*.
- [34] Kasper G. Larsen and Jelani Nelson. 2017. Optimality of the Johnson-Lindenstrauss lemma. In *Proceedings of the Symposium on Foundations of Computer Science*.
- [35] James R. Lee and Assaf Naor. 2005. Extending Lipschitz Functions via Random Metric Partitions. *Inventiones mathematicae* 160, 1 (2005), 59–95.
- [36] Nathan Linial, Eran London, and Yuri Rabinovich. 1995. The Geometry of Graphs and Some of its Algorithmic Applications. *Combinatorica* 15, 2 (1995), 215–245.
- [37] Konstantin Makarychev and Yuri Makarychev. 2016. Metric Extension Operators, Vertex Sparsifiers and Lipschitz Extendability. *Israel Journal of Mathematics* 212, 2 (2016), 913–959.
- [38] Konstantin Makarychev and Yuri Makarychev. 2016. A Union of Euclidean Metric Spaces is Euclidean. *Discrete Analysis* 14 (2016).
- [39] M. B. Marcus and G. Pisier. 1984. Characterizations of Almost Surely Continuous p -stable Random Fourier Series and Strongly Stationary Processes. *Acta mathematica* 152, 1 (1984), 245–301.
- [40] Jiří Matoušek. 2002. *Lectures on Discrete Geometry*. Springer.
- [41] Edward James McShane. 1934. Extension of Range of Functions. *Bull. Amer. Math. Soc.* 40, 12 (1934), 837–842.
- [42] Manor Mendel and Assaf Naor. 2006. Some applications of Balls extension theorem. *Proc. Amer. Math. Soc.* 134, 9 (2006), 2577–2584.

- [43] Assaf Naor, Yuval Peres, Oded Schramm, and Scott Sheffield. 2006. Markov chains in smooth Banach spaces and Gromov-hyperbolic metric spaces. *Duke Mathematical Journal* 134, 1 (2006), 165–197.
- [44] Assaf Naor and Gideon Schechtman. 2007. Planar Earthmover is not in L_1 . *SIAM J. Comput.* 37, 3 (2007), 804–826.
- [45] Huy L. Nguyễn. 2014. *Algorithms for High Dimensional Data*. Ph.D. Dissertation. Princeton University.
- [46] Rafail Ostrovsky and Yuval Rabani. 2007. Low Distortion Embedding for Edit Distance. *J. ACM* 54, 5 (2007), 23:1–23:16.
- [47] Juha Partanen and Jussi Väisälä. 1993. EXTENSION OF BILIPSCHITZ MAPS OF COMPACT POLYHEDRA. *Math. Scand.* 72, 2 (1993), 235–264.
- [48] Jussi Väisälä, Matti Vuorinen, and Hans Wallin. 1994. Thick sets and quasymmetric maps. *Nagoya Math. J.* 135 (1994), 121–148.
- [49] John von Neumann. 1928. Zur Theorie der Gesellschaftsspiele. *Math. Ann.* 100, 1 (1928), 295–320.