

Dimension-Accuracy Tradeoffs in Contrastive Embeddings for Triplets, Terminals & Top- k Nearest Neighbors

Vaggos Chatziafratis* Piotr Indyk†

Abstract

Metric embeddings traditionally study how to map n items to a target metric space such that distance lengths are not heavily distorted; but what if we only care to preserve the *relative order of the distances (and not their length)*? In this paper, we are motivated by the following basic question: given triplet comparisons of the form “item i is closer to item j than to item k ,” can we find low-dimensional Euclidean representations for the n items that respect those distance comparisons? Such order-preserving embeddings naturally arise in important applications—recommendations, ranking, crowdsourcing, psychometrics, nearest-neighbor search—and have been studied since the 1950s, under the name of *ordinal* or *non-metric* embeddings.

Our main results are:

- **Nearly-Tight Bounds on Triplet Dimension:** We introduce the natural concept of *triplet dimension* of a dataset, and surprisingly, we show that in order for an ordinal embedding to be triplet-preserving, its dimension needs to grow as $\frac{n}{2}$ in the worst case. This is optimal (up to constant) as $n - 1$ dimensions always suffice.
- **Tradeoffs for Dimension vs (Ordinal) Relaxation:** We then relax the requirement that every triplet should be exactly preserved and present almost tight lower bounds for the maximum ratio between distances whose relative order was inverted by the embedding; this ratio is known as (ordinal) *relaxation* in the literature and serves as a counterpart to (metric) distortion.
- **New Bounds on Terminal and Top- k -NNs Embeddings:** Going beyond triplets, we then study two well-motivated scenarios where we care about preserving specific sets of distances (not necessarily triplets). The first scenario is *Terminal Ordinal Embeddings* where we want to preserve relative distance orders to k given items (the “terminals”), and for that we present matching upper and lower bounds. The second scenario is *top- k -NNs Ordinal Embeddings*, where for each item we want to preserve the relative order of its k nearest neighbors, for which we present lower bounds.

To the best of our knowledge, these are some of the first tradeoffs on triplet-preserving ordinal embeddings and the first study of Terminal and Top- k -NNs Ordinal Embeddings.

*UC Santa Cruz. Part of this work was done while being supported by a FODSI postdoc fellowship at MIT and Northeastern.
†MIT.

1 Introduction

Given n items of interest endowed with some abstract notion of “distance” (not necessarily a metric), we often wish to represent them as a configuration of n points in some convenient target metric space, commonly a low-dimensional Euclidean space or a tree metric. Having such representations has been proven crucial for speeding up computation, reducing memory needs and has led to deep algorithmic and mathematical insights.

Such representations lie at the heart of many applications including nearest-neighbor searching, ad placement, recommendation systems, crowdsourcing, social networks, clustering, visualization and even psychometrics. Since distances encode interesting information about a dataset, the metric embeddings literature has studied methods to preserve those pairwise distances (either exactly or with distortion), and has yielded various tradeoffs between the faithfulness of the embedding and its dimensionality ([Mat13]). Our work is motivated by the following two observations:

- First, notice that many of the aforementioned applications *do not rely on the distances per se, but rather they rely on the relative order of those distances*. For example, in recommendation systems or online ad placements, the ranking of which n items to show is what matters, whereas pairwise distance lengths are of little importance. Moreover, at the heart of many of these applications is the fundamental problem of nearest-neighbor search ([AI08])—asking for the closest point from a dataset to a given query point q —which is intrinsically a question about relative orderings, rather than absolute distances.
- Second, the n items of interest may lie in an abstract space where even the notion of pairwise distance may be severely underspecified and hard to evaluate. This is especially common in psychometrics ([Tor52, Thu54, Kru64a, Kru64b]), where humans are asked to answer queries about their feelings, preferences etc., and in crowdsourcing marketplaces (e.g., Mechanical Turk) where “workers” are paid to provide responses to a series of questions. Because humans are surprisingly bad and inconsistent at answering *cardinal* questions (how *much* did you like this movie or this restaurant?), yet very fast and accurate at answering *ordinal* questions (did you enjoy A *more* than B ?), the deployed queries are usually (paired) comparisons ([Thu54]) between items, and as such they only provide an indirect access to some underlying notion of “distance” for the items.

Computational Task. The natural question that arises in the scenarios described above is whether there exist low-dimensional Euclidean representations that preserve the *relative ordering* of distances (first observation), given perhaps *incomplete* information about how underlying distances are related (second observation). More concretely, we are interested in the following basic question:

Given a collection of (triplet) comparisons of the form “ A is more similar to B than to C ”, are there low-dimensional embeddings that respect the relative order of distances (not their length)?

Our Contribution. Our main contribution is to derive several tradeoffs that arise between the faithfulness of the order-preserving embedding and its dimensionality. For the case where the embedding needs to respect all of the (triplet) comparisons (Sec. 3.1), we give nearly-tight bounds for the dimension. For the case where we allow some distances to be inverted (Sec. 3.2), we present almost tight lower bounds for the “ordinal relaxation” of the embedding (this is the analogue of the notion of “distortion” in metric embeddings, see definition below). Finally, we also study two new settings motivated by terminal embeddings (Sec. 4) and by top- k -NNs preservation (Sec. 5).

1.1 Related Work The aforementioned order-preserving embeddings are usually referred to in the literature as *ordinal*, or *non-metric* embeddings, or *monotone* maps. In this paper, the terms “contrastive embeddings” and “ordinal embeddings” are used interchangeably to refer to the goal of finding embeddings that preserve the relative order of distances (not their exact length); having access to such contrastive information is popular in contrastive learning ([SE05, SAG⁺22]).

We give here the definition:

DEFINITION 1.1. (ORDINAL EMBEDDING, ORDINAL DIMENSION) *Let $X = ([n], \delta)$ be any metric space on n points, and let $\|\cdot\|$ be a norm on \mathbf{R}^d . We say that $\phi : X \rightarrow (\mathbf{R}^d, \|\cdot\|)$ is an ordinal embedding if for every $x, y, z, w \in X$, we have the following: $\delta(x, y) < \delta(w, z) \iff \|\phi(x) - \phi(y)\| < \|\phi(w) - \phi(z)\|$. Moreover, the*

ordinal dimension of X is defined to be the smallest dimension of a Euclidean space into which X can be ordinally embedded.

The study on ordinal embeddings dates back to the early days of Multi-Dimensional Scaling (MDS) with some of the seminal works of the field in the 1950s and 1960s. At the time, MDS was heavily used in (mathematical) psychology and psychometrics, where the works of [Tor52], [She62, She74, CS74], [Kru64a, Kru64b] laid the foundations of many concepts in the field; these works attempted to formalize explicitly, the implicit connections between human-perceived similarities and differences among items (e.g., faces, tastes, odors etc.) via suitable data representations, and to propose methods for finding such representations obeying subjective descriptions.

In this early context from psychometrics, our basic question above for triplet comparisons asking “is item A more similar to B or to C ” would be equivalent to asking whether “ $\delta(A, B) < \delta(A, C)$ ”, which is an ordinal question that is easy for humans to answer (rather than exactly specifying values for the length of the distances). Later works made important steps towards understanding tradeoffs in ordinal embeddings, and the three most related to our work is the paper by [BL05] and by [ABD⁺08, BDH⁺08], who considered several special cases, where we have access to the *complete* set of all $\binom{n}{2}$ distance comparisons (i.e., $\delta(i, j) \leq \delta(k, l), \forall i, j, k, l \in X$) ([BL05, ABD⁺08]), or the case when the source and target metrics come from simple families: in [BDH⁺08] they provide approximation algorithms to embed unweighted graphs into a line metric and into a tree metric, and they also consider the embedding of unweighted trees into the line. For the case of embedding onto a line or a hierarchical tree (ultrametric), where the triplet comparisons are seen as a set of constraints with the goal of maximizing the number of constraints satisfied by the embedding, there are tight hardness of approximation results ([GMR08, CM23]). Moreover, several statistical properties and sample complexity bounds based on queries about distances on four points i, j, k, l have been studied in multiple works [AWC⁺07, TL14, KL14, GPvL19, GCY19].

An important notion for ordinal embeddings introduced by [ABD⁺08] is the notion of *ordinal relaxation*, which serves as the counterpart of the well-studied notion of distortion of metric embeddings (see e.g., [IMS17]):

DEFINITION 1.2. (ORDINAL RELAXATION) *Given an ordinal embedding $\phi : (X, \delta) \rightarrow (Y, \delta')$, we say it has multiplicative (ordinal) relaxation $\alpha \geq 1$, if $\alpha\delta(i, j) < \delta(k, l) \implies \delta'(i, j) < \delta'(k, l)$.*

To put it simply, significantly different distances (in the original space) should have their relative order preserved by the embedding,¹ or equivalently, relaxation is the maximum ratio between two distances whose relative order got inverted by the embedding. Minimum-relaxation ordinal embeddings were originally studied in [ABD⁺08] and [BDH⁺08], who established that ordinal embeddings have important differences from metric embeddings. They developed several approximation algorithms for ordinal embeddings on a line (1-dimensional Euclidean space) or on a tree, for various interesting cases such as source metrics induced by ultrametrics, or by the shortest-path metric of unweighted trees and unweighted graphs. Notice that by definition, for any source and target metrics, the optimal relaxation is at most the optimal distortion ([ABD⁺08]). Moreover, for any n -point metric space, observe that a general $O(\log n)$ upper bound on relaxation into $O(\log n)$ -dimensional Euclidean space follows easily by Bourgain’s metric embedding theorem ([Bou85]) coupled with the Johnson-Lindenstrauss lemma ([JLS86]). Another well-known fact is that n -dimensional Euclidean space suffices to ordinally embed any X , i.e., the ordinal dimension for any metric space X on n points is at most n (to be exact, $n - 1$ dimensions always suffice) ([BL05, ABD⁺08]).

1.2 Motivating Questions Despite their long history and the abovementioned general results, ordinal embeddings —both exact, or with relaxation— are not well-understood. It is easy to see that if we care about exact preservation of the distance ordering, then the Johnson-Lindenstrauss lemma is of no use in this case: the error parameter ϵ would have to be tiny as it needs to scale with the smallest distance gap (which could be $O(\frac{1}{n})$), essentially yielding no dimension reduction at all (recall, the final dimension would be $O(\frac{\log n}{\epsilon^2})$). Given this failure, we set out to address several basic questions on low-dimensional ordinal embeddings:

- **Q1** (Triplets, Exact & Relaxed): In many applications, we do not have access to the complete set of $\binom{n}{2}$ comparisons for all distances $\delta(i, j) \leq \delta(k, l)$ (as assumed in [BL05, ABD⁺08, BDH⁺08]), rather we are

¹As pointed out in the original work [ABD⁺08], in ordinal embeddings we want to respect distance equality, but in an ordinal embedding with relaxation 1, we may break ties.

given information on *triplets* of items of the form “item i is closer to item j than to item k ” or equivalently, $\delta(i, j) < \delta(i, k)$ for some abstract distance δ . Are there Euclidean representations with low dimension that respect the relative ordering on triplets, either exactly or with relaxation?

- **Q2** (k Terminals, Upper & Lower Bound): Given a set of k special items $T = \{t_1, \dots, t_k\}$ (the “terminals”), we want to find Euclidean representations so as to preserve all distance comparisons from the perspective of each terminal to the rest of the items. Are there algorithms that find such representations? What is the minimum dimension needed?
- **Q3** (Top- k -NNs, With & Without Mixed Comparisons): What if we only care to preserve the relative distances from each item i to its set $\text{NN}(i)$ containing the k nearest neighbors in the dataset, i.e., comparisons of the form $\delta(i, i'), \delta(i, i'')$ for $i', i'' \in \text{NN}(i)$? Is the minimum dimension needed different, if we also cared to preserve the distance ordering for mixed comparisons $\delta(i, i')$ vs $\delta(j, j')$ for $i' \in \text{NN}(i), j' \in \text{NN}(j), i \neq j$?
- **Q4** (Different Regimes for k): Is the answer to the above questions different for various regimes of the parameter k ? Is there a difference in the upper or lower bounds when $k = \Theta(n)$, or $k = o(n)$, or even constant independent of n (number of items)?

We believe that such questions are natural steps towards a better understanding of ordinal embeddings; they also are well-motivated from practical considerations, as we discuss in Sec. 1.5.

1.3 Our Results We provide several results for each of the questions listed above. Our work was primarily inspired by the lack of theoretical bounds for the basic problem of preserving triplets either exactly or approximately (Q1), and also by the lack of results for terminal and top- k -NNs ordinal embeddings (Q2 and Q3), in stark contrast to their counterparts in metric embeddings. Our main results can be summarized informally as follows:

- (R1) For preserving the order of all triplets exactly, we introduce the concept of *triplet dimension*, i.e., the minimum dimension needed by an ordinal embedding that respects all triplet comparisons. This is the natural analogue of ordinal dimension from Definition 1.1 but specialized for triplets (see also [RRŠ89] for other geometric notions of dimension). Perhaps surprisingly, we prove that the triplet dimension can grow linearly in n and may need to be at least $\frac{n}{2}$ (Theorem 3.1); this is tight because it is a folklore result that $n - 1$ dimensions always suffice. In the case where we allow for some of the triplet orderings to be inverted, we prove a lower bound on the relaxation of $\Omega\left(\frac{\log n}{\log d + \log \log n}\right)$ for any dimension d and n (Theorem 3.2). This is nearly tight (up to $\log \log n$ factors) because using Bourgain’s theorem and the JL lemma, we could obtain relaxation $O(\log n)$ for any metric into $O(\log n)$ -dimensional Euclidean space.
- (R2) For k terminals (k not necessarily fixed), we present a simple, yet optimal algorithm that allows us to embed the dataset into k -dimensional space such that for each of the k terminals, all relative distance orders to the remaining $n - k$ points are exactly preserved (Theorem 4.1). We complement this with tight lower bounds (Theorems 4.2 and 4.3) showing that $\Omega(k)$ dimensions are indeed necessary.
- (R3) For preserving the top- k -NNs of each point, we present an $\Omega(k)$ lower bound on the dimension (Theorem 5.1). As we will see, this is only slightly affected by whether or not we preserve the distance ordering for mixed comparisons $\delta(i, i')$ vs $\delta(j, j')$ for $i' \in \text{NN}(i), j' \in \text{NN}(j), i \neq j$.
- (R4) We show how the value of k actually affects the dimension needed for ordinal embeddings. Specifically, there is a difference for the lower bounds we can obtain for $k = \Theta(n)$ and for $k = o(n)$.

To the best of our knowledge, these are some of the first results for triplet, terminal, and top- k -NNs ordinal embeddings. In addition, our results are often tight, they extend prior works ([BL05, ABD⁺08, BDH⁺08]) and complement many of the empirical works for triplets ([SJ03, TLB⁺11, JN11, VDMW12, JJF⁺15, JJN16, KvL17, VHL⁺19, GCY19]) or other types of ordinal embeddings ([AWC⁺07, TL14, KL14, GPvL19]).

1.4 Our Techniques We build upon and extend the tools used in [BL05, ABD⁺08, BDH⁺08] who handled the complete case where all $\binom{n}{2}$ distance comparisons were available. Here, in order to handle relaxed ordinal embeddings for non-complete inputs, we have to come up with more elaborate constructions (see Theorem 3.2). Our results use previous constructions of dense high-girth graphs ([ES63, Sau70]) and rely on sampling edges from those graphs, and on several counting arguments depending on which case we deal with.

More specifically, we present three types of results: lower bounds for triplets with relaxation, lower bounds for various types of exact ordinal embeddings (triplet, terminal, top- k -NNs), and (optimal) upper bounds for k terminal embeddings.

The intuition behind our proofs for getting the relaxation lower bounds for triplets relies on finding a large number of metric spaces with significantly different behavior on triplet distances, at least from the perspective of a single vertex v . To find such metric spaces, we start from a high girth graph G (unweighted) and consider an appropriate number of subgraphs G' , chosen at random. Notice that if an edge (u, v) was present in G , but after sampling, it is not included in G' , then the distance u, v in G' suddenly becomes large (as much as the girth). Our derived tradeoffs for ordinal embeddings essentially stem from the known constructions of high girth graphs and any improvement on the latter yields improvement on our lower bounds.

Regarding exact ordinal embeddings, at their heart, our proofs are based on various counting arguments. We compare the total number of different orderings for the distances defined on triples of points against the total number of ways an ordinal embedding can embed the n points in d dimensions such that a given ordering on all triplets is satisfied. Similar arguments go through for the other types of embeddings we consider. We then observe that when the dimension d is “small” compared to n , there will be significantly more triplet orderings than distinct embeddings, which will give us the lower bound. A similar approach was carried out in [BL05] for complete instances where all distance comparisons $\delta(i, j) \leq \delta(k, l)$ were given as input.

Finally, for terminal ordinal embeddings we give a direct construction on where to embed the k terminals and the rest of the points such that distance orders from the rest of the points to the terminals are exactly preserved. We use k dimensions and this matches our lower bound for terminal ordinal embeddings.

1.5 Further Motivation and Related Work Regarding triplet-preserving ordinal embeddings (Q1), obtaining information about a dataset based on triplet comparisons like “which of j and k is closer to i ” is used in crowdsourcing and online platforms, to elicit user preferences and to perform downstream tasks such as clustering and nearest neighbor search. Important works, both theoretical and empirical, that focus on triplet embeddings under various settings include [SJ03, TLB⁺11, JN11, VDMW12, JJF⁺15, UDH15, JJN16, KvL17, KVH16, AU15, LHvL19, VHL⁺19, GCY19, FCM⁺20, HWvL20]. Triplet feedback of the form “ i is closer to j than to k ” is known to be much more reliable than absolute comparisons and much easier for humans to answer. For example, the question “are cats similar to tigers?” might yield conflicting answers, but the triplet query “are cats more similar to tigers or to dolphins?” is easier as humans have to pick the “odd-one-out” among the 3 alternatives; in addition, such triplet queries are also useful in the context of computational biology/phylogenetics for inferring ancestry relations in hierarchical clustering ([BGJ10, VD16, CNC18, EZK18, CMA21, CM23]), and in clustering via hyperbolic embeddings ([MZS⁺19, CGCR20]).

Regarding ordinal embeddings for terminals (Q2), these are useful in a scenario where we only care to preserve orders for a few k points to the rest of the $(n - k)$ items; this could arise for example, in a facility location or networking application, whenever we have a network comprising many clients and only k servers, and we want to have a simple data structure preserving the client-to-server service times (captured by the distance), but we do not care about client-to-client distance preservation. This is the analogous notion to metric embeddings with terminals that was introduced by [EFN17] and later studied in [MMMR18, NN19, CN22].

Regarding top- k -NNs ordinal embeddings (Q3), in many important applications of embeddings, preserving distance information about nearby points is much more important than preserving all distances. Indeed, it may be good enough to strictly maintain the order of the top- k nearest points, and for far away objects to just label them as “far”. In such scenarios it is natural to seek local embeddings that maintain only distances of close by neighbors. This has natural applications in ranking, search, and recommendations where often the few top results are viewed, and has been studied extensively in the metric embeddings literature, both in practice ([BN03, XSB06]) and in theory under the name of *local* (metric) embeddings ([ABN07, IN07]) or local versions of dimension reduction ([SS09]).

Finally, regarding the behaviour of the above questions as we vary k (Q4), we believe it is crucial to understand

the various tradeoffs between the parameter k , the number of points n , and the target space dimension d , as various applications may need different value ranges for the parameters.

2 Preliminaries

Let $[n]$ denote the set $\{1, 2, \dots, n\}$. Let $X = ([n], \delta)$ be a metric space on n points. Throughout our work, norm $\|\cdot\|$ is the standard Euclidean ℓ_2 norm, unless otherwise noted.

DEFINITION 2.1. (Triplet ORDINAL EMBEDDING, Triplet Dimension) *Let $\|\cdot\|$ be a norm on \mathbf{R}^d . We say that $\phi : X \rightarrow (\mathbf{R}^d, \|\cdot\|)$ is a triplet ordinal embedding if for every $x, y, z \in X$, we have the following: $\delta(x, y) < \delta(x, z) \iff \|\phi(x) - \phi(y)\| < \|\phi(x) - \phi(z)\|$. Moreover, the triplet dimension of X is defined to be the smallest dimension d of a triplet ordinal embedding of X . We sometimes say the embedding is triplet-preserving, or that it preserves all triplet orders.*

DEFINITION 2.2. (Terminal ORDINAL EMBEDDING) *Given $X = (V, \delta)$ with $|V| = n$ and a subset $T = \{t_1, \dots, t_k\} \subseteq V$ of k distinguished elements, which we call terminals, we say that $\phi : (X, T) \rightarrow (\mathbf{R}^d, \|\cdot\|)$ is a terminal ordinal embedding if it preserves the distance orders from the rest of the points (the clients in our previous example) to the terminals (the servers), i.e., if for every $t, t' \in T$ and $x, x' \in V \setminus T$, we have the following: $\delta(t, x) < \delta(t', x') \iff \|\phi(t) - \phi(x)\| < \|\phi(t') - \phi(x')\|$. The case $t = t'$ is also included.*

DEFINITION 2.3. (Top- k -NNs ORDINAL EMBEDDING) *Given X as above, define $\text{NN}(i)$ for $i \in X$ to be the set of $k \leq n - 1$ nearest neighbors of i according to δ . We say that $\phi : X \rightarrow (\mathbf{R}^d, \|\cdot\|)$ is a top- k -NNs ordinal embedding if it preserves the distance orders among all elements in $\text{NN}(i), \forall i$, i.e., if for every $i, j \in X$ and $i' \in \text{NN}(i), j' \in \text{NN}(j)$, we have the following: $\delta(i, i') < \delta(j, j') \iff \|\phi(i) - \phi(i')\| < \|\phi(j) - \phi(j')\|$. If in addition, we required $j \equiv i$ in the previous sentence, then we would get the problem of top- k -NNs ordinal embedding without mixed comparisons.*

3 Preserving Order on Triplets and the Triplet Dimension

In this section, we prove the first main result about triplet-preserving ordinal embeddings whose goal is to respect all triplet comparisons for the distances in the original space. Then, we show a tradeoff for the dimension vs the relaxation, analogous to Bourgain's embedding theorem providing a tradeoff for the dimension vs the distortion.

3.1 Lower Bound for Exact Triplet Preservation

THEOREM 3.1. *For every constant $\kappa > 0$, and for every large enough n (size of the dataset), no d -dimensional embedding in ℓ_2 can be a triplet ordinal embedding (see Definition 2.1), unless its dimension grows linearly as $d > \frac{n}{2+\kappa}$. (i.e., dimension must be roughly $n/2$ to preserve all triplets)*

Proof. First, recall the superfactorial function $G(z)$ as defined on integers:

$$G(n) = \begin{cases} 0 & \text{if } n = 0, -1, -2, \dots \\ \prod_{i=0}^{n-2} i! & \text{if } n = 1, 2, \dots \end{cases}$$

This is Barne's special G -function which is related to the gamma function as $G(z+1) = \Gamma(z)G(z)$, with $G(1) = 1$. Here we care about its asymptotic growth rate for integer n :

FACT 3.1. $G(n+1) = 0!1!\cdots(n-1)!$ and $\log G(n+1) = \frac{n^2}{2} \log n + o\left(\frac{n^2}{2} \log n\right) \approx \frac{n^2}{2} \log n$. (The symbol $f(n) \approx g(n)$ simply means that the two quantities have asymptotically the exact same behavior with the exact same leading constant in front of the dominant term, where we can ignore lower order terms.)

Observe that the number of distinct triplet-orderings on $n \geq 3$ items is at least as large as $G(n+1)$, since the first item can have any of the $(n-1)!$ permutations for its neighbors, the second item any of the $(n-2)!$ on its neighbors (excluding the first item) and so on. This is a good enough bound to prove the lower bound for the dimension as stated in the theorem.²

²We can compute the exact number of triplet orders, but the final lower bound for the dimension is asymptotically the same. The distinct triplet orders are $\frac{(n-1)!}{0!} \cdot \frac{n!}{2!} \cdot \frac{(n+1)!}{4!} \cdot \frac{(n+2)!}{6!} \cdots \frac{(2n-3)!}{(2n-4)!}$.

Triplet-Preserving Embeddings as Low-degree Polynomials. The second ingredient in the proof is to associate triplet-preserving embeddings with polynomials of degree 2 and count how many different configurations could exist for resolving distance comparisons among $\delta(i, j)$ and $\delta(i, k)$. The embedding assigns d coordinates to each of the n elements so we can define the $n \times d$ matrix, say X , just by stacking row-by-row the coordinates of each of the n elements; denote with x_i , the i^{th} row of this matrix. We say that the matrix X “realizes the triplet order of δ ” if the euclidean distances between pairs of rows $\|x_i - x_j\|$ and $\|x_i - x_k\|$ are consistent with δ , i.e., if for every $i, j, k \in [n]$, we have:

$$\delta(i, j) < \delta(i, k) \iff \|x_i - x_j\| < \|x_i - x_k\|$$

Notice that by squaring the euclidean distances, we can determine the inequality for a triplet (i, j, k) , by examining the sign of the following³ polynomial:

$$p_{ijk}(X) \equiv p_{(i,j),(i,k)}(X) = \|x_i - x_j\|^2 - \|x_i - x_k\|^2$$

Notice how the sign of this polynomial determines the relative ordering for the distances among pair (i, j) and pair (i, k) . In other words, there is a 1-to-1 correspondence between the *sign-patterns* of the polynomials $p_{ijk}(X)$ for $i, j, k \in [n]$ and the induced triplet orders.

DEFINITION 3.1. (SIGN PATTERNS) *Let p_1, p_2, \dots, p_m be real degree-2 polynomials over l variables and let a point $u = (u_1, \dots, u_l) \in \mathbb{R}^l$ be such that none of them vanishes. The sign-pattern of the polynomials at point u is the m -tuple $(\sigma_1, \dots, \sigma_m) \in (-1, +1)^m$ where $\sigma_i = \text{sign}(p_i(u))$. We use the notation $\text{signs}(p_1, p_2, \dots, p_m)$ to denote the total number of different sign-patterns that can be obtained from p_1, p_2, \dots, p_m as point u ranges over all points in \mathbb{R}^l .*

FACT 3.2. (FROM [AFR85]) *For any integer number β between 1 and m , the total number of sign-patterns of m polynomials (as above) is upper bounded by:*

$$(3.1) \quad \text{signs}(p_1, p_2, \dots, p_m) \leq 4\beta \cdot (8\beta - 1)^{l + \frac{m}{\beta} - 1}$$

Final Comparison: Orderings vs Signs. The final step is a comparison. On the one hand, we already saw in Fact 3.1 that $\log G(n+1) \approx \frac{n^2}{2} \log n$ and hence the number of distinct triplet orderings grows as $\frac{n^2}{2} \log n$. On the other hand, the number of polynomials is exactly $m = n \binom{n-1}{2} \approx n^3/2$, because they are indexed by a triplet (i, j, k) where distances between (i, j) and (i, k) are compared.

We set $n = c \cdot d$, for some constant c (as we will see any $c > 2$ suffices), and let the variables be $l = n \cdot d$ (every point gets assigned d coordinates). We also set the parameter $\beta = \mu n = \mu c d$ for sufficiently large constant μ . Taking the logarithm on both sides of (3.1) (and ignoring the lower-order terms):

$$\log(\text{signs}(p_1, p_2, \dots, p_m)) \leq (nd + \frac{n^3}{2\mu n}) \log(8\mu n)$$

The dominant term becomes $cd^2 \log d$. Observe that for any $c > 2$, this term is strictly smaller than $\frac{n^2}{2} \log n$ and so for $n = cd \iff d = \frac{n}{c} < \frac{n}{2}$, there will be at least two distinct triplet-orderings that get mapped to the same sign pattern. This implies that any embedding that uses $d < n/2$ dimensions cannot be triplet-preserving (for worst-case instances), concluding the theorem. \square

REMARK 1. *Our results can be extended to other ℓ_p normed spaces too. For example, we can derive exactly the same lower bound for the dimension $d > \frac{n}{2}$ for any ℓ_p space (fixed $p \geq 2$). The polynomials would have to be of degree p instead of degree 2; their sign patterns are at most $2\beta p \cdot (4\beta p - 1)^{l + \frac{m}{\beta} - 1}$ ([AFR85]), so p affects only lower order terms.*

³With slight abuse of notation, we also use X to denote the flattened version of the matrix, that is an l -dimensional variable with $l = n \cdot d$, and the resulting polynomials depend on (some) coordinates of X .

3.2 Tradeoff for Triplet Dimension vs Relaxation

THEOREM 3.2. *There is an absolute constant $c > 0$ (taking $c = 8$ suffices) such that for every integer d and integer n , there is a metric space T on n points such that the triplet relaxation of any ordinal embedding of T into d -dimensional Euclidean space is at least $\frac{\log n}{\log d + \log \log n + c} - 1$.*

Proof. Let $G = (V, E)$ be a high-girth graph whose girth is $g = \frac{\log n}{\log d + \log \log n + 8}$. We can assume that the number of edges $m = |E| \geq \frac{1}{4}n^{1+\frac{1}{g}} > 16nd \log n$, by known constructions ([Sau70]). We will construct a large number N of edge subgraphs of G , namely G_1, G_2, \dots, G_N (all have vertex set V), with the following crucial property:

$$(*) \quad \text{For every pair } G_i, G_j, \exists v \in V : E_{G_i}(v) \setminus E_{G_j}(v) \neq \emptyset \text{ and } E_{G_j}(v) \setminus E_{G_i}(v) \neq \emptyset$$

Here $E_{G_i}(v)$ denotes the set of edges incident to v within G_i . Property $(*)$ captures that from the perspective of vertex v , the subgraphs G_i, G_j are significantly different.

DEFINITION 3.2. (WITNESS, “FARAWAY” METRIC PAIR) *For a pair of subgraphs G_i, G_j , we say that a vertex v is a witness (of faraway metric pair) G_i, G_j if the following two relations hold: $E_{G_i}(v) \setminus E_{G_j}(v) \neq \emptyset$ and $E_{G_j}(v) \setminus E_{G_i}(v) \neq \emptyset$. The pair G_i, G_j is said to be a faraway metric pair if there exists such a witness vertex v .*

Restating property $(*)$, we shall say that every pair G_i, G_j is faraway, as such a pair differs significantly in their triplets from at least the perspective of the vertex v . Our next goal becomes how to generate lots of faraway metric pairs. We independently sample edges of G with probability $\frac{1}{2}$ and generate N subgraphs G_1, G_2, \dots, G_N . As we will see, we will pick $N = 2^{bm}$ for $b < \frac{1}{2} \log_2 (4/3)$.

LEMMA 3.1. *The set of subgraphs G_1, G_2, \dots, G_N generated as above satisfies property $(*)$ with high probability.*

Proof. Let us fix a pair G_i, G_j and compute the probability it is not a faraway pair. Just for now, let us assume that the graph G was a k -regular graph (every vertex had degree exactly k). This simplifies the exposition and we show how to drop this assumption next. Fix a vertex v . We have:

$$\Pr[v \text{ is not witness for } G_i, G_j] = \Pr[E_{G_i}(v) \setminus E_{G_j}(v) = \emptyset \text{ or } E_{G_j}(v) \setminus E_{G_i}(v) = \emptyset]$$

Each of the two events has probability $(\frac{3}{4})^k$ and to see why this is the case let's consider the first event $E_{G_i}(v) \setminus E_{G_j}(v) = \emptyset$ (the other event follows in the same way). Consider an edge $e = (v, w)$. If $e \in E_{G_i}(v)$ and $e \in E_{G_j}(v)$, or if $e \notin E_{G_i}(v)$, then $E_{G_i}(v) \setminus E_{G_j}(v) = \emptyset$. The latter events are disjoint and their union has probability $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} = \frac{3}{4}$, since each edge was independently included w.p. $\frac{1}{2}$. This means that:

$$\Pr[v \text{ is not witness for } G_i, G_j] \leq 2 \cdot (\frac{3}{4})^k$$

At this point, notice that we would almost be done with the proof of the lemma, if we could somehow ensure that the event of whether or not a vertex v is a witness for G_i, G_j was independent from other vertices v' being witnesses for G_i, G_j . This is because we could bound the probability $\Pr[\text{all } v \in V \text{ are not witnesses for } G_i, G_j] \leq 2^n (\frac{3}{4})^{kn}$ and using a union bound over the N^2 pairs of subgraphs, we would get the lemma. Unfortunately, the independence of witnesses does not hold (as edges share endpoints), nor does the k -degree regularity assumption.

Let us now show how to circumvent the regularity and the independence assumptions. We will use a sequential process to determine whether a vertex v is a witness or not, where we have to redefine the notion of a vertex neighborhood as we sequentially process the edges. Let us fix an ordering over the vertices of the graph v_1, v_2, \dots, v_n . Let's redefine the neighborhoods of each vertex in an incremental way as follows:

- First vertex v_1 gets all edges $E_G(v_1)$. Denote this set with $E'_G(v_1)$.
- Second vertex v_2 gets all edges in $E_G(v_2)$ except those already selected by v_1 , i.e., all edges $E_G(v_2) \setminus E_G(v_1)$. Denote this set with $E'_G(v_2)$.
- Third vertex v_3 gets all edges $E_G(v_3) \setminus (E_G(v_1) \cup E_G(v_2))$, denoted as $E'_G(v_3)$. We continue until the end, where the last vertex v_n has no edges left.

Notice that each edge of G contributes to exactly one neighborhood (the first one it is part of), so $\sum_{v \in V} |E'_G(v)| = m$, where m is the total number of edges in G . With this new definition, we can indeed write:

$$\Pr[\text{all } v \text{ are not witnesses for } G_i, G_j] \leq 2 \left(\frac{3}{4}\right)^{|E'_G(v_1)|} \cdot 2 \left(\frac{3}{4}\right)^{|E'_G(v_2)|} \cdots 2 \left(\frac{3}{4}\right)^{|E'_G(v_n)|} \leq 2^n \left(\frac{3}{4}\right)^m$$

Finally, we take a union bound over the N^2 subgraphs to show that they satisfy property (*) with high probability:

$$\Pr[\text{there exists pair } G_i, G_j \text{ that is not faraway}] \leq \binom{N}{2} 2^n \left(\frac{3}{4}\right)^m$$

This bound goes to 0, if we pick $N^2 \ll \left(\frac{4}{3}\right)^m$ (e.g., $N = 1.14^m$ suffices). This concludes the proof of the lemma. \square

In order to conclude the proof of the theorem, we still need to compare the number $N = 1.14^m > 1.14^{16nd \log n} > 8^{nd \log n}$ of faraway metrics we generated, with the total number of sign patterns of polynomials given in Fact 3.2. Again, we care about triplet distances so there are $n^3/2$ polynomials of degree 2 over dn variables, yielding a number of sign patterns $2^{(2+o(1))nd \log n} \ll 8^{nd \log n}$.

By the pigeonhole principle, two distinct metric spaces from our collection of N subgraphs, get mapped to the same sign pattern so the distance orders in their embeddings are the same. Given that the graph has girth g and using (*), this implies the relaxation in at least one of these embeddings is at least $g - 1$, completing the proof. \square

4 Terminal Ordinal Embeddings

In this section, we present tight bounds for *terminal ordinal embeddings* (see Definition 2.2), which is analogous to the notion of terminal embeddings studied recently in metric embeddings ([EFN17, MMMR18, NN19, CN22]).

4.1 Upper Bound for k terminals Let $T = \{t_1, \dots, t_k\}$ be the set of terminals, and $V \setminus T = \{v_1, \dots, v_{n-k}\}$ be the set of the other points. For each pair $t \in T, v \in V$, let $r(t, v)$ be a unique integer in the range $1, \dots, kn$ that specifies the rank of the distance between t and v among all kn such distances.

THEOREM 4.1. *Terminal ordinal embedding for k terminals can be done with k dimensions.*

Proof. Our embedding (into a k -dimensional space) is as follows:

- Each terminal $t_i \in T$ is mapped to $f(t_i) = -Me_i$, where M is a “large” number to be specified soon, and e_i has 1 on the i^{th} position and 0 elsewhere.
- Each vertex v is mapped into a k -dimensional point $f(v) = [r(t_1, v), r(t_2, v), \dots, r(t_k, v)]$.

LEMMA 4.1. *For large enough M (picking $M = k^3n^2$ suffices), our mapping preserves the distance orders between any pairs (t, v) and (t', v') .*

Proof. To see this, observe that

$$\|f(v) - f(t)\|^2 = \sum_{s \neq t} r(s, v)^2 + (r(t, v) + M)^2 = \sum_s r(s, v)^2 + 2r(t, v)M + M^2$$

Let $C(t, v) = \sum_s r(s, v)^2$, and observe that $C(t, v) \leq k \cdot (kn)^2 = k^3n^2$.

Thus, if we set $M = k^3n^2$, then the term $2r(t, v)M + M^2$ “dominates”, i.e., $\|f(v) - f(t)\|^2 < \|f(v') - f(t')\|^2$ if and only if $r(t, v) < r(t', v')$. \square

So our embedding preserves the correct relationship between all pairs between T and the rest of the points, and only uses k dimensions. \square

REMARK 2. *Note that our construction did not explicitly use that k is fixed so it works for larger values of k too. As we show next, for values of $k = o(n)$ this is tight, and for $k = \Theta(n)$ this is tight up to a constant of 2 compared to the lower bound. Moreover, our approach works for other ℓ_p norms by choosing appropriate M .*

4.2 Lower Bound for $k = \Theta(n)$ Terminals We first analyze the case where there are many items playing the role of terminals, specifically $k = \lambda n$ with $\lambda \in (0, 1]$. We prove the following:

THEOREM 4.2. *For every constant $\lambda \in (0, 1]$, and for every large enough n , no d -dimensional embedding in ℓ_2 can be a terminal ordinal embedding for $k = \lambda n$ terminals, unless its dimension grows linearly as $d > \frac{n}{c}$, where $c > \frac{2}{2\lambda - \lambda^2}$. Equivalently, dimension $d > k(1 - \frac{\lambda}{2})$ is needed for any terminal ordinal embedding on $k = \lambda n$ terminals.*

For $\lambda = 1$, we recover as a special case a known result of Bilu and Linial for monotone maps [BL05].

Proof. We have k terminals, so the total number of distinct distances to-be-preserved is:

$$\text{Distinct Distance Elements} = (n-1) + (n-2) + \dots + (n-k) \approx kn - \frac{k^2}{2}$$

There are roughly $\left(kn - \frac{k^2}{2}\right)!$ different orderings, so taking logarithms we have that:

$$(4.2) \quad \log(\#\text{orderings}) \approx \left(kn - \frac{k^2}{2}\right) \log \left(kn - \frac{k^2}{2}\right)$$

On the other hand, the number of paired distances we can compare is:

$$m = \text{Paired Terminal Distances} = \binom{kn - \frac{k^2}{2}}{2} \approx \frac{k^2 n^2 - k^3 n + \frac{k^4}{4}}{2}$$

For each paired terminal distance we associate a degree 2 polynomial on at most nd variables, the sign of which determines the outcome of the distance comparison. The logarithm of different sign patterns that can arise is bounded by:

$$(4.3) \quad \log(\text{signs}(p_1, p_2, \dots, p_m)) \approx \left(nd + \frac{m}{\mu n^x}\right) \log n^x$$

where μ can be chosen to be a large constant, and $\mu n^x \in [1, m]$. We can determine when the quantity in (4.2) is asymptotically larger than the quantity in (4.3): after setting $n = cd$ (for constant c), parameter $x = 2$ and choosing μ large enough, we ultimately get (for sufficiently large n):

$$\left(\lambda n^2 - \frac{\lambda^2 n^2}{2}\right) \log \left(\lambda n^2 - \frac{\lambda^2 n^2}{2}\right) \gg \left(cd^2 + \frac{k^2 n^2 - k^3 n + \frac{k^4}{4}}{2\mu n^2}\right) \log d^2$$

The notation \gg simply means that the dominant term in the left-hand side becomes larger than the dominant term in the right-hand side, for n chosen sufficiently large. The above happens for:

$$c^2 d^2 \left(\lambda - \frac{\lambda^2}{2}\right) 2 \log d > cd^2 2 \log d \iff c > \frac{2}{2\lambda - \lambda^2}$$

which finishes the proof of the theorem. □

4.3 Lower Bound for $k = o(n)$ Terminals The difference from the previous subsection is that here the number of terminals is sublinear in n (the total number of items). Interestingly, for technical reasons that will become clear very soon, there is a qualitative (and quantitative) difference compared to when $k = \Theta(n)$, that allows us to get tighter bounds compared to previously. Specifically, our lower bound for the sublinear case shows that if the dimension is $d < (1 - \epsilon)k$ for any $\epsilon > 0$, then the embedding will mess up some of the distance orders. Hence, this is a tighter bound compared to Theorem 4.2, where a term $\lambda/2$ needs to be subtracted.

THEOREM 4.3. *Let $k = \sqrt{n}$. For any constant $c > 1$ and dimension $d < \frac{k}{c} = \frac{\sqrt{n}}{c}$, no d -dimensional embedding can be a terminal ordinal embedding. Generally, this lower bound holds for any $k = \lambda n^{1-\beta}$ for $\lambda \in (0, 1]$ and $\beta \in (0, 1]$. Notice that β is strictly larger than 0, so k is sublinear in n .*

Proof. The proof is found in the Appendix A. □

4.4 Lower Bound for Terminal Embeddings without Inter-terminal Comparisons Here we show tightness of our k terminal ordinal embeddings, even if we ignore the inter-terminal distance pairs (t, v) vs $(t', v'), t \neq t'$.

THEOREM 4.4. *Even ignoring inter-terminal distance pairs (comparisons between (t, v) and $(t', v'), t \neq t'$), for the case of $k = \lambda n$ the dimension needs to grow as $d > k(1 - \frac{\lambda}{2})$ for any terminal ordinal embedding on $k = \lambda n$ terminals. For range of $k = \lambda n^{1-\beta}$ with $\lambda \in (0, 1]$ and $\beta \in (0, 1]$, no d -dimensional embedding for $d < \frac{k}{c} = \frac{\sqrt{n}}{c}$ can be a terminal ordinal embedding.*

Proof. The proof is found in the Appendix A. \square

REMARK 3. *Our construction in Theorem 4.1 is tight for $k = o(n)$: no less than $(1 - \epsilon)k$ dimensions can perform k terminal embeddings. For linearly many terminals $k = \lambda n$, there is still a small gap between upper and lower bound: upper is that $d = k$ dimensions suffice, lower is that at least $d > k(1 - \frac{\lambda}{2})$ dimensions are necessary.*

5 Top- k -NNs Ordinal Embeddings

Let $\text{NN}(i)$ denote the set of k nearest neighbors to point i according to the original distance δ . The top- k -NN ordinal embedding (see Definition 2.3) asks to preserve the ordering on the set of distances for $\text{NN}(i)$, for every i .

THEOREM 5.1. *Any top- k -NNs ordinal embedding needs to have a dimension that grows as $d = \Omega(k)$. Without mixed comparisons (Definition 2.3), the bound becomes $d = \Omega\left(\frac{k \log k}{\log n}\right)$.*

Proof. The proof is found in the Appendix A. \square

6 Conclusion

We studied several basic questions that arise in the context of ordinal embeddings (also contrastive embeddings) related to nearest-neighbor, ranking, recommendations, crowdsourcing and psychometrics, where we want to find euclidean representations that only respect the order among measured distances between n items, rather than the lengths of those distances. For the well-motivated problem of triplets where information of the form “ i is closer to j or to k ?” is provided, we give almost tight lower bounds for the necessary dimension so as to preserve the triplets relations either exactly or approximately. Going beyond triplets, we study the interesting scenario of terminal ordinal embeddings and we present matching upper and lower bounds. Finally, we present lower bounds for the top- k -nearest-neighbors ordinal embeddings problem.

References

- [ABD⁺08] Noga Alon, Mihai Bădoiu, Erik D Demaine, Martin Farach-Colton, MohammadTaghi Hajiaghayi, and Anastasios Sidiropoulos. Ordinal embeddings of minimum relaxation: general properties, trees, and ultrametrics. *ACM Transactions on Algorithms (TALG)*, 4(4):1–21, 2008.
- [ABN07] Ittai Abraham, Yair Bartal, and Ofer Neiman. Local embeddings of metric spaces. In *Proceedings of the thirty-ninth annual ACM Symposium on Theory of Computing*, pages 631–640, 2007.
- [AFR85] Noga Alon, Peter Frankl, and V Rodl. Geometrical realization of set systems and probabilistic communication complexity. In *26th Annual Symposium on Foundations of Computer Science (sfcs 1985)*, pages 277–280. IEEE, 1985.
- [AI08] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Communications of the ACM*, 51(1):117–122, 2008.
- [AU15] Ehsan Amid and Antti Ukkonen. Multiview triplet embedding: Learning attributes in multiple maps. In *International Conference on Machine Learning*, pages 1472–1480. PMLR, 2015.
- [AWC⁺07] Sameer Agarwal, Josh Wills, Lawrence Cayton, Gert Lanckriet, David Kriegman, and Serge Belongie. Generalized non-metric multidimensional scaling. In *Artificial Intelligence and Statistics*, pages 11–18. PMLR, 2007.
- [BDH⁺08] Mihai Bădoiu, Erik D Demaine, MohammadTaghi Hajiaghayi, Anastasios Sidiropoulos, and Morteza Zadimoghaddam. Ordinal embedding: Approximation algorithms and dimensionality reduction. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 21–34. Springer, 2008.
- [BGJ10] Jarosław Byrka, Sylvain Guillemot, and Jesper Jansson. New results on optimizing rooted triplets consistency. *Discrete Applied Mathematics*, 158(11):1136–1147, 2010.

[BL05] Yonatan Bilu and Nati Linial. Monotone maps, sphericity and bounded second eigenvalue. *Journal of Combinatorial Theory, Series B*, 95(2):283–299, 2005.

[BN03] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.

[Bou85] Jean Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel Journal of Mathematics*, 52(1):46–52, 1985.

[CGCR20] Ines Chami, Albert Gu, Vaggos Chatziafratis, and Christopher Ré. From trees to continuous embeddings and back: Hyperbolic hierarchical clustering. *Advances in Neural Information Processing Systems*, 33:15065–15076, 2020.

[CM23] Vaggos Chatziafratis and Konstantin Makarychev. Triplet reconstruction and all other phylogenetic csp's are approximation resistant. In *Foundations of Computer Science (FOCS)*, 2023.

[CMA21] Vaggos Chatziafratis, Mohammad Mahdian, and Sara Ahmadian. Maximizing agreements for ranking, clustering and hierarchical clustering via max-cut. In *International Conference on Artificial Intelligence and Statistics*, pages 1657–1665. PMLR, 2021.

[CN22] Yeshwanth Cherapanamjeri and Jelani Nelson. Terminal embeddings in sublinear time. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1209–1216. IEEE, 2022.

[CNC18] Vaggos Chatziafratis, Rad Niazadeh, and Moses Charikar. Hierarchical clustering with structural constraints. In *International conference on machine learning*, pages 774–783. PMLR, 2018.

[CS74] James P Cunningham and Roger N Shepard. Monotone mapping of similarities into a general metric space. *Journal of Mathematical Psychology*, 11(4):335–363, 1974.

[EFN17] Michael Elkin, Arnold Filtser, and Ofer Neiman. Terminal embeddings. *Theoretical Computer Science*, 697:1–36, 2017.

[ES63] Paul Erdos and Horst Sachs. Regukre graphen gegebener taillenweite mit minimaler knotenzahl. *Wittenberg Math.-Natur. Reihe* 12, 251–257, 1963.

[EZK18] Ehsan Emamjomeh-Zadeh and David Kempe. Adaptive hierarchical clustering using ordinal queries. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 415–429. SIAM, 2018.

[FCM⁺20] Bohan Fan, Diego Ihara Centurion, Neshat Mohammadi, Francesco Sgherzi, Anastasios Sidiropoulos, and Mina Valizadeh. Learning lines with ordinal constraints. *arXiv preprint arXiv:2004.13202*, 2020.

[GCY19] Nikhil Ghosh, Yuxin Chen, and Yisong Yue. Landmark ordinal embedding. *Advances in Neural Information Processing Systems*, 32, 2019.

[GMR08] Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 573–582. IEEE, 2008.

[GPvL19] Debarghya Ghoshdastidar, Michaël Perrot, and Ulrike von Luxburg. Foundations of comparison-based hierarchical clustering. *Advances in neural information processing systems*, 32, 2019.

[HWvL20] Siavash Haghiri, Felix A Wichmann, and Ulrike von Luxburg. Estimation of perceptual scales using ordinal embedding. *Journal of vision*, 20(9):14–14, 2020.

[IMS17] Piotr Indyk, Jiří Matoušek, and Anastasios Sidiropoulos. 8: low-distortion embeddings of finite metric spaces. In *Handbook of discrete and computational geometry*, pages 211–231. Chapman and Hall/CRC, 2017.

[IN07] Piotr Indyk and Assaf Naor. Nearest-neighbor-preserving embeddings. *ACM Transactions on Algorithms (TALG)*, 3(3):31–es, 2007.

[JJF⁺15] Kevin G Jamieson, Lalit Jain, Chris Fernandez, Nicholas J Glattard, and Rob Nowak. Next: A system for real-world development, evaluation, and application of active learning. *Advances in neural information processing systems*, 28, 2015.

[JJN16] Lalit Jain, Kevin G Jamieson, and Rob Nowak. Finite sample prediction and recovery bounds for ordinal embedding. *Advances in neural information processing systems*, 29, 2016.

[JLS86] William B Johnson, Joram Lindenstrauss, and Gideon Schechtman. Extensions of lipschitz maps into banach spaces. *Israel Journal of Mathematics*, 54(2):129–138, 1986.

[JN11] Kevin G Jamieson and Robert D Nowak. Low-dimensional embedding using adaptively selected ordinal data. In *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1077–1084. IEEE, 2011.

[KL14] Matthäus Kleindessner and Ulrike Luxburg. Uniqueness of ordinal embedding. In *Conference on Learning Theory*, pages 40–67. PMLR, 2014.

[Kru64a] Joseph B Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29(1):1–27, 1964.

[Kru64b] Joseph B Kruskal. Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29(2):115–129, 1964.

[Kvh16] Ramya Korlakai Vinayak and Babak Hassibi. Crowdsourced clustering: Querying edges vs triangles. *Advances*

in Neural Information Processing Systems, 29, 2016.

[KvL17] Matthias Kleindessner and Ulrike von Luxburg. Kernel functions based on triplet comparisons. *Advances in neural information processing systems*, 30, 2017.

[LHvL19] Michael Lohaus, Philipp Hennig, and Ulrike von Luxburg. Uncertainty estimates for ordinal embeddings. *arXiv preprint arXiv:1906.11655*, 2019.

[Mat13] Jiri Matoušek. Lecture notes on metric embeddings. Technical report, Technical report, ETH Zürich, 2013.

[MMMR18] Sepideh Mahabadi, Konstantin Makarychev, Yury Makarychev, and Ilya Razenshteyn. Nonlinear dimension reduction via outer bi-lipschitz extensions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1088–1101, 2018.

[MZS⁺19] Nicholas Monath, Manzil Zaheer, Daniel Silva, Andrew McCallum, and Amr Ahmed. Gradient-based hierarchical clustering using continuous representations of trees in hyperbolic space. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 714–722, 2019.

[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*, pages 1064–1069, 2019.

[RRŠ89] Jan Reiterman, Vojtech Rödl, and E Šinajová. Geometrical embeddings of graphs. *Discrete Mathematics*, 74(3):291–319, 1989.

[SAG⁺22] Nikunj Saunshi, Jordan Ash, Surbhi Goel, Dipendra Misra, Cyril Zhang, Sanjeev Arora, Sham Kakade, and Akshay Krishnamurthy. Understanding contrastive learning requires incorporating inductive biases. In *International Conference on Machine Learning*, pages 19250–19286. PMLR, 2022.

[Sau70] N Sauer. On the existence of regular n-graphs with given girth. *Journal of Combinatorial Theory*, 9(2):144–147, 1970.

[SE05] Noah A Smith and Jason Eisner. Contrastive estimation: Training log-linear models on unlabeled data. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 354–362, 2005.

[She62] Roger N Shepard. The analysis of proximities: multidimensional scaling with an unknown distance function. i. *Psychometrika*, 27(2):125–140, 1962.

[She74] Roger N Shepard. Representation of structure in similarity data: Problems and prospects. *Psychometrika*, 39(4):373–421, 1974.

[SJ03] Matthew Schultz and Thorsten Joachims. Learning a distance metric from relative comparisons. *Advances in neural information processing systems*, 16, 2003.

[SS09] Gideon Schechtman and Adi Shraibman. Lower bounds for local versions of dimension reductions. *Discrete & Computational Geometry*, 41(2):273–283, 2009.

[Thu54] Louis Leon Thurstone. The measurement of values. *Psychological review*, 61(1):47, 1954.

[TL14] Yoshikazu Terada and Ulrike Luxburg. Local ordinal embedding. In *International Conference on Machine Learning*, pages 847–855. PMLR, 2014.

[TLB⁺11] Omer Tamuz, Ce Liu, Serge Belongie, Ohad Shamir, and Adam Tauman Kalai. Adaptively learning the crowd kernel. *28th International Conference on Machine Learning (ICML)*, 2011.

[Tor52] Warren S Torgerson. Multidimensional scaling: I. theory and method. *Psychometrika*, 17(4):401–419, 1952.

[UDH15] Antti Ukkonen, Behrouz Derakhshan, and Hannes Heikinheimo. Crowdsourced nonparametric density estimation using relative distances. In *Third AAAI Conference on Human Computation and Crowdsourcing*, 2015.

[VD16] Sharad Vikram and Sanjoy Dasgupta. Interactive bayesian hierarchical clustering. In *International Conference on Machine Learning*, pages 2081–2090. PMLR, 2016.

[VDMW12] Laurens Van Der Maaten and Kilian Weinberger. Stochastic triplet embedding. In *2012 IEEE International Workshop on Machine Learning for Signal Processing*, pages 1–6. IEEE, 2012.

[VHL⁺19] Leena Chennuru Vankadara, Siavash Haghiri, Michael Lohaus, Faiz Ul Wahab, and Ulrike von Luxburg. Insights into ordinal embedding algorithms: A systematic evaluation. *arXiv preprint arXiv:1912.01666*, 2019.

[XSB06] Lin Xiao, Jun Sun, and Stephen Boyd. A duality view of spectral methods for dimensionality reduction. In *Proceedings of the 23rd international conference on Machine learning*, pages 1041–1048, 2006.

A Omitted Proofs

Proof. [Proof of Theorem 4.3] The proof proceeds in analogous manner as before, but the difference is that the term $kn - \frac{k^2}{2} \approx kn$, i.e., $\frac{k^2}{2} \ll kn$ for large n and so it won't affect the computations. We present the relevant calculations below. The important quantities we need to compare are:

- The logarithm of the total number of orderings: $\left(kn - \frac{k^2}{2} \right) \log \left(kn - \frac{k^2}{2} \right)$ vs
- The logarithm of sign patterns of polynomials: $\left(nd + \frac{m}{\mu n^x} \right) \log n^x$, with parameters μ, x : $\mu n^x \leq m$, where

$$m = \frac{k^2 n^2 - k^3 n + \frac{k^4}{4}}{2}, \text{ same as before.}$$

We can make several simplifications based on the fact that $k = \lambda n^{1-\beta} = o(n)$. Ignoring lower order terms we have that:

- $k^2 = o(kn)$ and also $m \approx \frac{k^2 n^2}{2} = \Theta(k^2 n^2)$.
- We can pick x such that $n^x \approx kn = o(n^2)$ and $\mu n^x \ll m$ for any constant μ .
- Let $k = c \cdot d$ for some constant c (to be determined below).

Taking these into account we need to compare: $cdn \log kn$ and $\left(nd + \frac{cdn}{2\mu}\right) \log kn$.

We can set μ large enough so that the former quantity is larger than the latter whenever $c > 1$, i.e., $k \geq cd$ for constant $c > 1$. \square

Proof. [Proof of Theorem 4.4] The proof follows almost the same calculations as in Theorem 4.2 and Theorem 4.3. The reason is the following: if we cared about only same-terminal distance pairs we get:

- $\#\text{orders} > (n-1)!(n-2)!\dots(n-k)!$
- Again the number of paired comparisons is at most $m \approx \frac{k^2 n^2}{2} = \Theta(n^2)$, so this is also the number of corresponding polynomials that determine the orderings.
- Number of orders is larger than $[(n-k)!]^k \implies \log \#\text{orders} > k \log[(n-k)!]$.
- The leading term again is $kn \log n$, and this is enough to beat the number of polynomials and get the tradeoffs we want. If we wanted to be more precise, we can use that $\log[(n-1)!(n-2)!\dots(n-k)!] \approx (kn - \frac{k^2}{2}) \log(n-k)$.

We can conclude that the same lower bounds (with same constants) hold. \square

Proof. [Proof of Theorem 5.1] Following the approach as in the previous proofs, we compare the following:

- Allowing for mixed comparisons, the number of possible orderings among distances in $\cup_i \text{NN}(i)$ for all the items i can be close to $(nk - o(nk))!$. Hence the logarithm is $\approx nk \log n$.
- Allowing for mixed comparisons, there are at most $m = \binom{nk}{2}$ paired distance comparisons (if we don't allow for mixed comparisons, m is even smaller). So the logarithm of the sign patterns for the corresponding polynomials is $(nd + \frac{m}{\mu n^x}) \log(\mu n^x)$ for any $x, \mu : \mu n^x \leq m$. For this calculation, it's sufficient to pick x : $n^x \approx kn = O(n^2)$ so $x \leq 2$.
- Comparing the logarithms we get that there are more orderings than sign patterns as long as: $nk \log n \gg ndx \log n$. So the dimension d must be at least $\Omega(k)$.
- Without mixed comparisons, the number of orderings per point is $k!$ since $|\text{NN}(i)| = k$. So we get in total at least $(k!)^{\Theta(n)}$ distance orderings (so $\log(k!)^{\Theta(n)} \approx nk \log k$). This implies that for the embedding to preserve distance orderings it has to use dimension $d = \Omega\left(\frac{k \log k}{\log n}\right)$.

\square