
Learning Nearest Neighbor Graphs from Noisy Distance Samples

Blake Mason *
University of Wisconsin
Madison, WI 53706
bmason3@wisc.edu

Ardhendu Tripathy *
University of Wisconsin
Madison, WI 53706
astripathy@wisc.edu

Robert Nowak
University of Wisconsin
Madison, WI 53706
rdnowak@wisc.edu

Abstract

We consider the problem of learning the *nearest neighbor graph* of a dataset of n items. The metric is unknown, but we can query an oracle to obtain a noisy estimate of the distance between any pair of items. This framework applies to problem domains where one wants to learn people’s preferences from responses commonly modeled as noisy distance judgments. In this paper, we propose an active algorithm to find the graph with high probability and analyze its query complexity. In contrast to existing work that forces Euclidean structure, our method is valid for general metrics, assuming only symmetry and the triangle inequality. Furthermore, we demonstrate efficiency of our method empirically and theoretically, needing only $\mathcal{O}(n \log(n) \Delta^{-2})$ queries in favorable settings, where Δ^{-2} accounts for the effect of noise. Using crowd-sourced data collected for a subset of the UT Zappos50K dataset, we apply our algorithm to learn which shoes people believe are most similar and show that it beats both an active baseline and ordinal embedding.

1 Introduction

In modern machine learning applications, we frequently seek to learn proximity/ similarity relationships between a set of items given only noisy access to pairwise distances. For instance, practitioners wishing to estimate internet topology frequently collect one-way-delay measurements to estimate the distance between a pair of hosts [9]. Such measurements are affected by physical constraints as well as server load, and are often noisy. Researchers studying movement in hospitals from WiFi localization data likewise contend with noisy distance measurements due to both temporal variability and varying signal strengths inside the building [4]. Additionally, human judgments are commonly modeled as noisy distances [26, 23]. As an example, Amazon Discover asks customers their preferences about different products and uses this information to recommend new items it believes are similar based on this feedback. We are often primarily interested in the *closest* or *most similar* item to a given one— e.g., the closest server, the closest doctor, the most similar product. The particular item of interest may not be known *a priori*. Internet traffic can fluctuate, different patients may suddenly need attention, and customers may be looking for different products. To handle this, we must learn the closest/ most similar item for *each* item. This paper introduces the problem of learning the *Nearest Neighbor Graph* that connects each item to its nearest neighbor from noisy distance measurements.

Problem Statement: Consider a set of n points $\mathcal{X} = \{x_1, \dots, x_n\}$ in a metric space. The metric is unknown, but we can query a stochastic oracle for an estimate of any pairwise distance. In as few queries as possible, we seek to learn a nearest neighbor graph of \mathcal{X} that is correct with probability $1 - \delta$, where each x_i is a vertex and has a directed edge to its nearest neighbor $x_{i^*} \in \mathcal{X} \setminus \{x_i\}$.

*Authors contributed equally to this paper and are listed alphabetically.

1.1 Related work

Nearest neighbor problems (from noiseless measurements) are well studied and we direct the reader to [3] for a survey. [6, 30, 25] all provide theory and algorithms to learn the nearest neighbor graph which apply in the noiseless regime. Note that the problem in the noiseless setting is *very* different. If noise corrupts measurements, the methods from the noiseless setting can suffer persistent errors. There has been recent interest in introducing noise via subsampling for a variety of distance problems [24, 1, 2], though the noise here is not actually part of the data but introduced for efficiency. In our algorithm, we use the triangle inequality to get tighter estimates of noisy distances in a process equivalent to the classical Floyd–Warshall [11, 7]. This has strong connections to the metric repair literature [5, 13] where one seeks to alter a set of noisy distance measurements as little as possible to learn a metric satisfying the standard axioms. [27] similarly uses the triangle inequality to bound unknown distances in a related but noiseless setting. In the specific case of noisy distances corresponding to human judgments, a number of algorithms have been proposed to handle related problems, most notably Euclidean embedding techniques, e.g. [17, 31, 23]. To reduce the load on human subjects, several attempts at an active method for learning Euclidean embeddings have been made but have only seen limited success [20]. Among the culprits is the strict and often unrealistic modeling assumption that the metric be Euclidean and low dimensional. In the particular case that the algorithm may query triplets (e.g., “is i or j closer to k ?”) and receive noisy responses, [22] develop an interesting, passive technique under general metrics for learning a relative neighborhood graph which is an undirected relaxation of a nearest neighbor graph.

1.2 Main contributions

In this paper we introduce the problem of identifying the *nearest neighbor graph* from noisy distance samples and propose ANNTri, an active algorithm, to solve it for general metrics. We empirically and theoretically analyze its complexity to show improved performance over a passive and an active baseline. In favorable settings, such as when the data forms clusters, ANNTri needs only $\mathcal{O}(n \log(n)/\Delta^2)$ queries, where Δ accounts for the effect of noise. Furthermore, we show that ANNTri achieves superior performance compared to methods which require much stronger assumptions. We highlight two such examples. In Fig. 2c, for an embedding in \mathbb{R}^2 , ANNTri outperforms the common technique of triangulation that works by estimating each point’s distance to a set of anchors. In Fig. 3b, we show that ANNTri likewise outperforms Euclidean embedding for predicting which images are most similar from a set of similarity judgments collected on Amazon Mechanical Turk. The rest of the paper is organized as follows. In Section 2, we further setup the problem. In Sections 3 and 4 we present the algorithm and analyze its theoretical properties. In Section 5 we show ANNTri’s empirical performance on both simulated and real data. In particular, we highlight its efficiency in learning from human judgments.

2 Problem setup and summary of our approach

We denote distances as $d_{i,j}$ where $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ is a distance function satisfying the standard axioms and define $x_{i^*} := \arg \min_{x \in \mathcal{X} \setminus \{x_i\}} d(x_i, x)$. Though the distances are unknown, we are able to draw independent samples of its true value according to a stochastic distance oracle, i.e. querying

$$Q(i, j) \text{ yields a realization of } d_{i,j} + \eta, \quad (1)$$

where η is a zero-mean subGaussian random variable assumed to have scale parameter $\sigma = 1$. We let $\hat{d}_{i,j}(t)$ denote the empirical mean of the values returned by $Q(i, j)$ queries made until time t . The number of $Q(i, j)$ queries made until time t is denoted as $T_{i,j}(t)$. A possible approach to obtain the nearest neighbor graph is to repeatedly query all $\binom{n}{2}$ pairs and report $x_{i^*}(t) = \arg \min_{j \neq i} \hat{d}_{i,j}(t)$ for all $i \in [n]$. But since we only wish to learn $x_{i^*} \forall i$, if $d_{i,k} \gg d_{i,i^*}$, we do not need to query $Q(i, k)$ as many times as $Q(i, i^*)$. To improve our query efficiency, we could instead adaptively sample to focus queries on distances that we estimate are smaller. A simple adaptive method to find the nearest neighbor graph would be to iterate over x_1, x_2, \dots, x_n and use a best-arm identification algorithm to find x_{i^*} in the i^{th} round.¹ However, this procedure treats each round independently, ignoring properties of metric spaces that allow information to be shared between rounds.

¹We could also proceed in a non-iterative manner, by adaptively choosing which among $\binom{n}{2}$ pairs to query next. However this has worse empirical performance and same theoretical guarantees as the in-order approach.

- Due to symmetry, for any $i < j$ the queries $Q(i, j)$ and $Q(j, i)$ follow the same law, and we can *reuse* values of $Q(i, j)$ collected in the i^{th} round while finding x_{j^*} in the j^{th} round.
- Using concentration bounds on $d_{i,j}$ and $d_{i,k}$ from samples from $Q(i, j)$ and $Q(i, k)$ collected in the i^{th} round, we can bound $d_{j,k}$ via the triangle inequality. As a result, we may be able to state $x_k \neq x_{j^*}$ without even querying $Q(j, k)$.

Our proposed algorithm ANNTri uses all the above ideas to find the nearest neighbor graph of \mathcal{X} . For general \mathcal{X} , the sample complexity of ANNTri contains a problem-dependent term that involves the order in which the nearest neighbors are found. For an \mathcal{X} consisting of sufficiently well separated clusters, this order-dependence for the sample complexity does not exist.

3 Algorithm

Our proposed algorithm (Algorithm 1) ANNTri finds the nearest neighbor graph of \mathcal{X} with probability $1 - \delta$. It iterates over $x_j \in \mathcal{X}$ in order of their subscript index and finds x_{j^*} in the j^{th} ‘round’. All bounds, counts of samples, and empirical means are stored in $n \times n$ symmetric matrices in order to share information between different rounds. We use Python array/Matlab notation to indicate individual entries in the matrices, for e.g., $\hat{d}[i, j] = \hat{d}_{i,j}(t)$. The number of $Q(i, j)$ queries made is queried is stored in the $(i, j)^{th}$ entry of T . Matrices U and L record upper and lower confidence bounds on $d_{i,j}$. U^Δ and L^Δ record the associated triangle inequality bounds. Symmetry is ensured by updating the $(j, i)^{th}$ entry at the same time as the $(i, j)^{th}$ entry for each of the above matrices. In the j^{th} round, ANNTri finds the correct x_{j^*} with probability $1 - \delta/n$ by calling SETri (Algorithm 2), a modification of the successive elimination algorithm for best-arm identification. In contrast to standard successive elimination, at each time step SETri only samples those points in the active set that have the fewest number of samples.

Algorithm 1 ANNTri

Require: n , procedure SETri (Alg. 2), confidence δ

- 1: Initialize \hat{d}, T as $n \times n$ matrices of zeros, U, U^Δ as $n \times n$ matrices where each entry is ∞ , L, L^Δ as $n \times n$ matrices where each entry is $-\infty$, NN as a length n array
- 2: **for** $j = 1$ **to** n **do**
- 3: **for** $i = 1$ **to** n **do** {find tightest triangle bounds}
- 4: **for all** $k \neq i$ **do**
- 5: Set $U^\Delta[i, k], U^\Delta[k, i] \leftarrow \min_\ell U_{i,k}^{\Delta_\ell}$, see (7)
- 6: Set $L^\Delta[i, k], L^\Delta[k, i] \leftarrow \max_\ell L_{i,k}^{\Delta_\ell}$, see (8)
- 7: NN[j] = SETri($j, \hat{d}, U, U^\Delta, L, L^\Delta, T, \xi = \delta/n$)
- 8: **return** The nearest neighbor graph adjacency list NN

Algorithm 2 SETri

Require: index j , callable oracle $Q(\cdot, \cdot)$ (Eq. (1)), six $n \times n$ matrices: $\hat{d}, U, U^\Delta, L, L^\Delta, T$, confidence ξ

- 1: Initialize active set $\mathcal{A}_j \leftarrow \{a \neq j : \max\{L[a, j], L^\Delta[a, j]\} < \min_k \min\{U[j, k], U^\Delta[j, k]\}\}$
- 2: **while** $|\mathcal{A}_j| > 1$ **do**
- 3: **for all** $i \in \mathcal{A}_j$ such that $T[i, j] = \min_{k \in \mathcal{A}_j} T[i, k]$ **do** {only query points with fewest samples}
- 4: Update $\hat{d}[i, j], \hat{d}[j, i] \leftarrow (\hat{d}[i, j] \cdot T[i, j] + Q(i, j)) / (T[i, j] + 1)$
- 5: Update $T[i, j], T[j, i] \leftarrow T[i, j] + 1$
- 6: Update $U[i, j], U[j, i] \leftarrow \hat{d}[i, j] + C_\xi(T[i, j])$
- 7: Update $L[i, j], L[j, i] \leftarrow \hat{d}[i, j] - C_\xi(T[i, j])$
- 8: Update $\mathcal{A}_j \leftarrow \{a \neq j : \max\{L[a, j], L^\Delta[a, j]\} < \min_k \min\{U[j, k], U^\Delta[j, k]\}\}$
- 9: **return** The index i for which $x_i \in \mathcal{A}_j$

3.1 Confidence bounds on the distances

Using the subGaussian assumption on the noise random process, we can use Hoeffding's inequality and a union bound over time to get the following confidence intervals on the distance $d_{j,k}$:

$$|\hat{d}_{j,k}(t) - d_{j,k}| \leq \sqrt{2 \frac{\log(4n^2(T_{j,k}(t))^2/\delta)}{T_{j,k}(t)}} =: C_{\delta/n}(T_{j,k}(t)), \quad (2)$$

which hold for all points $x_k \in \mathcal{X} \setminus \{x_j\}$ at all times t with probability $1 - \delta/n$, i.e.

$$\mathbb{P}(\forall t \in \mathbb{N}, \forall i \neq j, d_{i,j} \in [L_{i,j}(t), U_{i,j}(t)]) \geq 1 - \delta/n, \quad (3)$$

where $L_{i,j}(t) := \hat{d}_{i,j}(t) - C_{\delta/n}(T_{i,j}(t))$ and $U_{i,j}(t) := \hat{d}_{i,j}(t) + C_{\delta/n}(T_{i,j}(t))$. [10] use the above procedure to derive the following upper bound for the number of oracle queries used to find x_{j^*} :

$$\mathcal{O} \left(\sum_{k \neq j} \frac{\log(n^2/(\delta \Delta_{j,k}))}{\Delta_{j,k}^2} \right), \quad (4)$$

where for any $x_k \notin \{x_j, x_{j^*}\}$ the suboptimality gap $\Delta_{j,k} := d_{j,k} - d_{j,j^*}$ characterizes how hard it is to rule out x_k from being the nearest neighbor. We also set $\Delta_{j,j^*} := \min_{k \notin \{j, j^*\}} \Delta_{j,k}$. Note that one can use tighter confidence bounds as detailed in [12] and [18] to obtain sharper bounds on the sample complexity of this subroutine.

3.2 Computing the triangle bounds and active set $\mathcal{A}_j(t)$

Since $\mathcal{A}_j(\cdot)$ is only computed within SETri, we abuse notation and use its argument t to indicate the time counter private to SETri. Thus, the initial active set computed by SETri when called in the j^{th} round is denoted $\mathcal{A}_j(0)$. During the j^{th} round, the active set $\mathcal{A}_j(t)$ contains all points that have not been eliminated from being the nearest neighbor of x_j at time t . In what follows, we add a superscript Δ to denote a bound obtained via the triangle inequality, whose precise definitions are given in Lemma 3.1. We define x_j 's active set at time t as

$$\mathcal{A}_j(t) := \{a \neq j : \max\{L_{a,j}(t), L_{a,j}^\Delta(t)\} < \min_k \min\{U_{j,k}(t), U_{j,k}^\Delta(t)\}\}. \quad (5)$$

Assuming $L_{a,j}^\Delta(t)$ and $U_{j,k}^\Delta(t)$ are valid lower and upper bounds on $d_{a,j}$, $d_{j,k}$ respectively, (5) states that point x_a is active if its lower bound is less than the minimum upper bound for $d_{j,k}$ over all choices of $x_k \neq x_j$. Next, for any (j, k) we construct triangle bounds L^Δ, U^Δ on the distance $d_{j,k}$. Intuitively, for some reals g, g', h, h' , if $d_{i,j} \in [g, g']$ and $d_{i,k} \in [h, h']$ then $d_{j,k} \leq g' + h'$, and

$$d_{j,k} \geq |d_{i,j} - d_{i,k}| = \max\{d_{i,j}, d_{i,k}\} - \min\{d_{i,j}, d_{i,k}\} \geq (\max\{g, h\} - \min\{g', h'\})_+ \quad (6)$$

where $(s)_+ := \max\{s, 0\}$. The lower bound can be seen as true by Fig. 7 in the Appendix. Lemma 3.1 uses these ideas to form upper and lower bounds on distances by the triangle inequality. Note that this definition is inherently recursive as it may rely on past triangle inequality bounds to achieve the tightest possible result. We denote a triangle inequality upper and lower bounds on $d_{j,k}$ due to a point i at time t as $U_{j,k}^{\Delta_i}$ and $L_{j,k}^{\Delta_i}$ respectively.

Lemma 3.1. For all $k \neq 1$, $U_{1,k}^{\Delta_1}(t) = U_{1,k}^\Delta(t) = U_{1,k}(t)$. For any $i < j$ define

$$U_{j,k}^{\Delta_i}(t) := \min_{\max\{i_1, i_2\} < i} (\min\{U_{i,j}(t), U_{i,j}^{\Delta_{i_1}}(t)\} + \min\{U_{i,k}(t), U_{i,k}^{\Delta_{i_2}}(t)\}). \quad (7)$$

For all $k \neq 1$, $L_{1,k}^{\Delta_1}(t) = L_{1,k}^\Delta(t) = L_{1,k}(t)$. For any $i < j$ define

$$L_{j,k}^{\Delta_i}(t) := \max_{\max\{i_1, i_2, i_3, i_4\} < i} \left(\max\{L_{i,j}(t), L_{i,j}^{\Delta_{i_1}}(t), L_{i,k}(t), L_{i,k}^{\Delta_{i_2}}(t)\} \right. \\ \left. - \min\{U_{i,j}(t), U_{i,j}^{\Delta_{i_3}}(t), U_{i,k}(t), U_{i,k}^{\Delta_{i_4}}(t)\} \right)_+, \quad (8)$$

where $(s)_+ := \max\{s, 0\}$. If all the bounds obtained by SETri in rounds $i < j$ are correct then

$$d_{j,k} \in [L_{j,k}^\Delta(t), U_{j,k}^\Delta(t)], \quad \text{where} \quad L_{j,k}^\Delta(t) := \max_{i < j} L_{j,k}^{\Delta_i}(t) \quad \text{and} \quad U_{j,k}^\Delta(t) := \min_{i < j} U_{j,k}^{\Delta_i}(t).$$

The proof is in Appendix B.1. ANNTri has access to two sources of bounds on distances: concentration bounds and triangle inequality bounds, and as can be seen in Lemma 3.1, the former affects the latter. Furthermore, triangle bounds are computed from other triangle bounds, leading to the recursive definitions of $L_{j,k}^{\Delta_i}$ and $U_{j,k}^{\Delta_i}$. Because of these facts, triangle bounds are dependent on the order in which ANNTri finds each nearest neighbor. These bounds can be computed using dynamic programming and brute force search over all possible i_1, i_2, i_3, i_4 is not necessary. We note that the above recursion is similar to the Floyd-Warshall algorithm for finding shortest paths between all pairs of vertices in a weighted graph [11, 7]. The results in [27] show that the triangle bounds obtained in this manner have the minimum L_1 norm between the upper and lower bound matrices.

Extension to k -Nearest Neighbor Graphs: All algorithms and theory in this paper can be extended to the case of k -nearest neighbor graphs where one wishes to draw a directed edge from each point to all of its k -nearest neighbors. To modify ANNTri, one can change the subroutine SETri to be a variant of the KL-Racing algorithm by [21], for instance. Racing style algorithms are the natural extension of Successive-Elimination style algorithms to the top- k bandit setting. To achieve best complexity, in this case, one would again want to sample the distances with the only minimum number of calls to the oracle first, as in Line 3 of SETri. To make a statement similar to Theorem 4.4, i.e. to bound the complexity of learning k -NN graphs, it is necessary to alter the definition of the events $A_{j,k}$ so as to certify that a point has been eliminated from being a k -nearest neighbor as opposed to a 1-nearest neighbor in the current form. The suboptimality gaps $\Delta_{j,k}$ (and therefore $H_{j,k}$) would be defined differently for the k -nearest neighbor case leading to a different bound. A similar statement as Theorem 4.6 can likewise be achieved as long as $k < \log(n)$ and one should expect a complexity of $\mathcal{O}\left(kn \log(n) \overline{\Delta}^{-2}\right)$ for an appropriately defined $\overline{\Delta}^{-2}$.

4 Analysis

All omitted proofs of this section can be found in the Appendix Section B.

Theorem 4.1. *ANNTri finds the nearest neighbor for each point in \mathcal{X} with probability $1 - \delta$.*

4.1 A simplified algorithm

The following Lemma indicates which points must be eliminated initially in the j^{th} round.

Lemma 4.2. *If $\exists i : 2U_{i,j} < L_{i,k}$, then $x_k \notin \mathcal{A}_j(0)$ for ANNTri.*

Proof. $2U_{i,j} < L_{i,k} \iff U_{i,j} < L_{i,k} - U_{i,j} \leq L_{j,k}^{\Delta_i}$ □

Next, we define ANNEasy, a simplified version of ANNTri that is more amenable to analysis. Here, we say that x_k is eliminated in the j^{th} round of ANNEasy if i) $k < j$ and $\exists i : U_{i,j} < L_{j,k}$ (symmetry from past samples) or ii) $\exists i : 2U_{i,j} < L_{i,k}$ (Lemma 4.2). Therefore, x_j 's active set for ANNEasy is

$$\mathcal{A}_j = \{a \neq j : L_{a,k} \leq 2U_{j,k} \forall k \text{ and } L_{a,j} < \min_k U_{j,k}\}. \quad (9)$$

To define ANNEasy in code, we remove lines 3-6 of ANNTri (Algorithm 1), and call a subroutine SEEasy in place of SETri. SEEasy matches SETri (Algorithm 2) except that lines 1 and 8 are replaced with (9) instead. We provide full pseudocode of both ANNEasy and SEEasy in the Appendix A.1.1. Though ANNEasy is a simplification for analysis, we note that it empirically captures much of the same behavior of ANNTri. In the Appendix A.1.2 we provide an empirical comparison of the two.

4.2 Complexity of ANNEasy

We now turn our attention to account for the effect of the triangle inequality in ANNEasy.

Lemma 4.3. *For any $x_k \in \mathcal{X}$ if the following conditions hold for some $i < j$, then $x_k \notin \mathcal{A}_j(0)$.*

$$6C_{\delta/n}(1) \leq d_{i,k} - 2d_{i,j} \quad \text{and} \quad \{j, k\} \cap (\cup_{m < i} \{\ell : 2d_{m,i} < d_{m,\ell}\}) = \emptyset. \quad (10)$$

The first condition characterizes which x_k 's must satisfy the condition in Lemma 4.2 for the j^{th} round. The second guarantees that x_k was sampled in the i^{th} round, a necessary condition for forming triangle bounds using x_i .




Figure 1: Example datasets where triangle inequalities lead to provable gains.

Theorem 4.4. *Conditioned on the event that all confidence bounds are valid at all times, ANNEasy learns the nearest neighbor graph of \mathcal{X} in the following number of calls to the distance oracle:*

$$\mathcal{O} \left(\sum_{j=1}^n \sum_{k>j} \mathbb{1}_{[A_{j,k}]} H_{j,k} + \sum_{k<j} \mathbb{1}_{[A_{j,k}]} (H_{j,k} - \mathbb{1}_{[A_{k,j}]} H_{k,j})_+ \right). \quad (11)$$

In the above expression $H_{j,k} := \frac{\log(n^2/(\delta\Delta_{j,k}))}{\Delta_{j,k}^2}$ and $\mathbb{1}_{[A_{j,k}]} := 1$, if x_k does not satisfy the triangle inequality elimination conditions of (10) $\forall i < j$, and 0 otherwise.

The expression in (11) can be understood as a sum over the complexity of each of the n rounds, as specified by the outer sum. The complexity of each individual round is a sum of two terms. Consider the j^{th} round. The first term bounds the number of calls to $Q(j, k)$ for all $k > j$. In general $H_{j,k}$ calls are necessary, unless a triangle inequality bound allows for elimination of k without sampling, as given by $\mathbb{1}_{[A_{j,k}]}$. The second term bounds the number of calls to $Q(j, k)$ for all $k < j$. It has the same form as the first term, except we must now use past samples we may already have via symmetry of distances (provided the triangle inequality did not prevent us from querying $Q(k, j)$ in the previous round). The $(\cdot)_+$ operation prevents negative terms, since it may be the case that no additional samples are necessary, even if we don't use the triangle inequality for elimination.

In Theorem B.6, in the Appendix, we state the sample complexity when triangle inequality bounds are ignored by ANNTri, and this upper bounds (11). Whether a point can be eliminated by the triangle inequality depends both on the underlying distances and the order in which ANNTri finds each nearest neighbor (c.f. Lemma 4.3). In general, this dependence on the order is necessary to ensure that past samples exist and may be used to form upper and lower bounds. Furthermore, it is worth noting that even without noise the triangle inequality may not always help. A simple example is any arrangement of points such that $0 < r \leq d_{j,k} < 2r \forall j, k$. To see this, consider triangle bounds on any distance $d_{j,k}$ due to any $x_i, x_{i'} \in \mathcal{X} \setminus \{x_j, x_k\}$. Then $|d_{i,j} - d_{i,k}| \leq r < 2r \leq d_{i',j} + d_{i',k} \forall i, i'$ so $L_{i,j}^\Delta < U_{j,k}^\Delta \forall i, j, k$. Thus no triangle upper bounds separate from triangle lower bounds so no elimination via the triangle inequality occurs. In such cases, it is necessary to sample all $\mathcal{O}(n^2)$ distances. However, in more favorable settings where data may be split into clusters, the sample complexity can be much lower by using triangle inequality.

The order in which $\{x_{i^*}\}$ are found follows their subscript index, which is randomly chosen and fixed before starting the algorithm. As described above, different orders in which $\{x_i\}$ are processed can affect the query complexity of our algorithm. The best order that minimizes the total number of queries made in general depends on the true distance values. Even if the oracle is noiseless, there are datasets where the pair (i, j) with the smallest $d_{i,j}$ must be queried within the first n queries in order to identify the NN-graph using the minimum number of queries. Since this requirement cannot be ensured by any algorithm that only has access to information via a distance oracle, it is not possible to achieve the minimum number of queries in such examples.

4.3 Adaptive gains via the triangle inequality

We highlight two settings where ANNTri provably achieves sample complexity better than $\mathcal{O}(n^2)$ independent of the order of the rounds. Consider a dataset containing c clusters of n/c points each as in Fig. 1a. Denote the m^{th} cluster as \mathcal{C}_m and suppose the distances between the points are such that

$$\{x_k : d_{i,k} < 6C_{\delta/n}(1) + 2d_{i,j}\} \subseteq \mathcal{C}_m \forall i, j \in \mathcal{C}_m. \quad (12)$$

The above condition is ensured if the distance between any two points belonging to different clusters is at least a (δ, n) -dependent constant plus twice the diameter of any cluster.

Theorem 4.5. *Consider a dataset of \sqrt{n} clusters which satisfy the condition in (12). Then ANNEasy learns the correct nearest neighbor graph of \mathcal{X} with probability at least $1 - \delta$ in*

$$\mathcal{O}\left(n^{3/2}\overline{\Delta}^{-2}\right) \quad (13)$$

queries where $\overline{\Delta}^{-2} := \frac{1}{n^{3/2}} \sum_{i=1}^{\sqrt{n}} \sum_{j,k \in \mathcal{C}_i} \log(n^2/(\delta\Delta_{j,k}))\Delta_{j,k}^{-2}$ is the average number of samples distances between points in the same cluster.

By contrast, random sampling requires $\mathcal{O}(n^2\Delta_{\min}^{-2})$ where $\Delta_{\min}^{-2} := \min_{j,k} \log(n^2/(\delta\Delta_{j,k}))\Delta_{j,k}^{-2} \geq \overline{\Delta}^{-2}$. In fact, the value in (11) can be even lower if unions of clusters also satisfy (12). In this case, the triangle inequality can be used to separate *groups* of clusters. For example, in Fig. 1b, if $\mathcal{C}_1 \cup \mathcal{C}_2$ and $\mathcal{C}_3 \cup \mathcal{C}_4$ satisfy (12) along with $\mathcal{C}_1, \dots, \mathcal{C}_4$, then the triangle inequality can separate $\mathcal{C}_1 \cup \mathcal{C}_2$ and $\mathcal{C}_3 \cup \mathcal{C}_4$. This process can be generalized to consider a dataset that can be split recursively into subclusters following a binary tree of k levels. At each level, the clusters are assumed to satisfy (12). We refer to such a dataset as *hierarchical* in (12).

Theorem 4.6. *Consider a dataset $\mathcal{X} = \cup_{i=1}^{n/\nu} \mathcal{C}_i$ of n/ν clusters of size $\nu = \mathcal{O}(\log(n))$ that is hierarchical in (12). Then ANNEasy learns the correct nearest neighbor graph of \mathcal{X} with probability at least $1 - \delta$ in*

$$\mathcal{O}\left(n \log(n) \overline{\Delta}^{-2}\right) \quad (14)$$

queries where $\overline{\Delta}^{-2} := \frac{1}{n\nu} \sum_{i=1}^{n/\nu} \sum_{j,k \in \mathcal{C}_i} \log(n^2/(\delta\Delta_{j,k}))\Delta_{j,k}^{-2}$ is the average number of samples distances between points in the same cluster.

Expression (14) matches known lower bounds of $\mathcal{O}(n \log(n))$ on the sample complexity for learning the nearest neighbor graph from noiseless samples [30], the additional penalty of $\overline{\Delta}^{-2}$ is due to the effect of noise in the samples. An easy way to see the lower bound is to consider the fact that there are $\mathcal{O}(n^{n-1})$ unique nearest neighbor graphs so any algorithm will require $\mathcal{O}(\log(n^{n-1})) = \mathcal{O}(n \log(n))$ bits of information to identify the correct one. In Appendix C, we state the sample complexity in the average case, as opposed to the high probability statements above. The analog of the cluster condition (12) there does not involve constants and is solely in terms of pairwise distances (c.f. (33)).

5 Experiments

Here we evaluate the performance of ANNTri on simulated and real data. To construct the tightest possible confidence bounds for SETri, we use the law of the iterated logarithm as in [18] with parameters $\epsilon = 0.7$ and $\delta = 0.1$. Our analysis bounds the number of queries made to the oracle. We visualize the performance by tracking the empirical *error rate* with the number of queries made per point. For a given point x_i , we say that a method makes an error at the t^{th} sample if it fails to return x_{i^*} as the nearest neighbor, that is, $x_{i^*} \neq \arg \min_j \hat{d}[i, j]$. Throughout, we will compare ANNTri against random sampling. Additionally, to highlight the effect of the triangle inequality, we will compare our method against the same active procedure, but ignoring triangle inequality bounds (referred to as ANN in plots). All baselines may reuse samples via symmetry as well. We plot all curves with 95% confidence regions shaded.

5.1 Simulated Experiments

We test the effectiveness of our method, we generate an embedding of 10 clusters of 10 points spread around a circle such that each cluster is separated by at least 10% of its diameter in \mathbb{R}^2 as in shown in Fig. 2a. We consider Gaussian noise with $\sigma = 0.1$. In Fig. 2b, we present average error rates of ANNTri, ANN, and Random plotted on a log scale. ANNTri quickly learns x_{i^*} and has lower error with 0 samples due to initial elimination by the triangle inequality. The error curves are averaged over 4000 repetitions. All rounds were capped at 10^5 samples for efficiency.

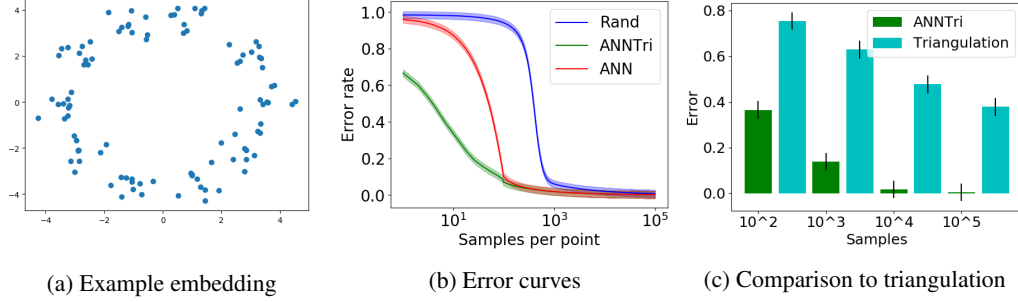


Figure 2: Comparison of ANNTri to ANN and Random for 10 clusters of 10 points separated by 10% of their diameter with $\sigma = 0.1$. ANNTri identifies clusters of nearby points more easily.

5.1.1 Comparison to triangulation

An alternative way a practitioner may use to obtain the nearest neighbor graph might be to estimate distances with respect to a few anchor points and then triangulate to learn the rest. [9] provide a comprehensive example, and we summarize in Appendix A.2 for completeness. The triangulation method is naïve for two reasons. First, it requires *much* stronger modeling assumptions than ANNTri—namely that the metric is Euclidean and the points are in a low-dimensional of known dimension. Forcing Euclidean structure can lead to unpredictable errors if the underlying metric might not be Euclidean, such as in data from human judgments. Second, this procedure may be more noise sensitive because it estimates squared distances. In the example in Section A.2, this leads to the additive noise being sub-exponential rather than subGaussian. In Fig. 2c, we show that even in a favorable setting where distances are truly sampled from a low-dimensional Euclidean embedding and pairwise distances between anchors are known exactly, triangulation still performs poorly compared to ANNTri. We consider the same 2-dimensional embedding of points as in Fig. 2a for a noise variance of $\sigma = 1$ and compare the ANNTri and triangulation for different numbers of samples.

5.2 Human judgment experiments

5.2.1 Setup

Here we consider the problem of learning from human judgments. For this experiment, we used a set \mathcal{X} of 85 images of shoes drawn from the UT Zappos50k dataset [32, 33] and seek to learn which shoes are most visually similar. To do this, we consider queries of the form “between i , j , and k , which two are most similar?”. We show example queries in Figs. 5a and 5b in the Appendix. Each query maps to a pair of triplet judgments of the form “is j or k more similar to i ?”. For instance, if i and j are chosen, then we may imply the judgments “ i is more similar to j than to k ” and “ j is more similar to i than to k ”. We therefore construct these queries from a set of triplets collected from participants on Mechanical Turk by [15]. The set contains multiple samples of all $85 \binom{84}{2}$ unique triples so that the probability of any triplet response can be estimated. We expect that i^* is most commonly selected as being more similar to i than any third point k . We take distance to correspond to the fraction of times that two images i, j are judged as being more similar to each other than a different pair in a triplet query (i, j, k) . Let $E_{i,k}^j$ be the event that the pair i, k are chosen as most similar amongst i, j , and k . Accordingly, we define the ‘distance’ between images i and j as

$$d_{i,j} := \mathbb{E}_{k \sim \text{Unif}(\mathcal{X} \setminus \{i,j\})} \mathbb{E}[\mathbb{1}_{E_{i,k}^j} | k]$$

where k is drawn uniformly from the remaining 83 images in $\mathcal{X} \setminus \{i, j\}$. For a fixed value of k ,

$$\mathbb{E}[\mathbb{1}_{E_{i,k}^j} | k] = \mathbb{P}(E_{i,k}^j | k) = \mathbb{P}(\text{“}i \text{ more similar to } j \text{ than to } k\text{”}) \mathbb{P}(\text{“}j \text{ more similar to } i \text{ than to } k\text{”}).$$

where the probabilities are the empirical probabilities of the associated triplets in the dataset. This distance is a quasi-metric on our dataset as it does not always satisfy the triangle inequality; but satisfies it with a multiplicative constant: $d_{i,j} \leq 1.47(d_{i,k} + d_{j,k}) \forall i, j, k$. Relaxing metrics to quasi-metrics has a rich history in the classical nearest neighbors literature [16, 29, 14], and ANNTri can be trivially modified to handle quasi-metrics. However, we empirically note that $< 1\%$ of the distances violate the ordinary triangle inequality here so we ignore this point in our evaluation.

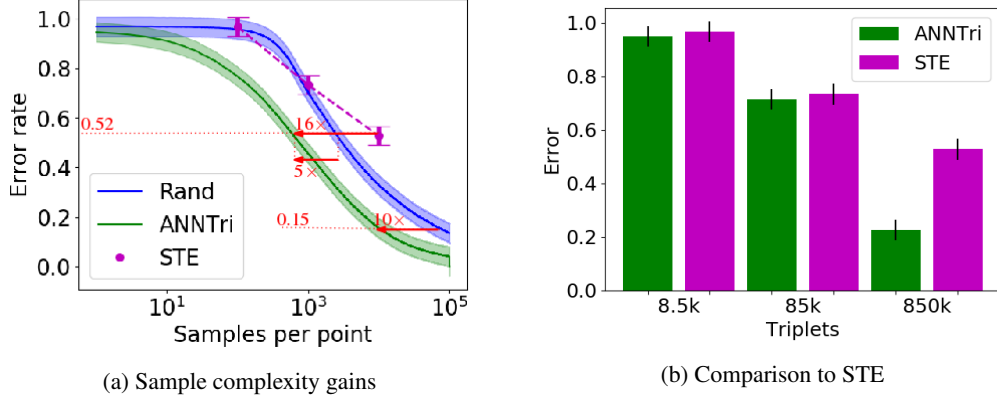


Figure 3: Performance of ANNTri on the Zappos dataset. ANNTri achieves superior performance over STE in identifying nearest neighbors and has 5 – 10x gains in sample efficiency over random.

5.2.2 Results

When ANNTri or any baseline queries $Q(i, j)$ from the oracle, we randomly sample a third point $k \in \mathcal{X} \setminus \{i, j\}$ and flip a coin with probability $\mathbb{P}(E_{i,k}^j)$. The resulting sample is an unbiased estimate of the distance between i and j . In Fig. 3a, we compare the error rate averaged over 1000 trials of ANNTri compared to Random and STE. We also plot associated gains in sample complexity by ANNTri. In particular, we see gains of 5 – 10x over random sampling, and gains up to 16x relative to ordinal embedding. ANNTri also shows 2x gains over ANN in sample complexity (see Fig. 6 in Appendix).

Additionally, a standard way of learning from triplet data is to perform ordinal embedding. With a learned embedding, the nearest neighbor graph may easily be computed. In Fig. 3b, we compare ANNTri against the state of the art STE algorithm [31] for estimating Euclidean embeddings from triplets, and select the embedding dimension of $d = 16$ via cross validation. To normalize the number of samples, we first perform ANNTri with a given max budget of samples and record the total number needed. Then we select a random set of triplets of the same size and learn an embedding in \mathbb{R}^{16} via STE. We compare both methods on the fraction of nearest neighbors predicted correctly. On the x axis, we show the total number of triplets given to each method. For small dataset sizes, there is little difference, however, for larger dataset sizes, ANNTri significantly outperforms STE. Given that ANNTri is active, it is reasonable to wonder if STE would perform better with an actively sampled dataset, such as [28]. Many of these methods are computationally intensive and lack empirical support [20], but we can embed using the full set of triplets to mitigate the effect of the subsampling procedure. Doing so, STE achieves 52% error, within the confidence bounds of the largest subsample shown in Fig. 3b. In particular, more data and more carefully selected datasets, may not correct for the bias induced by forcing Euclidean structure.

6 Conclusion

In this paper we solve the nearest neighbor graph problem by adaptively querying distances. Our method makes no assumptions beyond standard metric properties and is empirically shown to achieve sample complexity gains over passive sampling. In the case of clustered data, we show provable gains and achieve optimal rates in favorable settings. One interesting avenue for future work would be to specialize to the case of hyperbolic embeddings which naturally encode trees [8] and may be a more flexible way to describe hierarchical clusters as in Theorem 4.6. Implementations of ANNTri, ANN, and RANDOM can be found alongside a demo and summary slides at <https://github.com/blakemas/nngraph>.

Acknowledgments

The authors wish to thank Lalit Jain for many helpful discussions over the course of this work for which the paper is better and the reviewers for their helpful suggestions. This work was partially supported by AFOSR/AFRL grants FA8750-17-2-0262 and FA9550-18-1-0166.

References

- [1] Vivek Bagaria, Govinda M Kamath, Vasilis Ntranos, Martin J Zhang, and David Tse. Medoids in almost linear time via multi-armed bandits. *arXiv preprint arXiv:1711.00817*, 2017.
- [2] Vivek Bagaria, Govinda M Kamath, and David N Tse. Adaptive monte-carlo optimization. *arXiv preprint arXiv:1805.08321*, 2018.
- [3] Nitin Bhatia et al. Survey of nearest neighbor techniques. *arXiv preprint arXiv:1007.0085*, 2010.
- [4] Brandon M Booth, Tiantian Feng, Abhishek Jangalwa, and Shrikanth S Narayanan. Toward robust interpretable human movement pattern analysis in a workplace setting. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7630–7634. IEEE, 2019.
- [5] Justin Brickell, Inderjit S Dhillon, Suvrit Sra, and Joel A Tropp. The metric nearness problem. *SIAM Journal on Matrix Analysis and Applications*, 30(1):375–396, 2008.
- [6] Kenneth L Clarkson. Fast algorithms for the all nearest neighbors problem. In *24th Annual Symposium on Foundations of Computer Science (sfcs 1983)*, pages 226–232. IEEE, 1983.
- [7] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.
- [8] Andrej Cvetkovski and Mark Crovella. Hyperbolic embedding and routing for dynamic graphs. In *IEEE INFOCOM 2009*, pages 1647–1655. IEEE, 2009.
- [9] Brian Eriksson, Paul Barford, Joel Sommers, and Robert Nowak. A learning-based approach for ip geolocation. In *International Conference on Passive and Active Network Measurement*, pages 171–180. Springer, 2010.
- [10] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(Jun):1079–1105, 2006.
- [11] Robert W Floyd. Algorithm 97: shortest path. *Communications of the ACM*, 5(6):345, 1962.
- [12] Aurélien Garivier. Informational confidence bounds for self-normalized averages and applications. In *2013 IEEE Information Theory Workshop (ITW)*, pages 1–5. IEEE, 2013.
- [13] Anna C Gilbert and Lalit Jain. If it ain’t broke, don’t fix it: Sparse metric repair. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 612–619. IEEE, 2017.
- [14] Navin Goyal, Yury Lifshits, and Hinrich Schütze. Disorder inequality: a combinatorial approach to nearest neighbor search. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, pages 25–32. ACM, 2008.
- [15] Eric Heim, Matthew Berger, Lee Seversky, and Milos Hauskrecht. Active perceptual similarity modeling with auxiliary information. *arXiv preprint arXiv:1511.02254*, 2015.
- [16] Michael E Houle and Michael Nett. Rank-based similarity search: Reducing the dimensional dependence. *IEEE transactions on pattern analysis and machine intelligence*, 37(1):136–150, 2015.
- [17] Lalit Jain, Kevin G Jamieson, and Rob Nowak. Finite sample prediction and recovery bounds for ordinal embedding. In *Advances In Neural Information Processing Systems*, pages 2711–2719, 2016.
- [18] K. Jamieson and R. Nowak. Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *2014 48th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6, March 2014. doi: 10.1109/CISS.2014.6814096.

- [19] Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sebastien Bubeck. On finding the largest mean among many. *arXiv preprint arXiv:1306.3917*, 2013.
- [20] 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. In *Advances in Neural Information Processing Systems*, pages 2656–2664, 2015.
- [21] Emilie Kaufmann and Shivaram Kalyanakrishnan. Information complexity in bandit subset selection. In *Conference on Learning Theory*, pages 228–251, 2013.
- [22] Matthäus Kleindessner and Ulrike Von Luxburg. Lens depth function and k-relative neighborhood graph: versatile tools for ordinal data analysis. *The Journal of Machine Learning Research*, 18(1):1889–1940, 2017.
- [23] Joseph B Kruskal. Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29(2):115–129, 1964.
- [24] Daniel LeJeune, Richard G Baraniuk, and Reinhard Heckel. Adaptive estimation for approximate k-nearest-neighbor computations. *arXiv preprint arXiv:1902.09465*, 2019.
- [25] Jagan Sankaranarayanan, Hanan Samet, and Amitabh Varshney. A fast all nearest neighbor algorithm for applications involving large point-clouds. *Computers & Graphics*, 31(2):157–174, 2007.
- [26] Roger N Shepard. The analysis of proximities: multidimensional scaling with an unknown distance function. i. *Psychometrika*, 27(2):125–140, 1962.
- [27] Adish Singla, Sebastian Tschiatschek, and Andreas Krause. Actively learning hemimetrics with applications to eliciting user preferences. In *International Conference on Machine Learning*, pages 412–420, 2016.
- [28] Omer Tamuz, Ce Liu, Serge Belongie, Ohad Shamir, and Adam Tauman Kalai. Adaptively learning the crowd kernel. *arXiv preprint arXiv:1105.1033*, 2011.
- [29] Dominique Tschopp, Suhas Diggavi, Payam Delgosha, and Soheil Mohajer. Randomized algorithms for comparison-based search. In *Advances in Neural Information Processing Systems*, pages 2231–2239, 2011.
- [30] Pravin M Vaidya. Ano ($n \log n$) algorithm for the all-nearest-neighbors problem. *Discrete & Computational Geometry*, 4(2):101–115, 1989.
- [31] 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.
- [32] A. Yu and K. Grauman. Fine-grained visual comparisons with local learning. In *Computer Vision and Pattern Recognition (CVPR)*, Jun 2014.
- [33] A. Yu and K. Grauman. Semantic jitter: Dense supervision for visual comparisons via synthetic images. In *International Conference on Computer Vision (ICCV)*, Oct 2017.

Algorithm 3 ANNEasy

Require: n , procedure SEEasy, 4, confidence δ

- 1: Initialize \hat{d}, T as $n \times n$ matrices of zeros, U as $n \times n$ matrix where each entry is ∞ , L as $n \times n$ matrix where each entry is $-\infty$, NN as a length n array
 - 2: **for** $j = 1$ **to** n **do**
 - 3: NN[j] = SEEasy($j, \hat{d}, U, L, T, \xi = \delta/n$)
 - 4: **return** The nearest neighbor graph adjacency list NN
-

Algorithm 4 SEEasy

Require: index j , callable oracle $Q(\cdot, \cdot)$ (1), 4 $n \times n$ matrices: \hat{d}, U, L, T , confidence ξ

- 1: Initialize the active set $\mathcal{A}_j \leftarrow \{a \neq j : L[a, k] \leq 2U[j, k] \forall k \text{ and } L[a, j] < \min_k U[j, k]\}$
 - 2: **while** $|\mathcal{A}_j| > 1$ **do**
 - 3: **for all** $i \in \mathcal{A}_j$ such that $T[i, j] = \min_{k \in \mathcal{A}_j} T[i, k]$ **do** {only query points with fewest samples}
 - 4: Update $\hat{d}[i, j], \hat{d}[j, i] \leftarrow (\hat{d}[i, j] \cdot T[i, j] + Q(i, j)) / (T[i, j] + 1)$
 - 5: Update $T[i, j], T[j, i] \leftarrow T[i, j] + 1$
 - 6: Update $U[i, j], U[j, i] \leftarrow \hat{d}[i, j] + C_\xi(T[i, j])$
 - 7: Update $L[i, j], L[j, i] \leftarrow \hat{d}[i, j] - C_\xi(T[i, j])$
 - 8: Update $\mathcal{A}_j \leftarrow \{a \neq j : L[a, k] \leq 2U[j, k] \forall k \text{ and } L[a, j] < \min_k U[a, k]\}$
 - 9: **return** The index i for which $x_i \in \mathcal{A}_j$
-

Appendix

A Additional experimental results and details

A.1 Differences between ANNTri and ANNEasy

A.1.1 Pseudocode for ANNEasy and SEEasy

We begin by providing pseudocode for both ANNEasy and SEEasy as described in Section 4.1 in Algorithms 3 and 4.

A.1.2 Empirical differences in performance for ANNTri and ANNEasy

In Figure 4 we compare the empirical performance of ANNTri and ANNEasy. We compare their performance in the same setting as Figure 2a with 10 clusters of 10 points separated by their at least 10% of their diameter. The curves are averaged over 4000 independent trials and plotted with 95% confidence regions. As is indicated in the plot, ANNEasy has similar behavior as ANNTri, but achieves slightly worse performance.

A.2 Triangulation

In this section, we provide a brief review of triangulation to estimate Euclidean embeddings, similar to the presentation in [9]. The method is summarized as follows. Let \mathcal{X} be a set of n points in Euclidean d space and D be the associated Euclidean distance matrix where each entry is the square of the associated Euclidean distance. Let A be a set of anchor points. Without loss of generality, we take $A := \{x_1, \dots, x_{d+2}\}$. The $+2$ is to correct for the fact that Euclidean distance matrices have rank $d + 2$. Let $\mathbf{A} := D[1 : d + 2, 1 : d + 2]$ and $\mathbf{L} := D[1 : d + 2, 1 : n]$. Then it can easily be verified that $D = \mathbf{L} \mathbf{A}^{-1} \mathbf{L}^T$. To learn the entries in \mathbf{L} (as well as \mathbf{A}), sample the distance from each of the n points to the $d + 2$ anchors as many times as there is budget for and square the results. The empirical mean is a plugin estimator of the associated entry in \mathbf{L} and \mathbf{A} , and we take $\hat{\mathbf{L}}$ and $\hat{\mathbf{A}}$ to be their unbiased estimates. Therefore $\hat{D} := \hat{\mathbf{L}} \hat{\mathbf{A}}^{-1} \hat{\mathbf{L}}^T$ is an unbiased estimate of D . With \hat{D} , the nearest neighbor graph can easily be computed.




Figure 4: Comparison of error in identifying x_{i^*} ANNTri and the ANNEasy for 10 clusters of 10 points separated by 10% of their diameter with $\sigma = 0.1$.

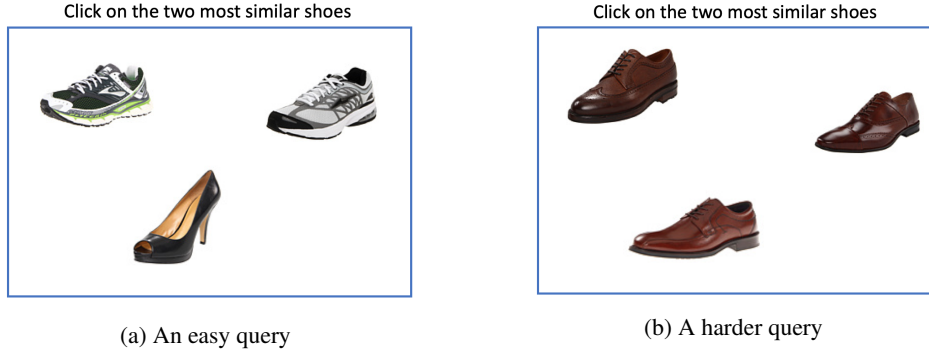


Figure 5: Two example zappos queries.

A.3 Additional experimental results for Zappos dataset

In Fig. 5 we show two example queries of the form “which pair are most similar of these three?”. Some queries are more straightforward whereas some are more subjective.

Additionally, in Fig. 6, we show the performance of ANNTri, ANN, and Random in identifying nearest neighbors from the Zappos data. In this setting, there is less of an advantage to using the triangle inequality due to the highly noisy and subjective nature of human judgments. Despite this, we still see a slight advantage to ANNTri over ANN. In particular, for moderate accuracy, there is a gain sample complexity of around 2x.

B Proofs and technical lemmas

B.1 Proof of Lemma 3.1

By symmetry for all $i < j$, we have existing samples of $Q(i, j)$ and $Q(i, k)$ and we use bounds based on these samples as well as past triangle inequality upper bounds on $d_{i,j}$ and $d_{i,k}$ due to $i_1 < i$ and $i_2 < i$ respectively. The upper bound is derived as follows:

$$d_{j,k} \leq d_{i,j} + d_{i,k} \leq \min\{U_{i,j}(t), U_{i,j}^{\triangle_{i_1}}(t)\} + \min\{U_{i,k}(t), U_{i,k}^{\triangle_{i_2}}(t)\} =: U_{j,k}^{\triangle_i}$$

Since we may form bounds based on all $i < j$ for which we have both samples of $Q(i, j)$ and $Q(i, k)$, we may optimize over i to get the tightest possible triangle inequality bounds on $d_{j,k}$.

Lower bounds are derived similarly. Again, intuitively, we may use past samples of both $Q(i, j)$ and $Q(i, k)$ and associated bounds to derive a lower bound on $d_{j,k}$. The form is slightly more complicated here since we have to worry about both upper and lower bounds on $d_{i,j}$ and $d_{i,k}$. These bounds may




Figure 6: Error rates for nearest neighbor identification on Zappos Data

either be from concentration bounds based on past samples directly or past triangle inequality upper and lower bounds on these distances due to points $i_1 - i_4 < i$.

$$\begin{aligned}
d_{j,k} &\geq |d_{i,j} - d_{i,k}| \\
&= \max\{d_{i,j}, d_{i,k}\} - \min\{d_{i,j}, d_{i,k}\} \\
&\geq (\max\{\max\{L_{i,j}(t), L_{i,j}^{\Delta_{i_1}}(t)\}, \max\{L_{i,k}(t), L_{i,k}^{\Delta_{i_2}}(t)\}\} \\
&\quad - \min\{\min\{U_{i,j}(t), U_{i,j}^{\Delta_{i_3}}(t)\}, \min\{U_{i,k}(t), U_{i,k}^{\Delta_{i_4}}(t)\}\})_+ \\
&= (\max\{L_{i,j}(t), L_{i,j}^{\Delta_{i_1}}(t), L_{i,k}(t), L_{i,k}^{\Delta_{i_2}}(t)\} \\
&\quad - \min\{U_{i,j}(t), U_{i,j}^{\Delta_{i_3}}(t), U_{i,k}(t), U_{i,k}^{\Delta_{i_4}}(t)\})_+
\end{aligned}$$

where $(s)_+ := \max\{s, 0\}$ and $i_1, i_2, i_3, i_4 < i$, (not necessarily unique) are chosen to optimize the bound. Similar to the upper bound, this holds with respect to any $i < j$ and we optimize over i . To ease presentation, let $UB'[i, j] := \min\{U_{i,j}, \min_{l < i} U_{i,j}^{\Delta_l}\}$ and $LB'[i, j] := \max\{L_{i,j}, \max_{l < i} L_{i,j}^{\Delta_l}\}$ be the tightest upper and lower bounds for $d_{i,j}$. For the lower bound, note that if the argument of $(\cdot)_+$ is negative, then any

$$\begin{aligned}
s &\in [\max\{LB'[i, j], LB'[i, k]\}, \min\{UB'[i, j], UB'[i, k]\}] \\
&= [LB'[i, j], UB'[i, j]] \cap [LB'[i, k], UB'[i, k]] \neq \emptyset
\end{aligned}$$

can be the value of both $d_{i,j}$ and $d_{j,k}$ as it lies in both their confidence intervals. Then points x_j, x_k can possibly be at the same location in the metric space, in which case $d_{j,k} = 0$. On the other hand if the RHS is positive, then x_j and x_k cannot be at the same location as $d_{i,j} \neq d_{i,k}$. In fact, the smallest possible value for $d_{j,k}$ occurs if x_i, x_j, x_k are collinear. This can be seen to be true from Figure 7. We finish with a quick lemma noting what can and cannot be eliminated via the triangle inequality.

Lemma B.1. *Conditioned on the good event that all bounds are correct at all times, the triangle inequality cannot be used to separate the two closest points to any given third point.*

Proof. Consider finding x_{i^*} . Let $d_{i,i^*} \leq d_{i,j} \leq d_{i,k} \forall k \neq i^*, j$. By the triangle inequality, $d_{i,i^*} \leq d_{i,j} + d_{j,i^*}$. Clearly, the RHS is no smaller than $d_{i,j}$. Since we are conditioning on all bounds being correct at all times, no upper bound on d_{i,i^*} from the triangle inequality can ever be smaller than $d_{i,j}$. Rearranging the inequality, we see that $d_{i,i^*} - d_{j,i^*} \leq d_{i,j}$. The LHS is no larger than d_{i,i^*} , and d_{i,i^*}




Figure 7: Pictorial justification for the lower bound in (6). True positions of points i, j, k are shown along with the upper and lower bounds for $d_{i,j}, d_{i,k}$ that are known to the algorithm. If the angle θ between \vec{ij} and \vec{ik} is known, the blue segment shows the lowest possible value for $d_{j,k}$ based on the bounds. The orange segment is the value in the RHS of (6). Without any information about θ , the three points could be collinear, in which case $d_{j,k}$ could equal the length of the orange segment.

is the only distance wrt x_i that is smaller than $d_{i,j}$ by assumption. Therefore, no lower bound on $d_{i,j}$ due to the triangle inequality is greater than $d_{i,i^*} < d_{i,j}$. \square

B.1.1 Helper Lemmas

Lemma B.2. Let $t \in \mathbb{N}$ index the rounds of the procedure *SETri* in finding x_{i^*} . Suppose all confidence intervals are valid, i.e., (3) is true. Then $\forall j \neq i$ and all t ,

$$L_{i,j}(t) \geq d_{i,j} - 2C_{\delta/n}(T_{i,j}(t)) \quad \text{and} \quad U_{i,j}(t) \leq d_{i,j} + 2C_{\delta/n}(T_{i,j}(t)). \quad (15)$$

Proof. If the good event (3) is true then for any pair (i, j) and time t we have

$$\hat{d}_{i,j}(t) < d_{i,j} + C_{\delta/n}(T_{i,j}(t)) \implies U_{i,j}(t) := \hat{d}_{i,j}(t) + C_{\delta/n}(T_{i,j}(t)) \leq d_{i,j} + 2C_{\delta/n}(T_{i,j}(t)).$$

A similar calculation can be done for $L_{i,j}(t)$ as well. \square

Lemma B.3. Let $j > i$, and let t_j be the time when x_j is last sampled in the i^{th} round and equivalently for t_k . Assume without loss of generality that $d_{i,j} < d_{i,k}$. If $d_{i,j}$ and $d_{i,k}$ are such that

$$4C_{\delta/n}(T_{i,j}(t_j)) + 2C_{\delta/n}(T_{i,k}(t_k)) \leq d_{i,k} - 2d_{i,j} \quad (16)$$

then *SETri* can eliminate $d_{j,k}$ without sampling it, i.e., $x_k \notin \mathcal{A}_j(0)$.

Proof. Focusing on the number of $Q(i, j)$ queries, we have that

$$2U_{i,j}(t_j) = 2(\hat{d}_{i,j}(t_j) + C_{\delta/n}(T_{i,j}(t_j))) \leq 2(d_{i,j} + 2C_{\delta/n}(T_{i,j}(t_j))), \quad (17)$$

the inequality in (17) is due to Lemma B.2, and using the number of $Q(i, k)$ queries,

$$L_{i,k}(t_k) \geq \hat{d}_{i,k}(t_k) - C_{\delta/n}(T_{i,k}(t_k)) \geq d_{i,k} - 2C_{\delta/n}(T_{i,k}(t_k)). \quad (18)$$

The first inequality in (18) is because if $k < j$ then there may have been more $Q(k, i)$ queries beyond the t_k number of $Q(i, k)$ queries made while finding x_{i^*} . Rearranging the equation in the Lemma statement,

$$2d_{i,j} + 4C_{\delta/n}(T_{i,j}(t_j)) \leq d_{i,k} - 2C_{\delta/n}(T_{i,k}(t_k)),$$

which implies that $2U_{i,j} \leq L_{i,k}$ from (17), (18). Hence from Lemma 4.2 $x_k \notin \mathcal{A}_j(0)$. \square

Lemma B.4. *There exists a dataset \mathcal{P} containing 2ν points such that for all $x_p \in \mathcal{P}$ and $\alpha > 0$ the set of suboptimality gaps $\Delta_{p,p'}$ is*

$$\left\{ 1 - \left(\frac{s-1}{\nu-1} \right)^\alpha : s \in \{1, 2, \dots, \nu-1\} \right\}. \quad (19)$$

Proof. Note that there are $\nu-1$ values given in (19) while there are $2\nu-2$ points in the cluster, excluding x_p and x_{p^*} . Each value in (19) is the suboptimality gap for two distinct points in $\mathcal{P} \setminus \{x_p, x_{p^*}\}$. We can construct such a dataset \mathcal{P} in the following manner.

We index these points as $p, p_1, p_2, \dots, p_{2\nu-1}$. Suppose the pairwise distance values are such that

$$d_{p,p_1} > d_{p,p_2} > \dots > d_{p,p_{\nu-1}} > d_{p,p_\nu} =: d_{p,p^*}, \text{ and } d_{p,p_{\nu+1}} < d_{p,p_{\nu+2}} < \dots < d_{p,p_{2\nu-1}} \text{ such that} \\ \forall s \in \{1, 2, \dots, \nu-1\} \text{ we have that } d_{p,p_{\nu-s}} = d_{p,p_{\nu+s}} \implies d_{p,p_i} = d_{p,p_{2\nu-i}}. \quad (20)$$

We can then construct a $2\nu \times 2\nu$ distance matrix D in the following manner. The first row of D is

$$D[0, :] := [0 \quad d_{p,p_1} \quad d_{p,p_2} \quad \dots \quad d_{p,p_{\nu-1}} \quad d_{p,p^*} \quad d_{p,p_{\nu+1}} \quad \dots \quad d_{p,p_{2\nu-2}} \quad d_{p,p_{2\nu-1}}].$$

The i th row of D is obtained by carrying out i circular shifts on the initial row $D[0, :]$ shown above. Thus D is a circulant matrix and we can see $D[i, j]$ and $D[j, i]$ to be as follows.

$$D[i, j] = \begin{cases} d_{p,p_{j-i}} & \text{if } j > i \\ d_{p,p_{2\nu-(i-j)}} & \text{if } j < i, \end{cases} \quad \text{and} \quad D[j, i] = \begin{cases} d_{p,p_{i-j}} & \text{if } i > j \\ d_{p,p_{2\nu-(j-i)}} & \text{if } i < j. \end{cases}$$

Then using (20) we have that $D[i, j] = D[j, i]$ for all $i \neq j$ and the diagonal entries are all 0. Thus D is symmetric. In addition, the distance values of the points to any point in the cluster take the same set of values. Suppose $d_{p,p^*} =: r > 0$ and $d_{p,p_1} = 2r$. Choose an $\alpha > 0$ and let

$$d_{p,p_{2\nu-i}} = d_{p,p_i} := r \left(2 - \left(\frac{s-1}{\nu-1} \right)^\alpha \right), \quad \forall s \in \{1, 2, \dots, \nu-1\}.$$

Then $D[i, j] \leq D[i, k] + D[k, j]$ for any three distinct i, j, k as the sum of any two elements is greater than $2r$, which is the largest element in D . Thus the distance values in D satisfy the triangle inequality, and D is a valid distance matrix. The suboptimality gaps for any point in the cluster is $\Delta_{p,p_i} = d_{p,p_i} - d_{p,p^*} = r(1 - ((i-1)/(\nu-1))^\alpha)$, choosing $r = 1$ finishes the required construction. \square

B.1.2 Proof of Theorem 4.1

Proof. ANNTri makes an error in finding the nearest neighbor for some point with probability $\mathbb{P}(\text{SETri is wrong for some } x_j, j \in \{1, 2, \dots, n\})$. We show that probability is at most $n\xi = \delta$, where the confidence level ξ for each execution of SETri is set to be δ/n . We use induction on $s \in \mathbb{N}$ to obtain that

$$\mathbb{P}(\forall j \in \{1, 2, \dots, s\}, k \neq j, \max\{L_{j,k}(t), L_{j,k}^\Delta(t)\} \leq d_{j,k} \leq \min\{U_{j,k}(t), U_{j,k}^\Delta(t)\}) \geq 1 - s\xi. \quad (21)$$

Consider the base case, i.e., point x_1 . From the initialization of ANNTri 1, $\min\{U_{1,k}(t), U_{1,k}^\Delta(t)\} = U_{1,k}(t)$, $\min\{L_{1,k}(t), L_{1,k}^\Delta(t)\} = L_{1,k}(t)$ for all $k \neq 1$. Using (3) we have $L_{1,k}(t) \leq d_{1,k} \leq U_{1,k}(t)$ with probability $1 - \delta/n$, and since ξ is δ/n the base case is true. Assume the hypothesis (21) is true for some s . We show that it is true for $s+1$ as well. We can bound the error event as follows.

$$\begin{aligned} & \mathbb{P}(\exists j \in \{1, \dots, s+1\}, k \neq j : d_{j,k} \notin [\max\{L_{j,k}(t), L_{j,k}^\Delta(t)\}, \min\{U_{j,k}(t), U_{j,k}^\Delta(t)\}]) \\ &= \mathbb{P}(\exists j \in \{1, \dots, s\}, k \neq j : d_{j,k} \notin [\max\{L_{j,k}(t), L_{j,k}^\Delta(t)\}, \min\{U_{j,k}(t), U_{j,k}^\Delta(t)\}]) \\ & \quad + \mathbb{P}\left(\{k \neq s+1 : d_{s+1,k} \notin [\max\{L_{s+1,k}(t), L_{s+1,k}^\Delta(t)\}, \min\{U_{s+1,k}(t), U_{s+1,k}^\Delta(t)\}]\} \right. \\ & \quad \left. \cap \{\forall j \in \{1, 2, \dots, s\}, k \neq j, \max\{L_{j,k}(t), L_{j,k}^\Delta(t)\} \leq d_{j,k} \leq \min\{U_{j,k}(t), U_{j,k}^\Delta(t)\}\} \right) \end{aligned} \quad (22)$$

From (21) the first summand in the RHS of (22) is at most $s\xi$. In the event corresponding to the second term, all the bounds used by SETri for $d_{j,k}, j \leq s, k \neq j$ are correct. Since $U_{s+1,\cdot}^\Delta$ and $L_{s+1,\cdot}^\Delta$ are both deterministically obtained (see (7), (8)) from them, they are correct as well. Thus

$$\begin{aligned} & \mathbb{P}(\max\{L_{s+1,k}(t), L_{s+1,k}^\Delta(t)\} \leq d_{s+1,k} \leq \min\{U_{s+1,k}(t), U_{s+1,k}^\Delta(t)\}) \\ &= \mathbb{P}(L_{s+1,k}(t) \leq d_{s+1,k} \leq U_{s+1,k}(t)) \geq 1 - \xi. \end{aligned}$$

Hence the second summand in the RHS of (22) is at most ξ . This proves (21) for $s+1$ and completes the induction.

Thus with probability $1 - n\xi = 1 - \delta$, the bounds obtained by SETri for finding $x_{j^*}, j \in \{1, \dots, n\}$ are all correct. We show that ANNTri correctly finds all nearest neighbors if the bounds are correct. For if not, suppose SETri returns the wrong nearest neighbor of x_j which happens only if x_{j^*} is not the last point in the active set. $x_{j^*} \notin \mathcal{A}$ because some other point $x_k \in \mathcal{A}$ eliminates it. Then $d_{j,k} < \min\{U_{j,k}, U_{j,k}^\Delta\} < \max\{L_{j,j^*}, L_{j,j^*}^\Delta\} < d_{j,j^*}$, which contradicts the fact that j^* is the nearest neighbor. \square

B.1.3 Proof of Lemma 4.3

Proof. Consider a point $x_i, i < j$ which satisfies the first part of (10). If $x_j \in \mathcal{A}_i(0)$ and $x_k \in \mathcal{A}_i(0)$, then neither x_j and x_k were eliminated without sampling when SEEasy was called for x_i and hence $T_{i,j} \geq 1$ and $T_{i,k} \geq 1$. Then we have that

$$4C_{\delta/n}(T_{i,j}(t_j)) + 2C_{\delta/n}(T_{i,k}(t_k)) \leq 6C_{\delta/n}(1) \leq d_{i,k} - 2d_{i,j}$$

and $x_k \notin \mathcal{A}_j(0)$ by Lemma B.3. The second part of (10) ensures that $\{x_j, x_k\} \subseteq \mathcal{A}_i(0)$ as shown next. The points eliminated from being the nearest neighbor of x_i using triangle inequality are $\mathcal{A}_i(0)^c = \cup_{m < i} \{\ell : 2U_{m,i} < L_{m,\ell}\}$. If the bounds obtained by SEEasy for all $m < i$ are correct,

$$\{\ell : 2U_{m,i} < L_{m,\ell}\} \subseteq \{\ell : 2d_{m,i} < d_{m,\ell}\} \implies \mathcal{A}_i(0)^c \subseteq \cup_{m < i} \{\ell : 2d_{m,i} < d_{m,\ell}\}.$$

Hence if the second condition of (10) is satisfied, then $\{j, k\} \subseteq \mathcal{A}_i(0)$ and we are done. \square

B.1.4 Proof of Theorem 4.4

Proof. Let x_j be the point on which SEEasy is called. Consider the case $j < k$. If $\mathbb{1}_{[A_{j,k}]} = 0$ then $x_k \notin \mathcal{A}_j(0)$ and no $Q(j, k)$ queries are made. Otherwise, x_k can be in the active set and from (4) at most $H_{j,k}$ samples of $d_{j,k}$ are taken. Now consider the case $k < j$. Samples of $d_{j,k}$ are only queried if $x_k \in \mathcal{A}_j(0)$. If $x_j \notin \mathcal{A}_k(0)$, i.e., x_j was eliminated when SEEasy was called for x_k then no $Q(k, j)$ queries made at that round. Again from (4) at most $H_{j,k}$ samples of $d_{j,k}$ are taken by SEEasy while finding x_{j^*} . If however $\mathbb{1}_{[A_{k,j}]} = 1$, then $Q(k, j)$ queries were made while finding x_{k^*} and let the number of those samples be $\#Q(k, j)$. Because of the sampling procedure of SEEasy, at most $(H_{j,k} - \#Q(k, j))_+$ queries are made for $d_{j,k}$. The total number of $Q(j, k)$ and $Q(k, j)$ queries is $\max\{H_{j,k}, \#Q(k, j)\}$, and since $\#Q(k, j) \leq H_{k,j}$, we get the result. \square

B.2 Details for Section 4.3

In this section, we consider a case where ANNTri achieves complexity that scales like $O(n^{1.5})$ as well as $O(n \log(n))$, the known optimal rate for the all nearest neighbors problem for noiseless data. To do this, we first prove a lemma about the complexity of learning with clustered data. In particular, we show that if the data comes from two well separated clusters, then the complexity of learning the nearest neighbor graph can be bounded as the complexity of learning the nearest neighbors of two points looking at the full dataset and the complexity of learning the remaining nearest neighbors graphs on each of the clusters.

Lemma B.5. *Consider $\mathcal{X} = \mathcal{C}_1 \cup \mathcal{C}_2$ where \mathcal{C}_1 and \mathcal{C}_2 both satisfy 12 for all i, j . Then ANNEasy learns the nearest neighbor graph of \mathcal{X} with probability at least $1 - \delta$ in at most*

$$O(|\mathcal{C}_1| + |\mathcal{C}_2| + \mathcal{H}_{\mathcal{C}_1} + \mathcal{H}_{\mathcal{C}_2}) \quad (23)$$

samples independent of the order in which it finds nearest neighbors where $\mathcal{H}_{\mathcal{C}_i}$ denotes the complexity of learning the nearest neighbor graph of cluster \mathcal{C}_i as bounded by B.6.

The above lemma implies that for the first point explored in each cluster, it is necessary to look at all other points in the dataset, but for all other points, it is only necessary to search within that point's respective cluster.

Proof. Choose a random order of points and fix it. Without loss of generality, we assume that $x_1 \in \mathcal{C}_1$. Let j_2 be the first point visited in \mathcal{C}_2 . Throughout, we will ignore reused samples since they only contribute at most a factor of 2 to the sample complexity as can be seen by Theorems 4.4 and B.6 and we seek an upper bound. Via standard analysis for successive elimination, x_{1^*} can be found in $\mathcal{O}\left(\sum_{j=2}^n H_{1,j}\right) = |\mathcal{C}_2| + \mathcal{O}\left(\sum_{j \in \mathcal{C}_1 \setminus \{x_1\}} H_{1,j}\right)$ samples with probability at least $1 - \delta/n$. For all $i = 2, \dots, j_2 - 1$,

$$\mathcal{A}_i(0)^c \supset \{\mathcal{A}_1(0) \cap \{k : d_{i,k} \geq 6d_{i,j} - 3d_{1,1^*}\}\} \supset \{\mathcal{X} \setminus \{x_1\} \cap \mathcal{C}_2\} = \mathcal{C}_2$$

which implies that $x_{i^*} \in \mathcal{C}_1$. For x_{j_2} we may trivially say that $\mathcal{A}_{j_2}(0)^c \supset \{j_2\}$ so $x_{j_2^*}$ can be learned in $\mathcal{O}\left(\sum_{l \neq j_2} H_{i,j}\right) = |\mathcal{C}_1| + \mathcal{O}\left(\sum_{j \in \mathcal{C}_2 \setminus \{x_{j_2}\}} H_{j_2,j}\right)$ samples with probability at least $1 - \delta/n$. We conclude by showing that for all remaining x_i , if $x_i \in \mathcal{C}_1$, then $\mathcal{A}_i(0) \subset \mathcal{C}_1$ and if $x_i \in \mathcal{C}_2$, then $\mathcal{A}_i(0) \subset \mathcal{C}_2$. Consider the case that $x_1 \in \mathcal{C}_1$. Suppose that $\exists x_j \in \mathcal{A}_i(0) \cap \mathcal{C}_2$. Then $2U_{1,i} > L_{1,j}$.

$$U_{i,1} = \hat{d}_{1,i} + C_{\delta/n}(T_{1,i}) \leq d_{1,i} + 2C_{\delta/n}(T_{1,i}) = d_{1,i} + 2C_{\delta/n}(1)$$

where the first inequality holds by B.2. Similarly,

$$L_{1,j} = \hat{d}_{1,j} - C_{\delta/n}(T_{1,j}) \geq d_{1,j} - 2C_{\delta/n}(T_{1,j}) \geq d_{1,j} - 2C_{\delta/n}(1)$$

Then $2(d_{1,i} + 2C_{\delta/n}(1)) \geq 2U_{1,i} > L_{1,j} \geq d_{1,j} - 2C_{\delta/n}(1) \implies d_{1,j} < 2d_{1,i} + 6C_{\delta/n}(1) \implies j \in \mathcal{C}_1$ which is a contradiction. A similar proof holds for $x_i \in \mathcal{C}_2$. It remains to argue that j_2 can be any number between 2 (by assumption that $x_1 \in \mathcal{C}_1$) and $|\mathcal{C}_1| + 1$ without affecting the bound on the complexity. By the assumption that \mathcal{C}_1 and \mathcal{C}_2 satisfy 12, out of cluster points can be eliminated in a single sample. Therefore, for any j_2 , $\sum_{l \in \mathcal{C}_1} H_{j_2,l} = |\mathcal{C}_1|$. Then we have that the total complexity is $\mathcal{O}(|\mathcal{C}_1| + |\mathcal{C}_2| + \mathcal{H}_{\mathcal{C}_1} + \mathcal{H}_{\mathcal{C}_2}) \forall j_2$. Since we have considered general orders of finding each nearest neighbor, we are done. \square

B.2.1 Proof of Theorem 4.5

Proof. By assumption, the dataset $\mathcal{X} = \cup_{i=1}^c \mathcal{C}_i$ with each cluster satisfies Equation 12. Therefore, for all m , $\mathcal{X} = \mathcal{C}_m \cup (\cup_{j \neq m} \mathcal{C}_j)$. By applying Lemma B.5, iteratively, we bound the complexity in terms of the complexity of learning the nearest neighbor graph of \mathcal{C}_m , the complexity of learning the nearest neighbor graph of $\cup_{j \neq m} \mathcal{C}_j$, and an additive penalty of n which accounts for the samples taken between the two. Since \mathcal{X} is a union of c clusters, this process may repeat c times. Therefore the total complexity can be bounded as

$$\mathcal{O}\left(cn + \sum_{i=1}^c \sum_{j,k \in \mathcal{C}_i} H_{j,k}\right)$$

Taking $c = \sqrt{n}$, we see that the above sum is $\mathcal{O}\left(n^{1.5} \overline{\Delta}^{-2}\right)$ where $\overline{\Delta}^{-2} = \frac{1}{c*n} \sum_{i=1}^c \sum_{j,k \in \mathcal{C}_i} H_{j,k}$ is the average number of times intra-cluster distances are sampled. By contrast, the complexity for random sampling is $\mathcal{O}(n^2 \Delta_{\min}^{-2})$ where $\Delta_{\min}^{-2} := \min_{j,k} H_{j,k}$. Comparing the two, we see that the latter is larger by at least a factor of $\mathcal{O}(\sqrt{n})$. \square

B.2.2 Proof of Lemma 4.6

Next we use Lemma B.5 to show that for datasets such that the clusters nest, we can achieve complexity scaling in $\mathcal{O}(n \log(n) \overline{\Delta}^{-2})$. In particular, we will recursively apply Lemma B.5 to show that clusters can be broken into subclusters and initial active sets shrink in diadic splits.

Proof. Before we prove the theorem, we begin by introducing some notation to make this proof concise. Recall that we have assumed that \mathcal{X} can be written as a hierarchy of clusters and sub clusters that form a balanced tree. We will denote the root of the tree with the full dataset as the 0^{th} level and

each split in that level will be indexed by $i = 1, \dots, 2^\ell$ where $\ell = 0, \dots, \log(n/\nu) - 1$ denotes the level. For notational ease, we take $\mathcal{C}_{0,1} \equiv \mathcal{X}$. $\mathcal{C}_{\ell,i}$ denotes the i^{th} cluster at the ℓ^{th} level of the tree which may be split into subclusters if $\ell < \log(n/\nu) - 1$. The idea will be to traverse the tree and split clusters into subclusters while keeping track of the number of between cluster samples that were necessary due to the bound in Lemma B.5. We let $\mathcal{H}_{\mathcal{C}_{\ell,i}}$ denote complexity of learning the nearest neighbor graph of $\mathcal{C}_{\ell,i}$.

Randomize the order and fix it. We will proceed by recursively applying Lemma B.5 to bound the complexity of learning the full nearest neighbor graph of a cluster in terms of learning it for each subcluster plus an additive penalty. By Lemma B.5 the complexity of finding the nearest neighbor graph of \mathcal{X} can be upper bounded as

$$\mathcal{O}(|\mathcal{C}_{1,1}| + |\mathcal{C}_{1,2}| + \mathcal{H}_{\mathcal{C}_{1,1}} + \mathcal{H}_{\mathcal{C}_{1,2}}) = \mathcal{O}(n + \mathcal{H}_{\mathcal{C}_{1,1}} + \mathcal{H}_{\mathcal{C}_{1,2}}).$$

We may again apply Lemma B.5 to $\mathcal{C}_{1,1}$ and $\mathcal{C}_{1,2}$ to bound their complexities as $\mathcal{O}(\frac{n}{2} + \mathcal{H}_{\mathcal{C}_{2,1}} + \mathcal{H}_{\mathcal{C}_{2,2}})$ and $\mathcal{O}(\frac{n}{2} + \mathcal{H}_{\mathcal{C}_{2,3}} + \mathcal{H}_{\mathcal{C}_{2,4}})$ respectively where $\mathcal{C}_{1,1} = \mathcal{C}_{2,1} \cup \mathcal{C}_{2,2}$ and $\mathcal{C}_{1,2} = \mathcal{C}_{2,3} \cup \mathcal{C}_{2,4}$. Therefore, similar to the above level, the total additive penalty for samples between clusters is n for the level. We may continue this process of splitting and paying the penalty of $n/2^\ell \times 2^\ell$ between cluster samples down to the bottom level $\ell = \log(n/\nu)$ with clusters of size ν .

Therefore, we may write the complexity as

$$\mathcal{O}\left(n \log\left(\frac{n}{\nu}\right) + \sum_{i=1}^{n/\nu} \sum_{j,k \in \mathcal{C}_{\log(n/\nu),i}} H_{j,k}\right). \quad (24)$$

Ignoring logarithmic factors, each complexity term $H_{j,k}$ is of the order $\mathcal{O}(\Delta_{j,k}^{-2})$. Therefore the entire summation is of the order

$$\mathcal{O}\left(n \log\left(\frac{n}{\nu}\right) + n\nu \overline{\Delta^{-2}}\right)$$

where $\overline{\Delta^{-2}} := \frac{1}{n\nu} \sum_{i=1}^{n/\nu} \sum_{j,k \in \mathcal{C}_i} \log(n^2/(\delta\Delta_{j,k}))\Delta_{j,k}^{-2}$ is the average complexity. Recalling that $\nu = \mathcal{O}(\log(n))$, we are done. \square

B.3 Sample complexity without using triangle inequality

Theorem B.6. *With probability $1 - \delta$, the number of oracle queries made by ANNTri and ANNEasy if all triangle bounds are ignored is at most*

$$\mathcal{O}\left(\sum_{i < j} \max\left\{\frac{\log(n^2/(\delta\Delta_{i,j}))}{\Delta_{i,j}^2}, \frac{\log(n^2/(\delta\Delta_{j,i}))}{\Delta_{j,i}^2}\right\}\right). \quad (25)$$

In the experiments, the process of using ANNTri and ignoring triangle inequality bounds is referred to as ANN.

Proof. In the case that triangle bounds are ignored, ANNTri and ANNEasy are the same. Consider the i^{th} round where we seek to identify x_{i^*} with probability $1 - \delta/n$. ANNTri has found x_{ℓ^*} for all $\ell < i$, in particular, it has evaluated $\hat{d}_{\ell,i}, U_{\ell,i}, L_{\ell,i}$. For every $x_j \neq x_{i^*}, x_j \in \mathcal{A}_i(0)$, we can bound the number of $Q(i, j)$ queries in the following manner. Suppose $j > i$ and $i^* > i$, so that at the beginning of the i^{th} round we have that $T_{i,j} = T_{i,i^*} = 0$. From (3), with probability $1 - \delta/n$, x_{i^*} is the last point in the active set. The point x_j is eliminated from the active set at time t_j if the following is true.

$$\begin{aligned} U_{i,i^*}(t_j) &\stackrel{(a)}{\leq} d_{i,i^*} + 2C_{\delta/n}(T_{i^*}(t_j)) < d_{i,j} - 2C_{\delta/n}(T_j(t_j)) \stackrel{(b)}{\leq} L_{i,j}(t_j), \\ &\implies 4C_{\delta/n}(t_j) < d_{i,j} - d_{i,i^*} = \Delta_{i,j}. \end{aligned} \quad (26)$$

Inequalities (a), (b) are due to Lemma B.2, and the fact that if j is eliminated at time t_j , then $T_{i,j}(t_j) = t_j$. From the property of the $C_{\delta/n}(\cdot)$ function, (26) is ensured when the number of samples of $d_{i,j}$ is

$$t_j \leq \left\lceil \kappa \frac{\log(n^2/(\delta\Delta_{i,j}/4))}{(\Delta_{i,j}/4)^2} \right\rceil.$$

We now consider the cases when at least one of i^*, j are less than i .

$i^* > i, j < i$: In this case, at the beginning of the i^{th} round $T_{i,j}$ is equal to the number of $Q(j, i)$ queries made (denoted as $\#Q(j, i)$) while finding x_{j^*} :

$$\#Q(j, i) \leq \left\lceil \kappa \frac{\log(n^2/(\delta\Delta_{j,i}/4))}{(\Delta_{j,i}/4)^2} \right\rceil.$$

If $\#Q(j, i) > t_j$, then no further $Q(i, j)$ queries are made in the i^{th} round, as argued next. Because the sampling procedure of SETri queries all points who have the minimum number of samples at current time, if a query $Q(i, j)$ is made at time $t + 1$, that implies $T_{i,i^*}(t) = \#Q(j, i)$. But then j is not in the active set at time t as

$$U_{i,i^*}(\#Q(j, i)) < U_{i,i^*}(t_j) < d_{i,j} - 2C_{\delta/n}(t_j) < d_{i,j} - 2C_{\delta/n}(\#Q(j, i)) = L_{i,j}(\#Q(j, i))$$

and hence $Q(i, j)$ is not made. If $\#Q(j, i) < t_j$, then x_j is eliminated when $t_j - \#Q(j, i)$ more samples of $d_{i,j}$ have been queried. Thus the total number of samples of $d_{i,j}$ is at most $\max\{t_j, \#Q(j, i)\}$.

The other two cases of 1) $i^* < i, j > i$, and 2) $i^* < i, j < i$ can be handled similarly. \square

C Average case performance of ANNEasy

We can obtain a different expression for the number of oracle queries if all the random quantities during a run of the algorithm take their expected values. In particular, Lemma 4.3 can be relaxed to the following.

Lemma C.1. *If all bounds obtained by SEEasy are correct and all the random quantities take their expected values, then for some $i < j$ such that $x_j \neq x_{i^*} \neq x_k$ if we have that*

$$d_{i,k} > 6d_{i,j} - 3d_{i,i^*}, \quad \text{and} \quad \{j, k\} \cap (\cup_{m < i} \{\ell : 2d_{m,i} < d_{m,\ell}\}) = \emptyset, \quad (27)$$

then $2U_{i,j} < L_{i,k}$ and hence $x_k \notin \mathcal{A}_j(0)$.

Proof. In the good event, the point x_{i^*} is the last element in the active set \mathcal{A}_i and points x_j, x_k have been eliminated from \mathcal{A}_i at some prior times t_j, t_k respectively. Both $t_j > 0$ and $t_k > 0$ as $\{x_j, x_k\} \subset \mathcal{A}_i(0)$ is ensured by the second part of the condition, as shown in the proof of Lemma 4.3. At time t_j , we have that

$$\min_{\ell} \hat{d}_{i,\ell} + C_{\delta/n}(t_j) \leq \min_{\ell} U_{i,\ell} \leq L_{i,j} \leq \hat{d}_{i,j} - C_{\delta/n}(t_j). \quad (28)$$

If all the random quantities take their expected values, then $\hat{d}_{i,\ell} = d_{i,\ell} \forall \ell \neq i$ and we have that

$$d_{i,i^*} + C_{\delta/n}(t_j) \leq d_{i,j} - C_{\delta/n}(t_j) \implies C_{\delta/n}(t_j) \leq \Delta_{i,j}/2. \quad (29)$$

Under the assumption, $\hat{d}_{i,j} = d_{i,j}$ and using the definition of its upper and lower confidence bounds, we get that $\mathbb{E}[L_{i,j}] \geq d_{i,j} - \Delta_{i,j}/2$ and $\mathbb{E}[U_{i,j}] \leq d_{i,j} + \Delta_{i,j}/2$. Similar bounds are true for x_k . Then

$$\begin{aligned} d_{i,k} > 6d_{i,j} - 3d_{i,i^*} &\implies d_{i,k} - \frac{d_{i,k} - d_{i,i^*}}{2} d_{i,k} - \frac{\Delta_{i,k}}{2} \\ &> 2 \left(d_{i,j} + \frac{d_{i,j} - d_{i,i^*}}{2} \right) = 2 \left(d_{i,j} + \frac{\Delta_{i,j}}{2} \right), \end{aligned}$$

which implies that $L_{i,k} = \mathbb{E}[L_{i,k}] > 2\mathbb{E}[U_{i,j}] = 2U_{i,j}$ and $x_k \notin \mathcal{A}_j(0)$. \square

If all the random quantities take their expected value, then using Lemma C.1 and the elimination criterion of ANNEasy (Lemma 4.2), the complement of the initial active set $\mathcal{A}_j(0)$ (also called the elimination set) can be characterized in the following manner.

$$\begin{aligned} \mathcal{A}_j(0)^c &= \cup_{i < j: j \in \mathcal{A}_i(0)} \{ \mathcal{A}_i(0) \cap \{k : 2U_{i,j} < L_{i,k}\} \} \\ &\supseteq \cup_{i < j: j \in \mathcal{A}_i(0)} \{ \mathcal{A}_i(0) \cap \{k : d_{i,k} > 6d_{i,j} - 3d_{i,i^*}\} \}. \end{aligned} \quad (30)$$

Replacing the indicator $\mathbb{1}_{[A_{j,k}]}$ in Theorem 4.4 with an indicator for the non-membership of point x_k in the set (30) gives us an upper bound to the sample complexity of ANNEasy that is valid when all random quantities take their expected values.

To gain an idea of the savings achieved by our algorithm in comparison to the random sampling, we evaluate the sample complexity expressions for an example dataset. The dataset we look at consists of c clusters, each cluster containing $n/c > 1$ points. The points are indexed such that the m th cluster is $\mathcal{C}_m := \{x_{\underline{m}}, x_{1+\underline{m}}, \dots, x_{\overline{m}}\}$, where

$$\underline{m} := 1 + (m-1)n/c \quad (31)$$

and

$$\overline{m} := mn/c \quad (32)$$

for all $m \in [c]$. Suppose the distances between the points are such that for any pair $\{x_i, x_j\} \subseteq \mathcal{C}_m$, the set of points

$$\{x_k : d_{i,k} < 6d_{i,j} - 3d_{i,i^*}\} \subseteq \mathcal{C}_m. \quad (33)$$

The above condition is ensured if the smallest distance between two points belonging to different clusters is at least six times the diameter of any cluster.

Lemma C.2. *Consider a dataset which satisfies the condition in (33). If all random quantities take their expected values, ANNEasy uses $O(\sqrt{n})$ fewer oracle queries than the random sampling baseline to learn the nearest neighbor graph.*

Proof. In the following we assume that all random quantities take their expected values. We can find the points that are definitely eliminated using the triangle inequality when ANNEasy is called using (30). The elimination set $\mathcal{A}_1(0)^{\mathbb{G}} = \{x_1\}$. For a point $x_i \in \mathcal{C}_1 \notin \mathcal{E}_1$, from (30), (33) we get that

$$\mathcal{A}_i(0)^{\mathbb{G}} \supseteq \{\mathcal{A}_1(0) \cap \{k : d_{1,k} > 6d_{1,i} - 3d_{1,1^*}\}\} \supseteq \{(\mathcal{X} \setminus \{x_1\}) \cap \mathcal{C}_1^{\mathbb{G}}\} = \mathcal{C}_1^{\mathbb{G}}.$$

Thus $\mathcal{A}_i(0) \subseteq \mathcal{C}_1$ for all $x_i \in \mathcal{C}_1$. Point $x_{\underline{m}}$ is the first point processed by ANNEasy in the m th cluster. Suppose there exists a point $x_j \in \mathcal{C}_m \cap \mathcal{A}_{\underline{m}}(0)^{\mathbb{G}}$, we show next that leads to a contradiction. Since $x_j \notin \mathcal{A}_{\underline{m}}(0)$, there is a point $x_i \in \mathcal{C}_{m'}$ with $i < j, m' < m$ such that $2U_{i,\underline{m}} < L_{i,j}$. Let $\text{Diam}(\mathcal{C}_m) := \max_{x_\ell, x_k \in \mathcal{C}_m} d_{\ell,k}$ be the diameter of cluster \mathcal{C}_m (similarly for $\mathcal{C}_{m'}$) and let $D(\mathcal{C}_{m'}, \mathcal{C}_m) := \min_{x_\ell \in \mathcal{C}_{m'}, x_k \in \mathcal{C}_m} d_{\ell,k}$ be the minimum inter-cluster distance. Since the random quantities take their expected values, we have that

$$U_{i,\underline{m}} \geq d_{i,\underline{m}} + \frac{d_{i,\underline{m}} - d_{i,i^*}}{2} \implies 2U_{i,\underline{m}} \geq 3D(\mathcal{C}_{m'}, \mathcal{C}_m) - \text{Diam}(\mathcal{C}_{m'}),$$

$$L_{i,j} \leq d_{i,j} - \frac{d_{i,j} - d_{i,i^*}}{2} \implies L_{i,j} \leq \frac{\text{Diam}(\mathcal{C}_{m'}) + D(\mathcal{C}_{m'}, \mathcal{C}_m) + \text{Diam}(\mathcal{C}_m)}{2} + \frac{\text{Diam}(\mathcal{C}_{m'})}{2}.$$

Using $2U_{i,\underline{m}} < L_{i,j}$ with the above inequalities implies that $2.5D(\mathcal{C}_{m'}, \mathcal{C}_m) < 2D(\mathcal{C}_{m'}) + 0.5D(\mathcal{C}_m)$, which is a contradiction as from (33) we have that $D(\mathcal{C}_{m'}, \mathcal{C}_m) \geq \max\{3D(\mathcal{C}_{m'}), 3D(\mathcal{C}_m)\}$. Thus we have that $\mathcal{C}_m \cap \mathcal{A}_{\underline{m}}(0)^{\mathbb{G}} = \emptyset$. For any $x_j \in \mathcal{C}_m, j \neq \underline{m}$ we have that $x_j \in \mathcal{A}_{\underline{m}}(0)$ and hence from (30),

$$\mathcal{A}_j(0)^{\mathbb{G}} \supseteq \{\mathcal{A}_{\underline{m}}(0) \cap \{k : d_{\underline{m},k} > 6d_{\underline{m},j} - 3d_{\underline{m},\underline{m}^*}\}\} \supseteq \mathcal{C}_m^{\mathbb{G}}.$$

Based on the above discussion, we have a lower bound on the number of points present in the elimination set $\mathcal{A}_j(0)^{\mathbb{G}}$ for any $x_j \in \mathcal{C}_m$. By choosing the following values for the indicator in (25)

$$\mathbb{1}_{[A_{j,k}]} = \begin{cases} 0 & \text{if } x_j \in \mathcal{C}_m \setminus \{x_{\underline{m}}\} \text{ and } x_k \notin \mathcal{C}_m, \\ 1 & \text{otherwise,} \end{cases}$$

we get the following upper bound to the number of oracle queries, where $x_{\overline{m}}$ is the last point in \mathcal{C}_m .

$$\mathcal{O}\left(\sum_{m=1}^c \left(\sum_{k > \underline{m}} H_{\underline{m},k} + \sum_{k < \underline{m}} H_{\underline{m},k} - \sum_{\ell=1}^{m-1} H_{\underline{m},\ell} + \sum_{\ell=1}^{m-1} (H_{\underline{m},\ell} - H_{\ell,\underline{m}}) + \sum_{p > \underline{m}} \sum_{q > p}^{\overline{m}} \max\{H_{p,q}, H_{q,p}\}\right)\right) \quad (34)$$

where \underline{m} and \overline{m} are defined in (31) and are functions of m . The number of terms in the sum above is $\mathcal{O}(cn + (n/c)^2)$. A uniform sampling baseline approach would have $\mathcal{O}(n^2)$ terms in its sample complexity. Letting $c = \sqrt{n}$ gives our result. \square

The above lemma ensures that we have $\mathcal{O}(\sqrt{n})$ fewer terms in the sample complexity expression for ANNEasy compared to random sampling if the dataset satisfies (33). We can get a more precise characterization of the savings in query complexity in terms of the $\Delta_{p,q}$ values. For instance, using a single-parameter model for the distribution of $\Delta_{p,q}$ as done in [19], we can directly use their Corollary 1 in our context.

Lemma C.3. *Consider a clustered dataset $\mathcal{X} = \cup_{m=1}^c \mathcal{C}_m$ whose points satisfy (33). Each cluster contains an even number $2\nu := n/c$ of points. For any $m \in [c]$ and $x_j \in \mathcal{C}_m$, suppose the suboptimality gaps $\Delta_{j,k}$ for all $x_k \in \mathcal{C}_m$ take one of the following values, parametrized by an $\alpha > 0$:*

$$1 - \left(\frac{s-1}{\nu-1} \right)^\alpha, \quad \text{where } s \in \{1, 2, \dots, \nu-1\}. \quad (35)$$

Note that there are $\nu-1$ values given in (35) while there are $2\nu-2$ points in the cluster, excluding x_j and x_{j^} . Each value in (35) is the suboptimality gap for two distinct points in \mathcal{C}_m . Ignoring log-factors, if $\alpha = 1$ ANNTri finds all nearest neighbors with probability $1 - \delta$ in $\mathcal{O}(n(\nu^2 + n))$ calls to the oracle, while uniform sampling requires $\mathcal{O}(n^2\nu^2)$ calls for the same guarantee.*

Proof. By putting the clusters far from each other, one can see that there exist $\mathcal{X} = \cup_{m=1}^c \mathcal{C}_m$ whose points satisfy (33). Lemma B.4 shows by explicit construction that the condition on the suboptimality gaps within each cluster as stated in (35) can also be satisfied. Note that (35) is the same parametrization as equation 3 in [19].

Consider the points in the m th cluster, i.e., points $x_{\underline{m}}$ through $x_{\overline{m}}$. The elimination set $\mathcal{A}_{\underline{m}}(0)^\mathbb{G}$ can be the singleton $\{x_{\underline{m}}\}$, but by Lemma C.1 for all $x_p \in \mathcal{C}_m \setminus \{x_{\underline{m}}\}$, $\mathcal{A}_p(0)^\mathbb{G} \supseteq \mathcal{C}_m^\mathbb{G}$. Finding x_{p^*} is a best arm identification problem among points within the cluster \mathcal{C}_m . The last term in (34) counts the total number of oracle queries made by ANNEasy to identify the nearest neighbors of all $x_p \in \mathcal{C}_m \setminus \{x_{\underline{m}}\}$. Thus the number of oracle queries made by ANNEasy for identifying x_{p^*} is at most $\sum_{q \neq p} H_{p,q}$, while uniform sampling will make $nH_{p,p'}$ queries, where $p' := \arg \min_{q \neq p^*} \Delta_{p,q}$.

Ignoring log-factors, the sample complexity for finding x_{p^*} for an $x_p \in \mathcal{C}_m$ by ANNEasy is

$$\tilde{\mathcal{O}} \left(\sum_{i=1}^{2\nu-1} \Delta_{p,p_i}^{-2} + \sum_{x_\ell \notin \mathcal{C}_m} \Delta_{p,\ell}^{-2} \right) = \tilde{\mathcal{O}} \left(\sum_{i=1}^{2\nu-1} \Delta_{p,p_i}^{-2} + n - 2\nu \right).$$

Corollary 1 of [19] lists the value of that sum for different choices of α , for e.g., if $\alpha = 1$ then the sample complexity is $\tilde{\mathcal{O}}(\nu^2 + n - 2\nu)$. On the other hand, for finding x_{p^*} uniform sampling would make $\tilde{\mathcal{O}}(n\Delta_{p,p'}^{-2})$, i.e., $\tilde{\mathcal{O}}(n(\nu-1)^2)$ queries. By construction of the dataset, finding the nearest neighbor of each point in \mathcal{X} is equally hard. Thus ANNTri would make $\tilde{\mathcal{O}}(n(\nu^2 + n - 2\nu))$ queries while uniform sampling would take $\tilde{\mathcal{O}}(n^2\nu^2)$ queries. \square

Note that our problem setting is inherently different from the noiseless setting where all x_{i^*} 's can trivially be learned in $\binom{n}{2}$ samples. Due to the presence of noise in our queries, many distances must be repeatedly queried so $\binom{n}{2}$ samples is insufficient.